EXCERPTED FROM

STEPHEN WOLFRAM
A NEW KIND OF SCIENCE

SECTION 9.15

The Phenomenon of Gravity

situations—save those associated with the overall expansion of the universe—the basic rules for the network at least on average just rearrange nodes and never change their number.

In traditional physics energy and momentum are always assumed to have continuous values. But just as in the case of position there is no contradiction with sufficiently small underlying discrete elements.

As I will discuss in the last section of this chapter, quantum mechanics tends to make one think of particles with higher momenta as being somehow progressively less spread out in space. So how can this be consistent with the idea that higher momentum is associated with having more nodes? Part of the answer probably has to do with the fact that outside the piece of the network that corresponds to the particle, the network presumably matches up to yield uniform space in much the same way as without the particle. And within the piece of the network corresponding to the particle, the effective structure of space may be very different—with for example more long-range connections added to reduce the effective overall distance.

The Phenomenon of Gravity

At an opposite extreme from elementary particles one can ask how the universe behaves on the largest possible scales. And the most obvious effect on such scales is the phenomenon of gravity. So how then might this emerge from the kinds of models I have discussed here?

The standard theory of gravity for nearly a century has been general relativity—which is based on the idea of associating gravity with curvature in space, then specifying how this curvature relates to the energy and momentum of whatever matter is present.

Something like a magnetic field in general has different effects on objects made of different materials. But a key observation verified experimentally to considerable accuracy is that gravity has exactly the same effect on the motion of different objects, regardless of what those objects are made of. And it is this that allows one to think of gravity as a general feature of space—rather than for example as some type of force that acts specifically on different objects.

In the absence of any gravity or forces, our normal definition of space implies that when an object moves from one point to another, it always goes along a straight line, which corresponds to the shortest path. But when gravity is present, objects in general move on curved paths. Yet these paths can still be the shortest—or so-called geodesics—if one takes space to be curved. And indeed if space has appropriate curvature one can get all sorts of paths, as in the pictures below.

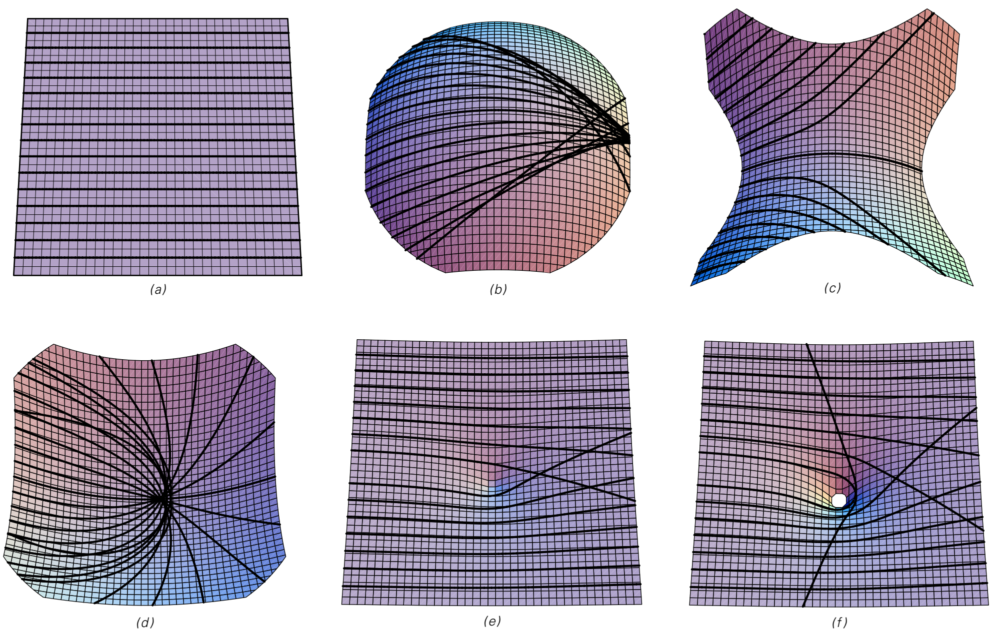

Examples of the effect of curvature in space on paths taken by objects. In each case all the paths shown start parallel, but do not remain so when there is curvature. The paths are geodesics which go the minimum distance on the surface to get to all the points they reach. (In general, the minimum may only be local.) Case (b) shows the top of a sphere, which is a surface of positive curvature. Case (c) shows the negatively curved surface $z = x^2 - y^2$, (d) a paraboloid $z = x^2 + y^2$, and (e,f) $z = 1/(r + \delta)$—a rough analog of curvature in space produced by a sphere of mass.

But in our actual universe what determines the curvature of space? The answer from general relativity is that the Einstein equations give conditions for the value of a particular kind of curvature in terms of the energy and momentum of matter that is present. And the point then is that the shortest paths in space with this curvature seem to be consistent with those followed by objects moving under the influence of gravity associated with the given distribution of matter.

For a continuous surface—or in general a continuous space—the idea of curvature is a familiar one in traditional geometry. But if the universe is at an underlying level just a discrete network of nodes then how does curvature work? At some level the answer is that on large scales the discrete network must approximate continuous space.

