

The Geometry and Physics of Knots

SIR MICHAEL ATIYAH

When a mathematician addresses a general scientific audience, even as enlightened an audience as attends a Royal Institution Discourse, he faces a daunting task. Mathematics can be such a highly technical and abstract subject that communicating its latest developments to a lay public presents formidable difficulties. Bearing this in mind I selected this evening's topic according to the following criteria. It should have:

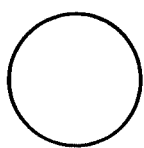
- (1) a simple visual content;
- (2) an interesting historical background;
- (3) a relation to physics;
- (4) a recent exciting story.

I picked on the subject of knots because it seemed to satisfy all these criteria, and on 11th July 1988 I sent my title off to the Royal Institution. Two weeks later, on 26th July, there was a spectacular further break-through which will be my main focus and makes the story even more exciting and topical than I had originally anticipated.

Knots, Links and Braids

Let me begin by introducing the dramatis personae of the evening: *knots* and their close relatives, *links* and *braids*. For a mathematician a knot is a closed piece of string, with no loose ends, exemplified in Figure 1.

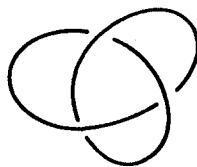
These pictures represent the piece of string laid flat on a table, so that at each crossing point one piece is 'under' and the other piece 'over' as indicated. Two knots are considered the same if we can manipulate one so that it looks like the other. The precise length (and thickness) of the string is irrelevant. Only its



Unknot



Unknot



Trefoil

Figure 1

'knottedness' is important. For example, the first two pictures (the circle and the figure-eight) represent the same knot – a rather uninteresting one referred to (for obvious reasons) as the 'unknot'. On the other hand, the trefoil knot cannot be disentangled and is not the same as the unknot. Of course in manipulating knots we are not allowed to cut and rejoin the string.

In every day parlance, a knotted piece of string usually has two free ends and we can with skill, untie the knot by threading the free ends through the knot. However, for a long piece of string this is a lengthy process and we might prefer (or be forced by other factors) to keep the free ends fixed and work directly on disentangling the knot. By this stage we are essentially back to the problem of knots in closed strings, which is why mathematicians have focussed on this point of view.

A link is like a knot except that it is made up of several closed pieces of string. Each piece may itself be knotted but, in addition, the different pieces may be 'linked' as illustrated by the following simple cases (Figure 2).



Figure 2

One reason why it is necessary to consider links as well as knots is that it is not easy, when presented with a diagrammatic picture (or by a real string tangle) to decide whether it consists of one or more pieces.

Finally a *braid* is a collection of strings, or 'strands', (with free ends) which may be entangled but which all 'move' in the same direction. Examples of braids with two or three strands are illustrated in Figure 3.

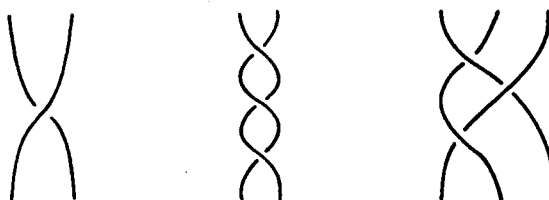


Figure 3

Note that each strand is always moving upwards. As with knots and links two braids are considered the same if we can manipulate one into the other while keeping the ends fixed and always maintaining the upwardness of each braid. The exact positions of the end points of the braid are also considered irrelevant.



Figure 4

Any braid can be converted into a link in a standard manner by simply connecting up the initial and final points as indicated in Figure 4.

A little manipulation shows that this particular example represents a trefoil knot. More generally every knot or link arises this way from a suitable braid.

As these few examples illustrate, the problem of deciding whether two plane diagrams represent the same knot or link is a difficult one. It corresponds essentially to the practical difficulty in trying to disentangle a complicated piece of string. By disregarding questions concerning the length and thickness of the string the problem is not so much one of *geometry* as of *topology*. In fact the study of knots is the archetype of a topological problem.

