The Mathematics of Three-dimensional Manifolds

Topological study of these higher-dimensional analogues of a surface suggests the universe may be as convoluted as a tangled loop of string. It now appears most of the manifolds can be analyzed geometrically

by William P. Thurston and Jeffrey R. Weeks

Thousands of years ago many people thought the earth was flat. The flatness of the earth's surface must have seemed self-evident to anyone who looked out across an ocean or a prairie, and it was argued, not entirely unreasonably, that the surface of the earth must either be infinite or have an edge. It is now understood, of course, how such a basic misconception could have come about: even from a few thousand meters above the ground a small part of the roughly spherical earth looks like a small part of a plane. What is less often appreciated is that an unlimited number of terrestrial shapes would give rise to the same local observations. For example, it would be consistent with such local observations for the earth to have the shape of an irregular blob or of a doughnut.

Investigations in the branch of mathematics called topology make it clear that we confront an analogous situation when we attempt to describe the overall form of the universe based on the limited view from our point in space. An observer on the earth cannot conclude that the universe retains the geometric structure of ordinary Euclidean space to indefinite distances, although there is still no evidence to the contrary. If the structure of the universe is not Euclidean, what are the alternatives? One familiar idea is that space may be "curved" in much the same way as a surface can be curved. The three-dimensional curvature of space and a closely related concept, the four-dimensional curvature of space and time, have become important ideas in astronomy and cosmology because of the key role they play in Einstein's general theory of relativity.

Nevertheless, the determination of curvature alone is not enough to specify what can be called, loosely speaking, the shape of the universe. Certain kinds of possible three-dimensional structure for the universe can be specified by analogy with the two-dimensional surfaces, but the analogy only begins to suggest the richness of form that is introduced by a third dimension. Indeed, since space and time are treated in the theory of relativity as a single entity called space-time, one might suppose that the appropriate mathematical structure of the universe must be a four-dimensional one. There is good reason to believe, however, that the structure of four-dimensional space-time is governed by the structure of three-dimensional space alone. Hence in order to investigate the overall structure of the universe without prejudice one must begin to understand the kinds of three-dimensional structure that could give rise to the observed universe. The structures are called three-dimensional manifolds, or three-manifolds for short.

The study of three-manifolds is, in a sense, a generalization of the study of two-manifolds, or surfaces. Topologists have known how to describe and classify all possible two-manifolds for more than a century, but the systematic classification of all three-manifolds remains an unsolved problem due to the exceedingly complex forms to which some three-manifolds give rise. A mathematical procedure called surgery suggests a measure of the complexity. Surgery makes it possible to construct a three-manifold from any tangled loop of string, no matter how knotted or convoluted the tangle. Imagine confronting two snarled masses of fishing line and trying to determine whether or not they are tangled in exactly the same way. Unless it is possible to classify such tangles of line in a systematic way, there is no hope that three-manifolds can be analyzed either. Until recently, therefore, mathematicians saw little reason for thinking a systematic theory of three-manifolds could be devised.

That pessimistic assessment must now be reconsidered. Investigations by one of us (Thurston) into the geometry of three-manifolds show there is a pattern that may lead to an understanding of all possible three-manifolds. All known three-manifolds fit the pattern, and as a result their twisting and winding can be described in geometric terms.

The theory of manifolds arose in the 19th century out of a need to understand quantitative relations geometrically. For example, the set of solutions to an equation that has two variables can be plotted as a set of points in the plane. Each point represents a pair of values for the variables that make the equation true; typically, the set of points is a curve or a set of curves. Similarly, the set of solutions to an equation that has three variables can often be plotted as a two-dimensional surface in three-dimensional space, such as the surface of a sphere. For equations with more than three variables one can describe the set of solutions geometrically in much the same way: it is a higher-dimensional manifold in a still higher-dimensional space. Although one cannot visualize such objects directly, mathematicians have developed conceptual tools for the study of equations that lead to higher-dimensional manifolds.

Topology cannot actually solve equations. What it provides is a mathematical vocabulary—adjectives and nouns—that allow a set of solutions to be discussed in a general way without actually being specified. Thus, although the manifold of points that makes up the set of solutions to an equation has a precise and unambiguous shape, the topology of the manifold is not constrained by the properties of that shape. Instead the topology encompasses whatever properties are retained when the manifold is deformed in an arbitrary way, as long as the deformation is done without cutting, tearing or puncturing.