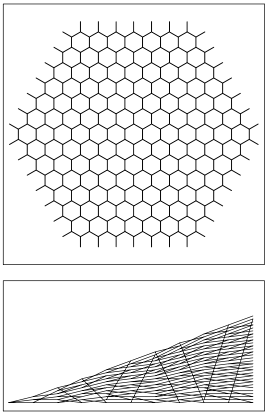

A hexagonal array corresponding to flat two-dimensional space.

But it turns out that one can actually also recognize curvature in the basic structure of a network. If one has a simple array of hexagons—as in the picture on the left—then this can readily be laid out flat on a two-dimensional plane. But what if one replaces some of these hexagons by pentagons? One still has a fundamentally two-dimensional surface. But if one tries to keep all edges the same length the surface will inevitably become curved—like a soccer ball or a geodesic dome.

So what this suggests is that in a network just changing the pattern of connections can in effect change the overall curvature. And indeed the pictures below show a succession of networks that in effect have curvatures with a range of negative and positive values.

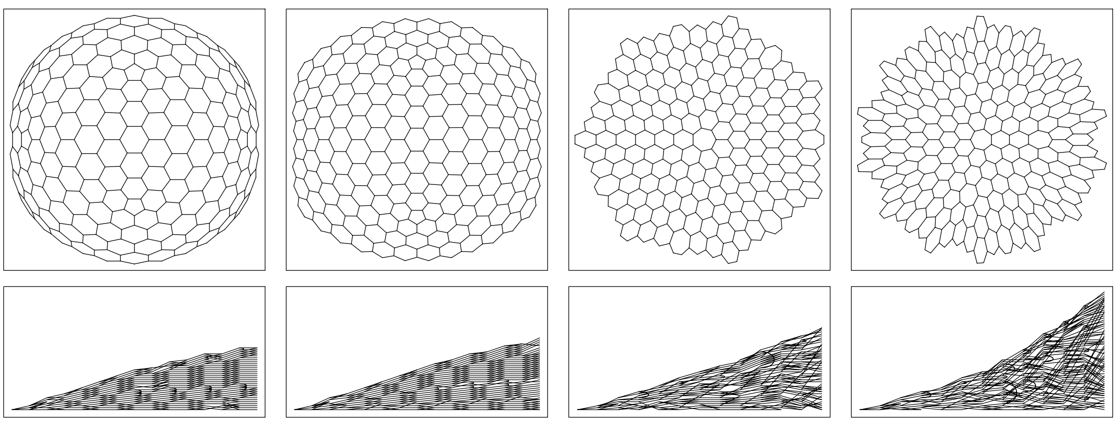

Networks with various limiting curvatures. If every region in the network is in effect a hexagon—as in the picture at the top of the page—then the network will behave as if it is flat. But if pentagons are introduced, as in the cases on the left, the network will increasingly behave as if it has positive curvature—like part of a sphere. And if heptagons are introduced, as in the cases on the right, the network will behave as if it has negative curvature. In the bottom row of pictures, the networks are laid out as on page 479, so that successive heights give the number of nodes at successive distances r from a particular node. In the limit of large r, this number is approximately $r^2 (1 - k r^2 + \ldots)$ where k turns out to be exactly proportional to the curvature.

But how can we determine the curvature from the structure of each network? Earlier in this chapter we saw that if a network is going to correspond to ordinary space in some number of dimensions d, then this means that by going r connections from any given node one must reach about r^{d-1} nodes. But it turns out that when curvature is present it leads to a systematic correction to this.

In each of the pictures on the facing page the network shown can be thought of as corresponding to two-dimensional space. And this means that to a first approximation the number of nodes reached must increase linearly with r. But the bottom row of pictures show that there are corrections to this. And what happens is that when there is positive curvature—as in the pictures on the left—progressively fewer than r nodes end up being reached. But when there is negative curvature—as on the right—progressively more nodes end up being reached. And in general the leading correction to the number of nodes reached turns out to be proportional to the curvature multiplied by r^{d+1}.

So what happens in more than two dimensions? In general the result could be very complicated, and could for example involve all sorts of different forms of curvature and other characteristics of space. But in fact the leading correction to the number of nodes reached is always quite simple: it is just proportional to what is called the Ricci scalar curvature, multiplied by r^{d+1}. And already here this is some suggestion of general relativity—for the Ricci scalar curvature also turns out to be a central quantity in the Einstein equations.

But in trying to see a more detailed correspondence there are immediately a variety of complications. Perhaps the most obvious is that the traditional mathematical formulation of general relativity seems to rely on many detailed properties of continuous space. And while one expects that sufficiently large networks should in some sense act on average like continuous space, it is far from clear at first how the kinds of properties of relevance to general relativity will emerge.

If one starts, say, from an ordinary continuous surface, then it is straightforward to approximate it as in the picture on the right by a collection of flat faces. And one might think that the edges of these faces would define a network of the kind I have been discussing.

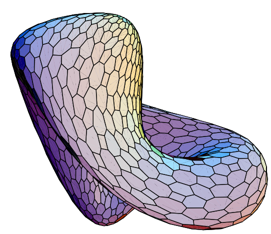

A surface approximated by flat faces whose edges form a trivalent network.