History of Knot Theory

Knots have attracted attention since the earliest times, as in the classical story of Alexander the Great and the Gordian knot. It was not however until the nineteenth century that it began to be considered scientifically. The notion of linkage is of fundamental importance in connection with electromagnetic induction, a fact which was fully appreciated by Maxwell. It is appropriate to recall that, here at the Royal Institution, Faraday demonstrated that an electric current along a wire produces an external magnetic field whose lines of forces link around the wire (Figure 5).

Such ideas may have been, in part, behind the ambitious theory of *Vortex Atoms* put forward by Lord Kelvin around 1867 [2]. At this period the ultimate nature of matter was a great mystery (it still is!) and Kelvin had the magnificent idea that

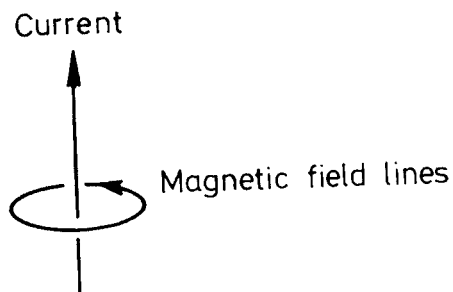
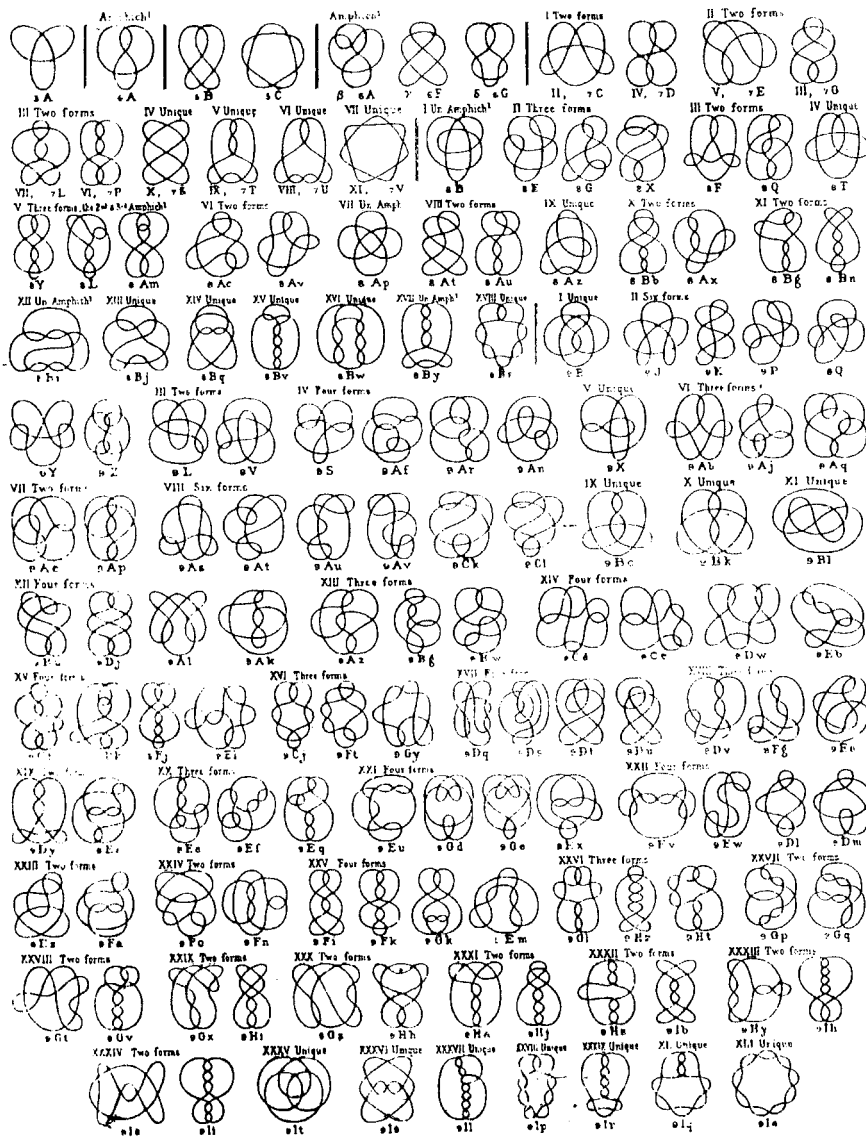


Figure 5

THE FIRST SEVEN ORDERS OF KNOTTINESS.