A doughnut can be deformed into a coffee cup by making a concave depression in the surface of the doughnut and then enlarging the depression while shrinking the rest of the doughnut. As the old joke goes, a topologist is a person who cannot tell a (one-hole) doughnut from a (one-handle) coffee cup. On the other hand, the topologist does distinguish the surface of a doughnut from the surface of a glass without handles, because there is no way one shape can be continuously deformed to yield the other. It may seem that by allowing arbitrary deformations topology discards most of the interesting features of a manifold. In many mathematical questions, however, topological information plays a significant role.

The first substantial contributions to the topological theory of three-manifolds were made around the turn of the century by Henri Poincaré, Max Dehn and Poul Heegaard. One difficulty with the study of three-manifolds is that direct visualization must partially give way to abstract representation. Many surfaces can be visualized because they

TOPOLOGICAL STRUCTURE of the universe need not conform to the structure of infinite, three-dimensional Euclidean space. The mathematical theory of three-dimensional manifolds, or three-manifolds, demonstrates that space may "curve back on itself" in an infinite variety of ways. One possible model for the topology of space is the three-manifold discovered by Herbert Seifert and C. Weber, who is now at the University of Geneva, in 1932. The manifold cannot be depicted from without because to do so one would have to view it from a fourth or higher dimension. Nevertheless, it can be visualized in a more limited sense as a dodecahedron whose opposite faces are mathematically glued together, or identified. The colored, ruled bars moving into and out of the faces of the dodecahedron indicate how the gluings are to be carried out: one member of each pair of faces is matched to its counterpart after a rotation of three-tenths of a turn about the axis perpendicular to the two faces. Although parts of the bars are shown as ghosted images outside the dodecahedron, the bars do not really exist there because it is assumed that only points inside the dodecahedron exist. When one of the bars moves toward one of the faces of the dodecahedron, it disappears at that face and reappears at the opposite face as if it were entering the dodecahedron from another direction. If the structure of the universe is that of the Seifert-Weber manifold, the universe is finite but will expand forever.

