EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 9.16

Quantum Phenomena

emerges from the same underlying network—or in effect from the structure of space. And indeed even in traditional general relativity one can try avoiding introducing matter explicitly—for example by imagining that everything we call matter is actually made up of pure gravitational energy, or of something like gravitational waves.

But so far as one can tell, the details of this do not work out—so that at the level of general relativity there is no choice but to introduce matter explicitly. Yet I suspect that this is in effect just a sign of limitations in the Einstein equations and general relativity.

For while at a large scale these may provide a reasonable description of average behavior in a network, it is almost inevitable that closer to the scale of individual connections they will have to be modified. Yet presumably one can still use the Einstein equations on large scales if one introduces matter with appropriate properties as a way to represent small-scale effects in the network.

In the previous section I suggested that energy and momentum might in effect be associated with the presence of excess nodes in a network. And this now potentially seems to fit quite well with what we have seen in this section. For if the underlying rule for a network is going to maintain to a certain approximation the same average number of nodes as flat space, then it follows that wherever there are more nodes corresponding to energy and momentum, this must be balanced by something reducing the number of nodes. But such a reduction is exactly what is needed to correspond to positive curvature of the kind implied by the Einstein equations in the presence of ordinary matter.

Quantum Phenomena

From our everyday experience with objects that we can see and touch we develop a certain intuition about how things work. But nearly a century ago it became clear that when it comes to things like electrons some of this intuition is no longer correct. Yet there has developed an elaborate mathematical formalism in quantum theory that successfully reproduces much of what is observed. And while some aspects of this formalism remain mysterious, it has increasingly come to be believed that any fundamental theory of physics must somehow be based on it.

Yet the kinds of programs I have discussed in this book are not in any obvious way set up to fit in with this formalism. But as we have seen a great many times in the course of the book, what emerges from a program can be very different from what is obvious in its underlying rules. And in fact it is my strong suspicion that the kinds of programs that I have discussed in the past few sections will actually in the end turn out to show many if not all the key features of quantum theory.

To see this, however, will not be easy. For the kinds of constructs that are emphasized in the standard formalism of quantum theory are very different from those immediately visible in the programs I have discussed. And ultimately the only reliable way to make contact will probably be to set up rather complete and realistic models of experiments—then gradually to see how limits and idealizations of these manage to match what is expected from the standard formalism. Yet from what we have seen in this chapter and earlier in this book there are already some encouraging signs that one can identify.

At first, though, things might not seem promising. For my model of particles such as electrons being persistent structures in a network might initially seem to imply that such particles are somehow definite objects just like ones familiar from everyday experience. But there are all sorts of phenomena in quantum theory that seem to indicate that electrons do not in fact behave like ordinary objects that have definite properties independent of us making observations of them.

So how can this be consistent? The basic answer is just that a network which represents our whole universe must also include us as observers. And this means that there is no way that we can look at the network from the outside and see the electron as a definite object. Instead, anything we deduce about the electron must come from processes that explicitly go on inside the network.

But this is not just an issue in studying things like electrons: it is actually a completely general feature of the models I have discussed. And in fact, as we saw earlier in this chapter, it is what allows them to support meaningful notions of even such basic concepts as time. At a more formal level, it also implies that everything we can observe can be captured by a causal network. And as I will discuss a little below, I suspect that the idea of causal invariance for such a network will then be what turns out to account for some key features of quantum theory.

The basic picture of our universe that I have outlined in the past few sections is a network whose connections are continually updated according to some simple set of underlying rules. In the past one might have assumed that a system like this would be far too simple to correspond to our universe. But from the discoveries in this book we now know that even when the underlying rules for a system are simple, its overall behavior can still be immensely complex.

And at the lowest level what I expect is that even though the rules being applied are perfectly definite, the overall pattern of connections that will exist in the network corresponding to our universe will continually be rearranged in ways complicated enough to seem effectively random.

Yet on a slightly larger scale such randomness will then lead to a certain average uniformity. And it is then essentially this that I believe is responsible for maintaining something like ordinary space—with gradual variations giving rise to the phenomenon of gravity.

But superimposed on this effectively random background will then presumably also be some definite structures that persist through many updatings of the network. And it is these, I believe, that are what correspond to particles like electrons.

As I discussed in the last two sections, causal invariance of the underlying rules implies that such structures should be able to move at a range of uniform speeds through the background. Typically properties like charge will be associated with some specific pattern of connections at the core of the structure corresponding to a particle, while the energy and momentum of the particle will be associated with roughly the number of nodes in some outer region around the core.

So what about interactions? If the structures corresponding to different particles are isolated, then the underlying rules will make them persist. But if they somehow overlap, these same rules will usually make some different configuration of particles be produced.