* See first note, p. 324 below, 1898.

To face p. 334

Figure 6

atoms might consist of knotted vortex tubes of the ether. His arguments in favour of this possibility can be summarised as follows:

- (1) *Stability*. The stability of matter could be explained by the topological stability of knots, *i.e.* under continuous deformation knots remain essentially the same.
- (2) *Variety*. The large number (actually infinite) of different knots could account for the different chemical elements.
- (3) *Vibrations*. The vortex tubes could presumably vibrate and this might explain spectral lines.

As a twentieth century footnote we might add a further argument:

- (4) *Transmutation*. At very high energies atoms can change into other atoms just as knots can, if we allow some cutting and recombination, change into other knots.

Kelvin's theory was, for a decade or so, taken very seriously and Maxwell's verdict was that "it satisfies more of the conditions than any atom hitherto imagined". If Kelvin's theory was on the right lines then a classification of knots was clearly going to be an essential ingredient and P.G. Tait, one of Kelvin's collaborators, spent more than 10 years studying and tabulating knots. He enumerated knots by the number of crossings of a plane diagram and produced tables of the distinct knots arising. This turned out to be a monumental task. Tait studied knots with up to 11 crossings (a sample page of his tables is copied in Figure 6). For 10 crossings there are 165 different knots while more recent computer tabulations for 13 crossings produce over 10,000 different knots. Perhaps it is fortunate for chemists that Kelvin's theory was eventually discarded!

In the course of his investigations Tait made a number of empirical discoveries which have subsequently been christened as Tait's conjectures. These conjectures appear highly plausible but resisted all attempts at proof by mathematicians for a whole century. Very recently, as a result of the exciting new ideas I am reporting on, many of Tait's conjectures have now been established. I will explain the simplest of Tait's conjectures. For this I need two notions. First an *alternating* knot diagram is one where, in following the path of knot, we meet crossings which are alternately 'over' and 'under'. A diagram which is not alternating tends to be one that can be simplified (see Figure 7), so that it might seem reasonable to concentrate on knots given by alternating diagrams (non-alternating knots, *i.e.* knots which cannot be represented by an alternating diagram, do exist but require many crossings and are hard to draw). Next, consider a schematic knot diagram as in Figure 8 in which a single crossing point separates the knot into two separate parts, schematically indicated by boxes. Clearly such a diagram can be simplified by a simple half-twist which removes the central crossing, to give Figure 9.



Figure 7

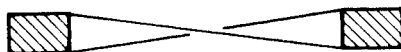


Figure 8

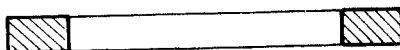


Figure 9

If all superfluous crossings are eliminated in this way, we end up with a *reduced* knot diagram. The Tait conjecture asserts that, for a *reduced alternating* knot diagram, the number of crossings is an *invariant* of the knot, *i.e.* it is independent of the particular diagram representing the knot (as long as this is reduced and alternating). Since a given knot can be represented by many different reduced alternating diagrams the fact that the number of crossings is always the same is by no means obvious.