But in fact, such a network has vastly less information. For given just a set of connections between nodes, there is no obvious way even to know which of these connections should be associated with the same face—let alone to work out anything like angles between faces.

Yet despite this, it turns out that all the geometrical features that are ultimately of relevance to general relativity can actually be determined in large networks just from the connectivity of nodes.

One of these is the value of the so-called Ricci tensor, which in effect specifies how the Ricci scalar curvature is made up from different curvature components associated with different directions.

As indicated above, the scalar curvature associated with a network is directly related to how many nodes lie within successive distances r of a given node on the network—or in effect how many nodes lie within successive generalized spheres around that node. And it turns out that the projection of the Ricci tensor along a particular direction is then just related to the number of nodes that lie within a cylinder oriented in that direction. But even just defining a consistent direction in a network is not entirely straightforward. But one way to do it is simply to pick two points in the network, then to say that paths in the network are going in the same direction if they are segments of the same shortest path between those points. And with this definition, a region that approximates a cylinder can be formed just by setting up spheres with centers at every point on the path.

But there is now another issue to address: at least in its standard formulation general relativity is set up in terms of properties not of three-dimensional space but rather of four-dimensional spacetime. And this means that what is relevant are properties not so much of specific networks representing space, but rather of complete causal networks.

And one immediate feature of causal networks that differs from space networks is that their connections go only one way. But it turns out that this is exactly what one needs in order to set up the analog of a spacetime Ricci tensor. The idea is to start at a particular event in the causal network, then to form what is in effect a cone of events that can be reached from there. To define the spacetime Ricci tensor, one considers—as on page 516—a sequence of spacelike slices through this cone and asks how the number of events that lie within the cone increases as one goes to successive slices. After t steps, the number of events reached will be proportional to t^d. But there is then a correction proportional to t^{d+2}, that has a coefficient that is a combination of the spacetime Ricci scalar and a projection of the spacetime Ricci tensor along what is in effect the time direction defined by the sequence of spacelike slices chosen.

So how does this relate to general relativity? It turns out that when there is no matter present the Einstein equations simply state that the spacetime Ricci tensor—and thus all of its projections—are exactly zero. There can still for example be higher-order curvature, but there can be no curvature at the level described by the Ricci tensor.

So what this means is that any causal network whose behavior obeys the Einstein equations must at the level of counting nodes in a cone have the same uniform structure as it would if it were going to correspond to ordinary flat space. As we saw a few sections ago, many underlying replacement rules end up producing networks that are for example too extensively connected to correspond to ordinary space in any finite number of dimensions. But I suspect that if one has replacement rules that are causal invariant and that in effect successfully maintain a fixed number of dimensions they will almost inevitably lead to behavior that follows something close to the Einstein equations.

Probably the situation is somewhat analogous to what we saw with fluid behavior in cellular automata in Chapter 8—that at least if there are underlying rules whose behavior is complicated enough to generate significant effective randomness, then almost whenever the rules lead to conservation of total particle number and momentum something close to the ordinary Navier-Stokes equation behavior emerges.

So what about matter?

As a first step, one can ask what effect the structure of space has on something like a particle—assuming that one can ignore the effect of the particle back on space. In traditional general relativity it is always assumed that a particle which is not interacting with anything else will move along a shortest path—or so-called geodesic—in space.

But what about an explicit particle of the kind we discussed in the previous section that exists as a structure in a network? Given two nodes in a network, one can always identify a shortest path from one to the other that goes along a sequence of individual connections in the network. But in a sense a structure that corresponds to a particle will normally not fit through this path. For usually the structure will involve many nodes, and thus typically require many connections going in more or less the same direction in order to be able to move across the network.

But if one assumes a certain uniformity in networks—and in particular in the causal network—then it still follows that particles of the kind that we discussed in the previous section will tend to move along geodesics. And whereas in traditional general relativity the idea of motion along geodesics is essentially an assumption, this can now in principle be derived explicitly from an underlying network model.

One might have thought that in the absence of matter there would be little to say about gravity—since after all the Einstein equations then say that there can be no curvature in space, at least of the kind described by the Ricci tensor. But it turns out that there can still be other kinds of curvature—described for example by the so-called Riemann tensor—and these can in fact lead to all sorts of phenomena. Examples include familiar ones like inverse-square gravitational fields around massive objects, as well as unfamiliar ones like gravitational waves.

But while the mathematical structure of general relativity is complicated enough that it is often difficult to see just where in spacetime effects come from, it is usually assumed that matter is somehow ultimately required to provide a source for gravity. And in the full Einstein equations the Ricci tensor need not be zero; instead it is specified at every point in space as being equal to a certain combination of energy and momentum density for matter at that point. So this means that to know what will happen even in phenomena primarily associated with gravity one typically has to know all sorts of properties of matter.

But why exactly does matter have to be introduced explicitly at all? It has been the assumption of traditional physics that even though gravity can be represented in terms of properties of space, other elements of our universe cannot. But in my approach everything just 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