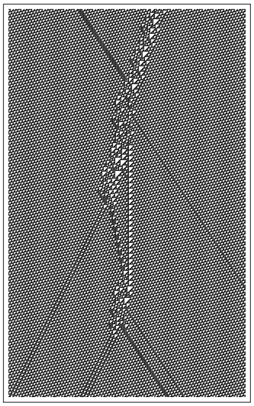

A collision between localized structures in the rule 110 class 4 cellular automaton.

At some level the situation will no doubt be a little like in the evolution of a typical class 4 cellular automaton, as illustrated on the left. Given some initial set of persistent structures, these can interact to produce some intermediate pattern of behavior, which then eventually resolves into a final set of structures that again persist.

In the intermediate pattern of behavior one may also be able to identify some definite structures. Ones that do not last long can be very different from ones that would persist forever. But ones that last longer will tend to have properties progressively closer to genuinely persistent structures. And while persistent structures can be thought of as corresponding to real particles, intermediate structures are in many ways like the virtual particles of traditional particle physics.

So this means that a picture like the one on the left above can be viewed in a remarkably literal sense as being a spacetime diagram of particle interactions—a bit like a Feynman diagram from particle physics.

One immediate difference, however, is that in traditional particle physics one does not imagine a pattern of behavior as definite and determined as in the picture above. And indeed in my model for the universe it is already clear that there is more going on. For any process like the one in the picture above must occur on top of a background of apparently random small-scale rearrangements of the network. And in effect what this background does is to introduce a kind of random environment that can make many different detailed patterns of behavior occur with certain probabilities even with the same initial configuration of particles.

The idea that even a vacuum without particles will have a complicated and in some ways random form also exists in standard quantum field theory in traditional physics. The full mathematical structure of quantum field theory is far from completely worked out. But the basic notion is that for each possible type of particle there is some kind of continuous field that exists throughout space—with the presence of a particle corresponding to a simple type of structure in this field.

In general, the equations of quantum field theory seem to imply that there can be all sorts of complicated configurations in the field, even in the absence of actual particles. But as a first approximation, one can consider just short-lived pairs of virtual particles and antiparticles. And in fact one can often do something similar for networks. For even in the planar networks discussed on page 527 a great many different arrangements of connections can be viewed as being formed from different configurations of nearby pairs of non-planar persistent structures.

Talking about a random background affecting processes in the universe immediately tends to suggest certain definite relations between probabilities for different processes. Thus for example, if there are two different ways that some process can occur, it suggests that the total probability for the whole process should be just the sum of the probabilities for the process to occur in the two different ways.

But the standard formalism of quantum theory says that this is not correct, and that in fact one has to look at so-called probability amplitudes, not ordinary probabilities. At a mathematical level, such amplitudes are analogous to ones for things like waves, and are in effect just numbers with directions. And what quantum theory says is that the probability for a whole process can be obtained by linearly combining the amplitudes for the different ways the process can occur, then looking at the square of the magnitude of the result—or the analog of intensity for something like a wave.

So how might this kind of mathematical procedure emerge from the types of models I have discussed? The answer seems complicated. For even though the procedure itself may sound straightforward, the constructs on which it operates are actually far from easy to define just on the basis of an underlying network—and I have seen no easy way to unravel the various limits and idealizations that have to be made.

Nevertheless, a potentially important point is that it is in some ways misleading to think of particles in a network as just interacting according to some definite rule, and being perturbed by what is in essence a random background. For this suggests that there is in effect a unique history to every particle interaction—determined by the initial conditions and the configuration that exists in the random background.

But the true picture is more complicated. For the sequence of updates to the underlying network can be made in any order—yet each order in effect gives a different detailed history for the network. But if there is causal invariance, then ultimately all these different histories must in a sense be equivalent. And with this constraint, if one breaks some process into parts, there will typically be no simple way to describe how the effect of these parts combines together.

And for at least some purposes it may well make sense to think explicitly about different possible histories, combining something like amplitudes that one assigns to each of them. Yet quite how this might work will certainly depend on what feature of the network one tries to look at.

It has always been a major issue in quantum theory just how one tells what is happening with a particular particle like an electron. From our experience with everyday objects we might think that it should somehow be possible to do this without affecting the electron. But if the only things we have are particles, then to find out something about a given particle we inevitably have to have some other particle—say a photon of light—explicitly interact with it. And in this interaction the original particle will inevitably be affected in some way.

And in fact just one interaction will certainly not be enough. For we as humans cannot normally perceive individual particles. And indeed there usually have to be a huge number of particles doing more or less the same thing before we successfully register it.