Invariants of Knots

After Kelvin's Theory of Vortex Atoms was discarded the study of knots ceased to be a part of physics and was relegated to the world of mathematics. In the twentieth century Topology grew into a major discipline and many powerful techniques were developed. In particular some of these techniques could be applied to the classification of knots, and a major success was the discovery in 1928 of the *Alexander polynomial* of a knot. Named after the American mathematician J.W. Alexander (and not the Macedonian general!) this is an *invariant* of a knot, *i.e.* it is something which can be calculated algebraically from any diagram of the knot but is independent of the diagram chosen. An invariant of this type can be very useful in distinguishing different knots. All one has to do is to calculate the invariant for two knots and see if the answers are different. In this case, the knots have to be different. In the contrary case, when the invariants give the same answer, no conclusion can be drawn and more work must be done (*i.e.* more subtle invariants have to be investigated).

Despite the great advances made in Topology in the latter part of the twentieth century, progress in knot theory since the discovery of the Alexander polynomial has been slow. It was therefore a great surprise to all the experts when, in 1984, the New Zealand mathematician Vaughan Jones discovered another polynomial invariant of knots [1]. This is now called the Jones polynomial, and it has many remarkable properties. Superficially similar to the Alexander polynomial it is, in fact, fundamentally different. In particular, the Jones polynomial can (sometimes) distinguish a knot from its mirror image: for example, it distinguishes right-handed

and left-handed trefoil knots. The Alexander polynomial on the other hand can never distinguish mirror images.

It may be helpful at this stage to illustrate all this with a few formulae. For any knot (or link) let us denote the Alexander polynomial as $A(t)$ and the Jones polynomial as $V(t)$, where t is a formal variable. For the standard unknotted circle, $A(t) = V(t) = 1$ while for the trefoil knot

$$A(t) = t - 1 + t^{-1}$$

$$V(t) = -t^4 + t^3 + t$$

For the other trefoil knot (the mirror image) t gets replaced by its inverse t^{-1} . Clearly $A(t)$ is unchanged, but $V(t)$ is altered. Note that both $A(t)$ and $V(t)$ are polynomials in t and t^{-1} with integer coefficients. This is always the case for any knot. These coefficients are essentially numerical invariants of the knot. Arranging them in a polynomial is a convenient algebraic device for organising them.

The Jones invariants have proved to be very powerful tools in the study of knots. They have led to the proof of many of Tait's century old conjectures, including the one about reduced alternating knots I described above. They have also been generalised so that we now have a whole (infinite) collection of Jones-type polynomial invariants of knots. On the other hand, these invariants do not fit easily into the standard framework of conventional topological theory. They have more to do with ideas and techniques from various branches of theoretical physics, including statistical mechanics and quantum field theory.

This is a very brief description of the situation as it stood in the summer of 1988. There was an enormous amount of activity connecting physics with knot theory, but the essential reasons behind it all remained mysterious. Then, on 26th July 1988 at the International Conference on Mathematical Physics in Swansea, Edward Witten, from Princeton, suddenly hit on the key idea [3]. We now have a much better and more fundamental understanding of the relation between quantum physics and knot theory. The purpose of my Discourse is to explain this new insight without resorting to mathematical technicalities.

Force-free Regions

In classical electromagnetism a charged particle is acted on by a force field. In a region of space where there are no such forces a particle travels freely. However, quantum mechanics alters this naive picture as is clearly demonstrated by the famous Bohm-Aharonov effect. This concerns a beam of electrons travelling in a region external to a solenoid. The solenoid produces a magnetic field inside but the external region is shielded and has no magnetic field. Classically nothing should happen to the electrons but experimentally a diffraction pattern is produced. Quantum mechanically this is interpreted by saying that the wave-function of an electron undergoes a phase-shift on going round the solenoid.

Thus, although there are no forces in the region external to the solenoid, there are definite physical effects. Moreover the phase-shift is proportional to the strength of the solenoid.

This phenomenon is difficult to understand so it may be helpful to consider an analogous geometrical situation. In fact, according to Einstein's theory of General Relativity, gravitational force can be interpreted as curvature of space-time. A region where space-time is flat has therefore no gravitational force. Our example will be one where, outside some central core (analogue of the solenoid) space is flat but straight lines exhibit a kind of phase-shift. For simplicity, the example will be two-dimensional and concerned with the geometry of surfaces.