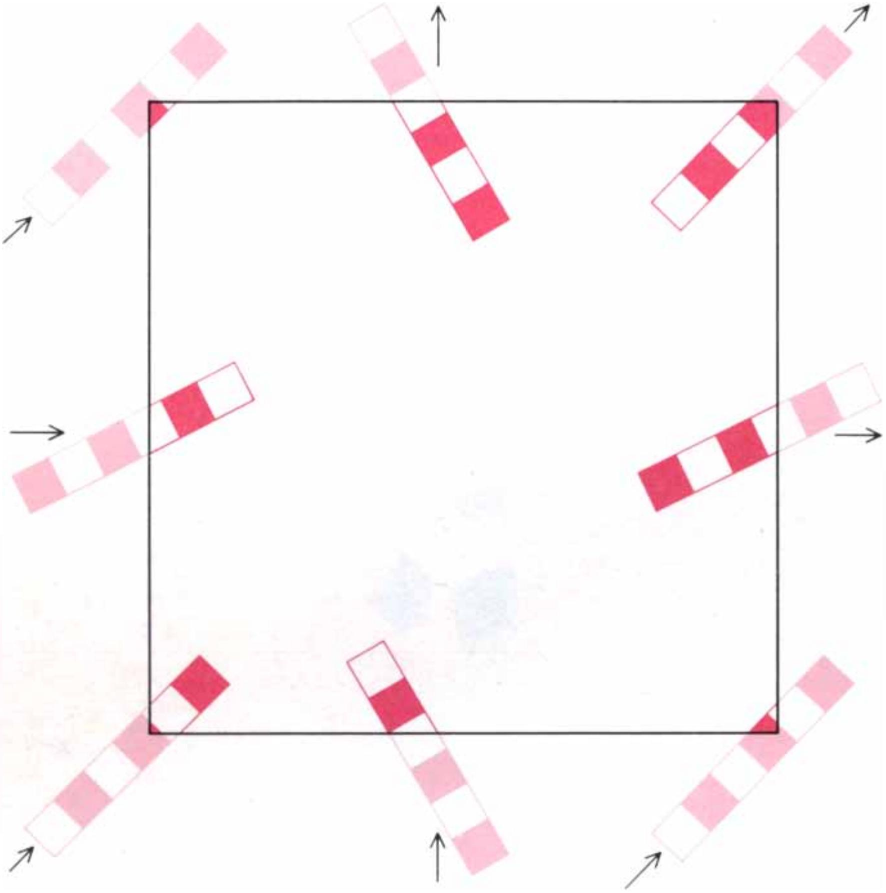

TWO-DIMENSIONAL MANIFOLD known as the two-torus can be represented as a square whose opposite edges are abstractly glued together. In other words, the top edge of the square is identified with the bottom edge and the left edge is identified with the right edge. If a ruled bar moves off the right edge, it reappears at the left edge; if the bar moves off the top edge, it reappears at the bottom. The motion is similar to that of objects in many video games. When the edges are abstractly glued together, all four vertexes of the square coincide in the manifold; when a ruled bar moves toward a vertex, pieces of it reappear at the other three vertexes.

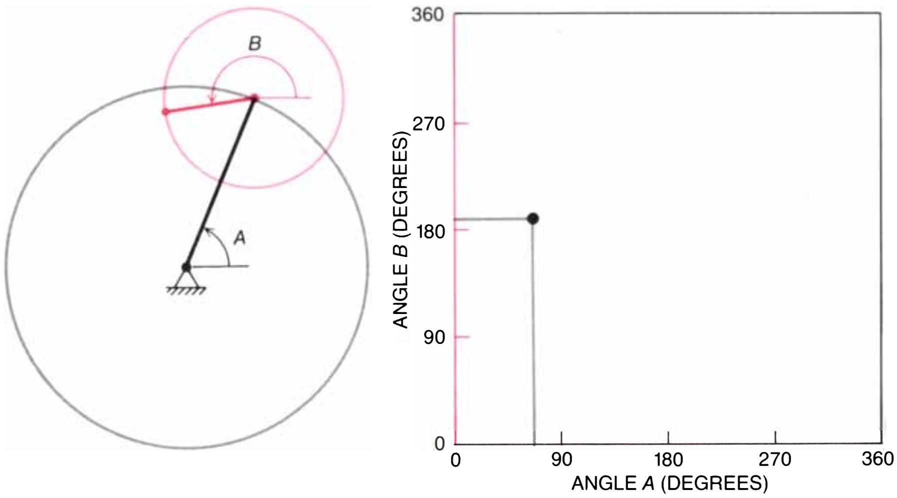

DOUBLE-CRANK LINKAGE SYSTEM is made up of two rigid bars pinned together at one point; the end of one bar is held fixed. The bars are free to rotate about the pins as long as the movement is confined to the plane of the page (*left*). Every possible configuration of the two bars can be given by a point plotted on two perpendicular coordinate axes: one axis gives the angle between the direction of the first bar and a fixed direction, and the second axis gives the corresponding angle for the second bar. The set of all the plotted points, which represent all possible positions of the linkage, is called the configuration space of the linkage (*right*). Because the configuration of the double crank does not change if the angle of either bar is changed by 360 degrees, the configuration space is a square bounded by the lines that represent zero-degree and 360-degree rotations for each bar. Points on opposite sides of the square represent identical configurations of the linkage; in other words, the configuration space is a two-torus.

can be seen externally from the third dimension, a dimension that is one higher than the dimension of the surface. The extra dimension gives the surface enough room to bend around and close up with itself. One might try to visualize a three-manifold externally, as if one were viewing it from a space with four or more dimensions, but it turns out such contortions are not necessary.

In the 19th century mathematicians found that two-manifolds can be represented as polygons whose edges are to be glued together, or in other words identified with each other in a specified way. In the novel *Flatland,* published in 1884, Edwin A. Abbott describes a two-dimensional creature living entirely within the plane. Consider the movements of such a creature on a two-manifold with a more exotic topology, namely a square whose opposite edges are identified. When the creature moves off the top edge of the square, it reappears at the bottom; when it moves off the right edge, it reappears at the left. Intrinsically, therefore, the top of the square is glued to the bottom and the right edge is glued to the left. It is worth noting that many video games operate on the same principle: when a figure moves off the top edge of the screen, it reappears at the bottom, and so on.

For a square it is a simple matter to carry out the gluings. Attaching the top of the square to the bottom gives rise to a cylinder open at both ends, and gluing the open ends leads to a one-hole doughnut. After the edges have been glued the seams are erased; the Flatlander cannot tell where the gluings were made. The doughnut and the square (with edges properly identified) count topologically as the same abstract manifold, namely the two-torus.