Most often the way this is made to happen is by setting up some kind of detector that is initially in a state that is sufficiently unstable that just a single particle can initiate a whole cascade of consequences. And usually such a detector is arranged so that it evolves to one or another stable state that has sufficiently uniform properties that we can recognize it as corresponding to a definite outcome of a measurement.

At first, however, such evolution to an organized state might seem inconsistent with microscopic reversibility. But in fact—just as in so many other seemingly irreversible processes—all that is needed to preserve reversibility is that if one looks at sufficient details of the system there can be arbitrary and seemingly random behavior. And the point is just that in making conclusions about the result of a measurement we choose to ignore such details.

So even though the actual result that we take away from a measurement may be quite simple, many particles—and many events—

will always be involved in getting it. And in fact in traditional quantum theory no measurement can ultimately end up giving a definite result unless in effect an infinite number of particles are involved.

As I mentioned above, ordinary quantum processes can appear to follow different histories depending on what scheme is used to decide the order in which underlying rules are applied. But taking the idealized limit of a measurement in which an infinite number of particles are involved will probably in effect establish a single history.

And this implies that if one knew all of the underlying details of the network that makes up our universe, it should always be possible to work out the result of any measurement. I strongly believe that the initial conditions for the universe were quite simple. But like many of the processes we have seen in this book, the evolution of the universe no doubt intrinsically generates apparent randomness.

And the result is that most aspects of the network that represents the current state of our universe will seem essentially random. So this means that to know its form we would in essence have to sample every one of its details—which is certainly not possible if we have to use measurements that each involve a huge number of particles.

One might however imagine that as a first approximation one could take account of underlying apparent randomness just by saying that there are certain probabilities for particles to behave in particular ways. But one of the most often quoted results about foundations of quantum theory is that in practice there can be correlations observed between particles that seem impossible to account for in at least the most obvious kind of such a so-called hidden-variables theory.

For in particular, if one takes two particles that have come from a single source, then the result of a measurement on one of them is found in a sense to depend too much on what measurement gets done on the other—even if there is not enough time for information travelling at the speed of light to get from one to the other. And indeed this fact has often been taken to imply that quantum phenomena can ultimately never be the result of any definite underlying process of evolution.

But this conclusion depends greatly on traditional assumptions about the nature of space and of particles. And it turns out that for the kinds of models I have discussed here it in general no longer holds.

And the basic reason for this is that if the universe is a network then it can in a sense easily contain threads that continue to connect particles even when the particles get far apart in terms of ordinary space.

The picture that emerges is then of a background containing a very large number of connections that maintain an approximation to three-dimensional space, together with a few threads that in effect go outside of that space to make direct connections between particles.

If two particles get created together, it is reasonable to expect that the tangles that represent their cores will tend to have a few connections in common—and indeed this for example happens for lumps of non-planarity of the kind we discussed on page 527. But until there are interactions that change the structure of the cores, these common connections will then remain—and will continue to define a thread that goes directly from one particle to the other.

But there is immediately a slight subtlety here. For earlier in this chapter I discussed measuring distance on a network just by counting the minimum number of successive individual connections that one has to follow in order to get from one point to another. Yet if one uses this measure of distance then the distance between two particles will always tend to remain fixed as the number of connections in the thread.

But the point is that this measure of distance is in reality just a simple idealization of what is relevant in practice. For the only way we end up actually being able to measure physical distances is in effect by looking at the propagation of photons or other particles. Yet such particles always involve many nodes. And while they can get from one point to another through the large number of connections that define the background space, they cannot in a sense fit through a small number of connections in a thread. So this means that distance as we normally experience it is typically not affected by threads.

But it does not mean that threads can have no effect at all. And indeed what I suspect is that it is precisely the presence of threads that leads to the correlations that are seen in measurements on particles.

It so happens that the standard formalism of quantum theory provides a rather simple mathematical description of these correlations. And it is certainly far from obvious how this might emerge from detailed mechanisms associated with threads in a network. But the fact that this and other results seem simple in the standard formalism of quantum theory should not be taken to imply that they are in any sense particularly fundamental. And indeed my guess is that most of them will actually in the end turn out to depend on all sorts of limits and idealizations in quantum theory—and will emerge just as simple approximations to much more complex underlying behavior.

In its development since the early 1900s quantum theory has produced all sorts of elaborate results. And to try to derive them all from the kinds of models I have outlined here will certainly take an immense amount of work. But I consider it very encouraging that some of the most basic quantum phenomena seem to be connected to properties like causal invariance and the network structure of space that already arose in our discussion of quite different fundamental issues in physics.

And all of this supports my strong belief that in the end it will turn out that every detail of our universe does indeed follow rules that can be represented by a very simple program—and that everything we see will ultimately emerge just from running this program.