Let me begin by recalling that in plane Euclidean geometry the sum of the interior angles of any triangle is 180° . On a sphere, for spherical triangles made out of great circle arcs, the angle sum exceeds 180° . Moreover this excess depends inversely on the radius of the sphere: the smaller the radius the more curved the sphere and the bigger the excess. Now consider instead the surface of a cone. We can slit the cone along a line OA through the vertex and then open it up so as to get a plane region as indicated in Figure 10.

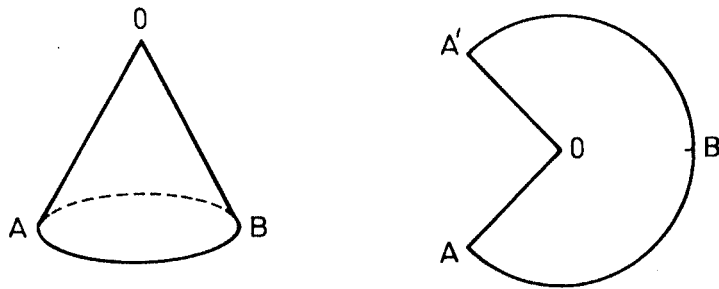


Figure 10

Consider now triangles drawn on the surface of the cone. We can examine these on the opened up cone. From this (see Figure 11) it is clear that: (i) a triangle not containing the vertex in its interior has angle sum 180° ; (ii) a triangle containing the vertex in its interior has angle sum exceeding 180° (by an amount α).

If the angle α is small the cone is very nearly a flat plane and the angle sum is nearly 180° . On the other hand for large values of α the cone is sharply pointed and the angle sum exceeds 180° by a large amount.

The essence of this example is that a cone, away from the vertex is really flat and that its local geometry is Euclidean. However, the global geometry of a cone, represented by triangles containing the vertex in their interior, is definitely not Euclidean. The vertex is a singularity which produces these non-Euclidean effects.

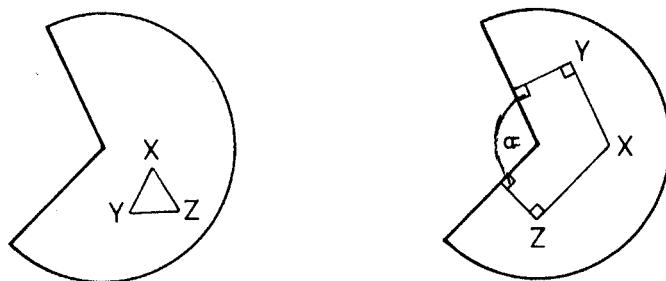


Figure 11

Unlike a normal 'source' the vertex produces no field of force (local curvature). Its effect is purely global or topological: it enables us to distinguish between triangles which contain the vertex and those which do not. Moreover the vertex has a 'strength', the angle α , which determines the scale of its effect.

To compare this with the Bohm-Aharonov effect we should reduce the dimension of the latter by taking a plane slice orthogonal to the solenoid. The solenoid (if ideally thin) would then produce the analogue of the cone's vertex. The phase-shift is the analogue of the angle excess α .

Since the quantum wave-function of an electron is a difficult notion to grasp we can try to understand the Bohm-Aharonov effect in simpler terms, analogous to our geometric example, by thinking of an electron as a particle with some additional or internal structure. For example, if we restrict ourselves to a fixed plane slice orthogonal to the solenoid, we could picture an electron as having a (very small) irregular shape (e.g. a triangle). Rotating this shape would then constitute phase-shifts, and we could interpret the effect of the solenoid as producing a definite rotation of the electron shape as it went round the solenoid.

This simplified picture can be generalised to allow for a number of parallel solenoids. The external force-free region would then, in a plane slice, affect our small electron shapes by producing an appropriate rotation depending on how many of the solenoids the electron went round. The rotations would simply add up.

All of this is really preparatory to a major generalisation of a new kind which I shall describe in the next section and which will bring us closer to the Jones polynomials.