As the video games demonstrate, however, it is not necessary to do the gluing to get an intuitive understanding of the two-torus. With a little practice it is just as easy to follow the motion of an object on the square, where the gluings are specified only in an abstract sense. Abstract gluing brings within the compass of geometric intuition a great many manifolds that would otherwise be difficult to visualize. What is most important for us is that the gluing trick can readily be generalized to bring geometric intuition to bear on the understanding of three-manifolds.

Consider the three-manifold generated from a rectangular block of space, such as the space inside a room. Abstractly glue the front wall of the room to the back wall, the left wall to the right wall and the floor to the ceiling. If the gluings were actually done, one would have to imagine the room bending around and joining itself in a fourth dimension. All that is needed for the description of the manifold, however, is

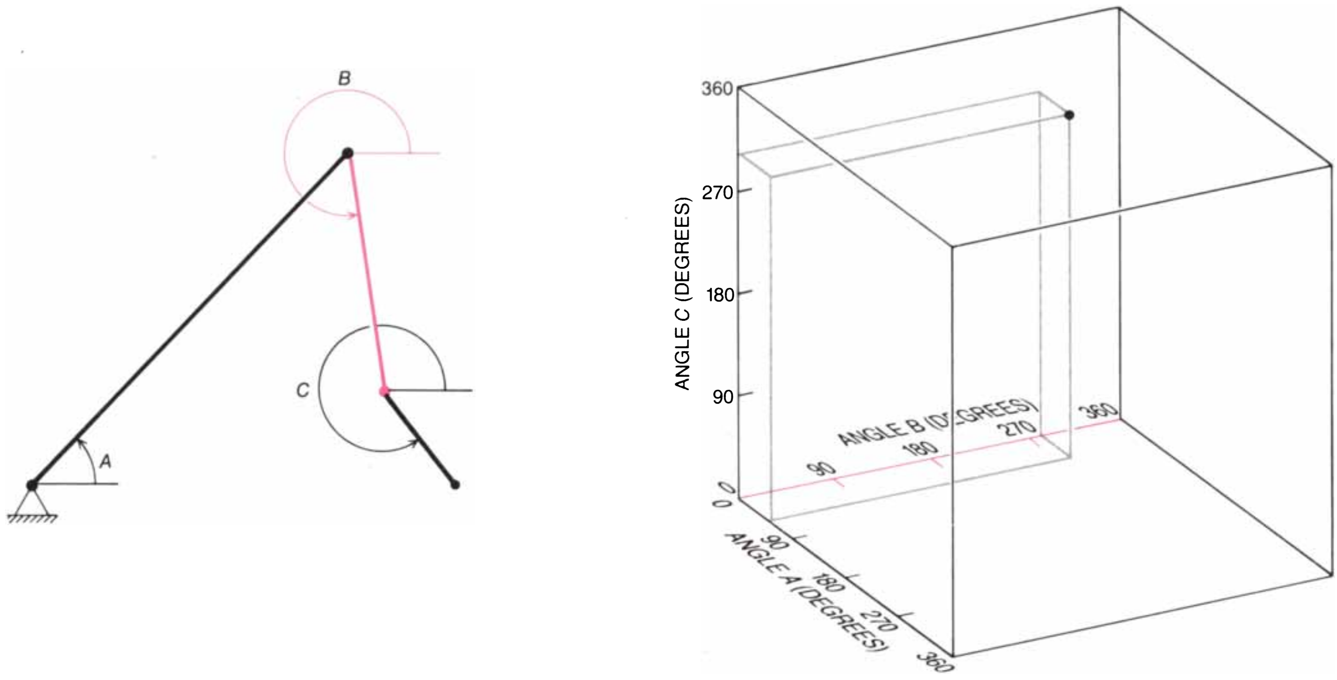

THREE-MANIFOLD analogous to the two-torus arises from the set of all possible configurations of a triple crank whose motion is confined to a plane (*left*). If the angles between each bar and a fixed direction are plotted on three mutually perpendicular axes, every possible position of the triple crank can be plotted as a point in a cube (*right*). The configuration of the linkage does not change with any complete rotation of a bar. Thus in the configuration space every face of the cube that corresponds to a rotation of 360 degrees is abstractly identified with its opposite face, which corresponds to a rotation of zero degrees. The resulting three-manifold is called the three-torus.

given by the procedure for abstract gluing. If an object within the manifold is moved toward the front wall, it disappears at that wall and reappears on the back wall; similarly, the object disappears at the right wall as it reappears on the left wall and disappears at the ceiling as it reappears on the floor. Evidently the motion is strikingly similar to the motion of an object within the two-torus; the manifold is the three-dimensional analogue of the two-torus, and so it is called the three-torus.