Higher-dimensional Phases

Phase in the quantum theory of the electron is an angle. However for important generalisations, which are now standard in modern elementary particle theory, phase can become a rotation in a three-dimensional (or higher-dimensional)

'internal' space. Now rotations in three-dimensions are determined by an *axis* of rotation and an angle. The possible variability of the axis produces important new features. In particular, the composite of two rotations with different axes is a complicated operation and is not just given by adding angles as in the planar case. Moreover, the order in which the two rotations are performed affects the answer.

Suppose we now consider a solenoid in such a more general theory. Our test particle (playing the role of the electron) in the external region will then undergo some definite three-dimensional 'internal' rotation, described by an axis and an angle. Next let us consider a number of such solenoids. For simplicity take them all parallel and having the same strength or angle. However, their axes are unrestricted (the axis is related to 'internal' space and is not tied to the direction of the solenoid in real space).

Taking a plane slice across these solenoids, the external region in which our test particle moves has no forces but is characterised by the axes associated to the various solenoids. Each choice of axes gives a physically different force-free region or 'vacuum' outside the solenoids.

So far we have been describing an essentially static situation. Now we will allow things to change with time, by letting the solenoids move around. For simplicity we stick to our plane slice and idealise its intersection with each solenoid to a point. Our motion is then represented by some motion of these points in the plane. We will insist: (i) the points always remain distinct; (ii) at the end of the motion the set of points returns to its original position (possibly permuting the individual points).

Since we have discarded the third spatial dimension (along the solenoids) it is convenient to re-use the third dimension as time, and to represent the motion of our set of points in the plane by its space-time graph. Such a graph gives precisely a *braid* as described above, the number of strands being the number of solenoids.

Since the external region is force-free we can think of our moving points as drifting around independently. This makes it plausible that *only the topology of the motion, essentially the braid, is significant.*

We can now ask how our vacuum behaves under this motion. How does it propagate? More explicitly, given an initial choice of axes for the different points (representing the solenoids) what is the final set of axes arising from a given motion or braid?

It is at this stage that we have to remember that we are dealing with a quantum theory. This means that we can only expect probabilistic answers to our questions. Thus given an initial set of axes A and a final set B we can ask for the probability that, starting with A we end up with B.

For example, suppose we have four points (representing four solenoids). Taking the X, Y, Z axes for each point gives a total number of $4 \times 3 = 12$ choices for A. Similarly there are twelve choices for B. Hence our probabilities are described by a

12×12 matrix P_{AB} . Each braid on four strands is then supposed to produce such a 12×12 matrix P .

The matrix P describes how the 'quantum vacuum' is affected by the motion defined by the relevant braid. In particular the single number

$$\text{Trace } P = \sum_A P_{AA}$$

is an important numerical invariant. When divided by the number of choices for A it gives the probability that the initial and final vacuum state coincide.

So far I have attempted to show how braids appear naturally in modern physical theories and that appropriate numerical invariants have a physical meaning as quantum probabilities. In the next section I will take the last step by passing from braids to links and knots. This depends on the use of ideas from relativity theory.

Relativistic Invariance

I have already introduced knots, links and braids and I indicated that, by closing up a braid, we obtain a link and in fact all links arise in this way. However, many different braids can give rise to the same link. In particular the number of braid strands can vary. Thus to construct a link invariant we can start from some braid invariant, but this will have to be rather special, satisfying various constraints, if it is to give a link invariant.

In the previous section I indicated how quantum theory ideas can lead to braid invariants, the probabilities of 'vacuum evolution': more precisely the traces of the probability matrices give numerical invariants. It is natural to ask what further constraints are required for these to give braid invariants. In physical terms the answer is given by *relativistic invariance*.

Let me recall, at this stage, that the fundamental principle of relativity theory is that physics should be described in terms of four dimensional space/time, with time on an essentially equal footing with the three spatial dimensions. For a given physical theory to be relativistically invariant is a strong constraint, and is not always evident when the theory is described by a dynamic description involving a time evolution. It depends very much on the detailed laws or evolution equations of the theory.