If ordinary concepts of space and physical reality are momentarily set aside, one can readily imagine living in a three-torus. Look at the back wall and the line of sight passes through that wall and returns from the opposite point on the front wall. What you see is a copy of yourself from behind. Look to the right and you see a copy of yourself from the left; look down at the floor and you see the top of your head. Indeed, because the line of sight continues crossing the room in all directions, you see what appear to be infinitely many copies of yourself and the room, all arranged in a rectangular lattice. The optical effect is similar to the one created by a room whose walls, floor and ceiling are covered with mirrors. The difference is that there are no reflections reversing the images of the room; instead all the images are direct copies of the original.

Does the fact that astronomers have not observed such peculiar visual effects imply the universe cannot be a three-torus? No. The universe is between 10 and 20 billion years old. If it were a three-torus, say, 60 billion light-years across, no light would have had enough time to complete a round trip. Another possibility is that observational astronomy has already recorded light that has traveled all the way around the universe: if the universe is a three-torus, one of the distant galaxies we observe may be our own. The possibility would be hard to verify because the image of our galaxy would be formed from light that left its source billions of years ago and spent the intervening time crossing the universe. What could be seen, given unlimited resolution of the image, would be the Milky Way in its earliest stages of evolution, as it looked when the light was emitted. Such a universe has a finite volume but no boundary of any kind.

Similar models of the possible spatial structure of the universe can be derived from other polyhedrons as well as from the cube. In each case the best way to understand the manifold is to imagine certain faces of the polyhedrons abstractly glued together. Two such manifolds are readily constructed from the regular dodecahedron. The 12 faces of the dodecahedron are regular pentagons arranged in pairs in such a way that the members of each pair are parallel and on diametrically opposite sides of the dodecahedron.

In the first dodecahedral three-manifold one member of each pair of pentagons is identified with the opposite pentagon by rotating the first member one-tenth of a turn counterclockwise about the axis perpendicular to its surface. The manifold is called the Poincaré manifold because it is equivalent to a three-manifold discovered by Poincaré in 1902. (Poincaré was unaware, however, that the manifold could be made from a dodecahedron.) The second dodecahedral manifold arises when each pentagon is glued to its opposite counterpart after a counterclockwise rotation of three-tenths of a turn. The resulting manifold is called the Seifert-Weber dodecahedral space, after Herbert Seifert and C. Weber, now at the University of Geneva, who discovered the manifold in 1932 [*see illustration on page 109*]. Like the three-torus, both manifolds would give rise to a universe with a finite volume but no boundary or edge.

Vast numbers of additional models for the large-scale structure of space can be constructed in a similar way. Because most polyhedrons are irregular, most three-manifolds arise from abstract gluing of the faces of irregular polyhedrons. The description of the gluing can become quite complicated when the number of faces is large.

It may appear there is a certain unreality to the exercise. The nonspecialist would probably grant that the true topology of space is a worthy object of speculation but would likely wonder how far such speculation should be carried. In particular, the nonspecialist might well question the utility of studying "spaces" in the plural. To the topologist such objections miss the mark because they focus only on the metaphorical content of topology. The study of topology can certainly be motivated by problems that arise in other contexts, but topology itself is a theory of pure form, not a theory of the real world. If the structure of space were somehow settled tomorrow, no topologist would give up the study of abstract spaces.

This assertion need not imply that topology is irrelevant to the real world. On the contrary, like other branches of mathematics, topology has many strong and substantive connections with the world, but the connections are indirect. If a particular metaphor, such as the spatial one, starts to wear thin, it is best to abandon the metaphor rather than to abandon the study of the form to which the metaphor gave rise. Experience has shown repeatedly that a mathematical theory with a rich internal structure generally turns out to have significant implications for the understanding of the real world, often in ways no one could have envisioned before the theory was developed. A theory would never reach the mature stage of development in which such applications are recognized if it were constantly burdened with overworked metaphors.

In order to illustrate the scope of topological analysis, it is useful to abandon the cosmological metaphor for the time being in favor of a more terrestrial one. Consider a mechanical system of bars and linkages, such as the one that connects a key in a manual typewriter to the type element. We shall discuss only planar linkages, or in other words assemblies of rigid bars pinned to one another in such a way that all the bars move in only one plane. There must also be at least one anchor point, or fixed base, to which the bars are pinned.

The aim of a theory of mechanical linkages is to analyze the possible motions of a linkage. There are many real mechanical devices to which the analysis applies, and they need not have much physical resemblance to a collection of bars joined together. The study of linkages was much in vogue in the second half of the 19th century, when there was interest in the problem of finding a linkage in which at least one point moves in a straight line. It seemed that a solution to the problem would lead to many practical applications, such as the design of a power train for a steam locomotive. Although a number of elegant theoretical solutions to the problem

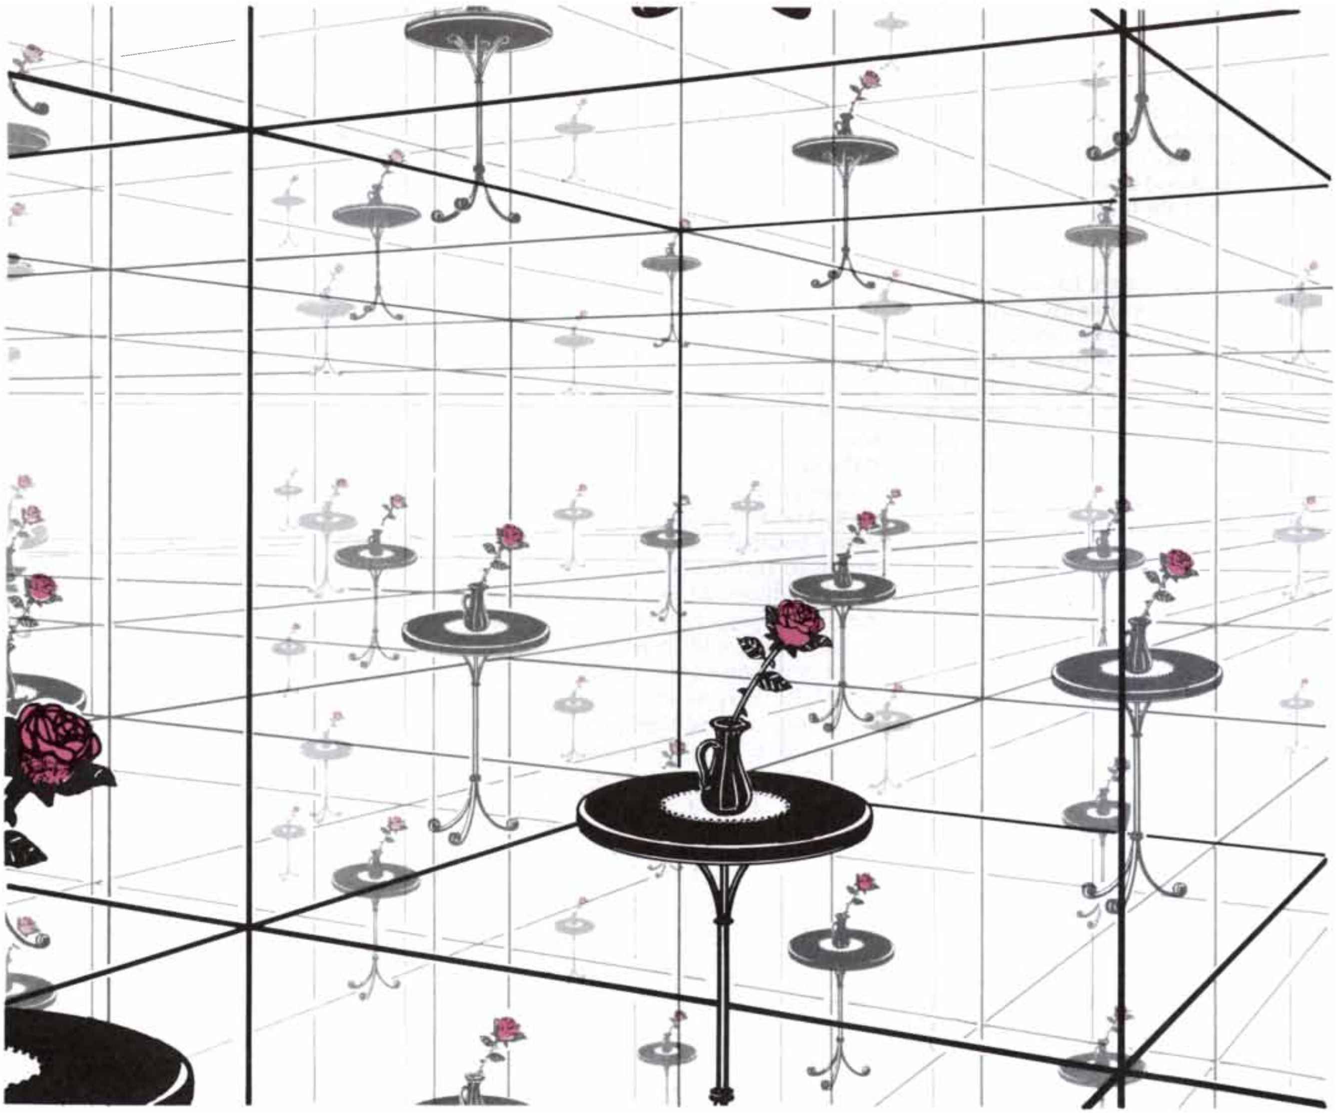

OBSERVER'S VIEW inside a three-torus is similar to the view inside a room whose walls, floor and ceiling are covered with mirrors; there is, however, no mirror reversal of the images. The line of sight passes into, say, the right wall and emerges from the left wall; looking right, therefore, the observer sees the room as it would appear from the left wall. Similarly, looking forward the observer sees the room as it would appear from the back wall, and looking up the observer sees the room as it would appear from the floor. Since