There is a general approach due to Feynman which describes a physical theory directly in a space/time version and makes manifest relativistic invariance. This approach is based on the Feynman path integral which essentially assigns probabilities to all possible motions.

The theory which I was attempting to describe in the previous section was in dynamic or evolution form. However, Witten [3] has given a Feynman integral description of this theory and hence demonstrated its relativistic invariance.

Because we ignored one of the three spatial dimensions (along the solenoids) all our dimensions are one less than usual. Thus space/time is now just three dimensional and our braid was just a graph in this reduced space/time. Relativistic invariance means that, in this three-dimensional picture, time has no special significance. But this essentially brings us back to links, because the distinguishing feature of a braid is that all its strands move upwards (with time).

Witten's version of the theory therefore leads at once to link invariants and Witten has shown that these invariants are just the Jones invariants described earlier. To be precise the quantities (quantum probabilities) which Witten's theory assigns to a link are values of the corresponding Jones polynomial $V(t)$ at special values of the variable t . However, there are infinitely many such values (depending on the 'strengths' of the solenoids) and so this information determines the complete polynomial $V(t)$.

Conclusion

Let me try to sum up the story on the physical interpretation of the Jones invariants of knots and links. We start by considering a number of points in the plane each carrying a generalised magnetic flux. This flux has a strength or 'angle' which is fixed but it has also a variable axis. A choice of axis at each point determines the external vacuum – a force-free region in which physical effects are nevertheless present.

Allowing this situation to evolve in time leads to a braid, as a space-time graph of the motion of the points. The external vacuum evolves and this evolution is described by quantum probabilities. These probabilities yield invariants of the braid.

Finally, a relativistically invariant version of this story yields invariants of links. Thus the Jones invariants acquire a physical meaning as expectation values in a relativistically invariant quantum field theory. Moreover this theory is based on the generalisations of electromagnetism in which phase angles are replaced by rotations in higher-dimensional 'internal' space.

Even a superficial reading of the preceding paragraphs brings out the remarkable fact that this interpretation of the Jones invariant of links involves *all* the most sophisticated aspects of modern theoretical physics: quantum theory, relativity and 'generalised' electromagnetism. By contrast the Alexander polynomial can be simply interpreted using only classical electromagnetism. The fact that the Jones polynomial can distinguish mirror images is directly related to the physical theories involved being themselves sensitive to orientation, and this in turn is related to the lack of 'parity conservation' in fundamental physical processes.

It is striking that the development of knot theory has thus come full circle. It began as a putative branch of physics, with Kelvin's theory of vortex atoms.

Thrown back at the mathematicians knot theory developed slowly over the next hundred years, until it exploded into activity with the discovery of the Jones polynomial. Now we understand the Jones polynomial in terms of present day physics, vastly more sophisticated than the physics of Kelvin's time.

The relevance of the Jones polynomial for real physics, as opposed to the formal model physics I have been discussing, is still not clear. However, in more general terms, topological ideas are now playing a fundamental role in all current physical theories. The deep reason for this is certainly connected with the stability ideas that motivated Kelvin. The difference is that physicists today do not regard atoms as the ultimate entities: they are digging deeper. The ideas that Kelvin was trying to apply at the atomic level are now, in a new guise, being applied at an even more fundamental level. In a sense Kelvin was on the right track and a long way ahead of his time. He would undoubtedly be both amused and impressed at the way, a century later, knots have re-entered physics.

References

- 1 Jones, V.F.R. 1987. Hecke algebra representations of braid groups and link polynomials. *Ann. Math.*, **126**, 335.
- 2 Thompson, W.H. 1869. On vortex atoms. *Trans. R. Soc. Edin.*, **34**, 15–24.
- 3 Witten, E. 1988. Some geometrical applications of quantum field theory. *Proc. Int. Congress of Mathematical Physics*, Swansea.