the line of sight continues indefinitely across the three-torus, the room appears to be an infinite rectangular lattice extending in all directions. The three-torus is not infinite, however, because the images in the infinite rectangular array are really all images of the same thing.

were found, none of them proved to be a mechanically practical design.

A mechanical linkage can be represented mathematically by a set of line segments in the plane; at some intersections of the lines there may be pivot points, or pins, for the linkage. In the mathematical theory one assumes the lines and pivot points can pass freely through one another. The problem of constructing a physical model whose bars and pins replicate the motions of the idealized linkage is not a trivial one, but it is secondary to the mathematical analysis. It turns out that for any mathematical version of a linkage there is a physical linkage that executes the same motion, although the physical linkage may be much more complicated than its theoretical counterpart and may look quite different.

The set of all possible positions for a mechanical linkage is called the configuration space of the linkage; in most cases it is a topological manifold. Consider the simplest possible linkage, made up of a single bar pinned to an anchor point at one end but otherwise free to move in a plane. The moving end of the bar traces a circle in space, and each point on the circle corresponds to only one position of the linkage. The configuration space is a circle, which can also be viewed as a straight line segment whose ends are abstractly glued together. The circle is a one-dimensional manifold analogous to the two-torus, and every point in the manifold is identified with one position of the linkage.

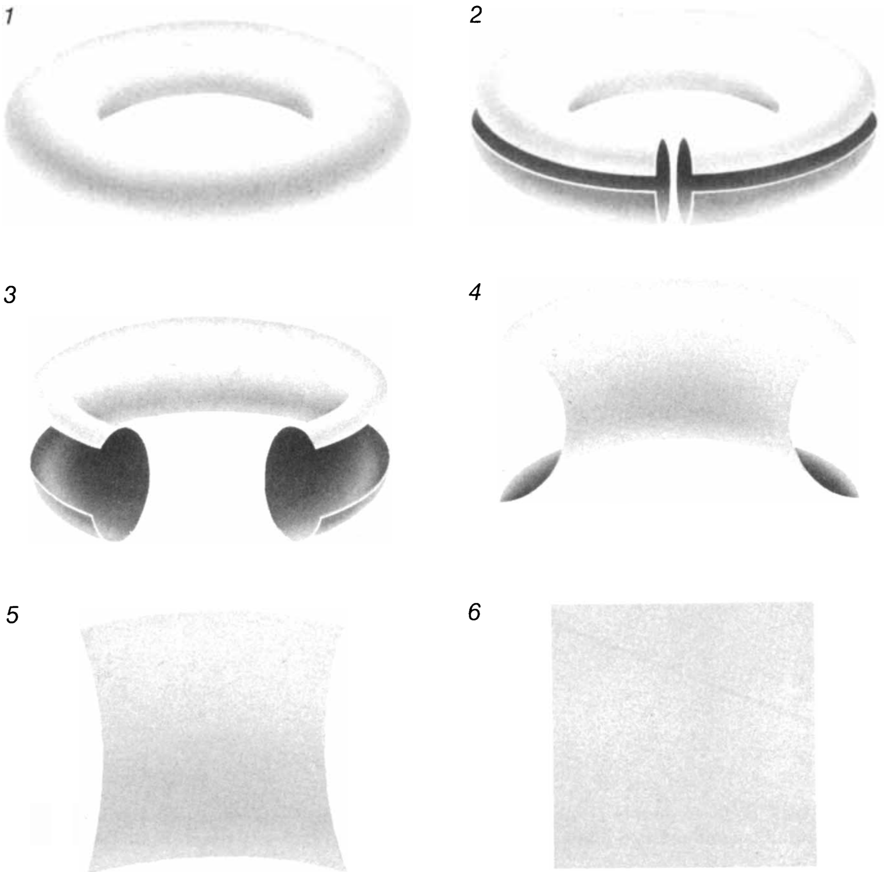

DOUGHNUT WITH ONE HOLE can be slit open and stretched into a square. If the opposite edges of the square are abstractly glued, the resulting surface is topologically equivalent to the doughnut. Since the square is flat like the plane, its geometry is Euclidean; hence from a topological point of view the one-hole doughnut is said to admit a Euclidean geometry.

By pinning another bar to the end of the first bar one obtains a double crank, a mechanical linkage with two degrees of freedom. If the second bar in the linkage is shorter than the first, the free end of the second bar can reach any point in a ring centered on the anchor point. The ring is bounded on the outside by a circle whose radius is the sum of the lengths of the two bars and bounded on the inside by a circle whose radius is the difference between the two lengths. If the second bar is the same length as the first, the free end of the second bar can reach any point in a circle whose radius is equal to the sum of the lengths of the two bars. If the second bar is longer than the first, the trace of the free end is also a ring whose inside radius is equal to the difference between the two lengths. One must not mistake these sets of points, however, for the configuration space of the linkage. The reason is that knowledge about the position of the end point of the second bar does not uniquely determine the configuration of the linkage. For every point reached by the free end of the second bar, the elbow of the double crank can bend in either of two ways.

In order to analyze the configuration space correctly it is easier to consider the possible configurations of the double crank without regard for the position of the free end. Every configuration can be described by two angles, namely the angle between each bar and some fixed direction (say to the right), measured in a counterclockwise sense. The two angles range freely and independently from zero to 360 degrees, but for each bar the angle zero degrees is identified with the angle 360 degrees. If the two angles are plotted on mutually perpendicular coordinate axes in the plane, every point within the square bounded by the lines labeled zero and 360 degrees for each angle corresponds to a different configuration of the double crank. Furthermore, every configuration of the double crank is represented by a point within the square. Since zero and 360 degrees are identified, the top edge of the square is identified with the bottom edge and the left edge is identified with the right edge. The configuration space is the two-torus.

If a third bar is added to the free end of the double crank, any position of the resulting triple crank can be described by giving the three angles of the bars. The angles are again measured counterclockwise from a fixed direction, and again they range freely and independently from zero to 360 degrees. The angles zero and 360 degrees are identified as before. When the three angles are plotted on three mutually perpendicular coordinate axes, every possible position of the triple crank is represented by a unique point in a cube whose opposite pairs of faces are abstractly glued together. The configuration space of the triple crank is therefore equivalent to the three-torus.

All the configuration spaces we have described so far lead to polygons or polyhedrons whose edges or faces can be glued together without distortion. There is no topological rule, however, that forbids abstract gluing when the edges or surfaces fail to be geometrically congruent; in fact, the examples we have given are quite special in that the gluings do not require distortion of the parts to be glued. Consider a mechanical linkage made up of three double cranks, each one fixed to a vertex of an equilateral triangle and pinned at its moving end to the moving ends of the other two double cranks [*see illustration on next page*]. In order to understand the configuration space of the linkage, first plot the set of possible positions for the central pin. Each double crank keeps the

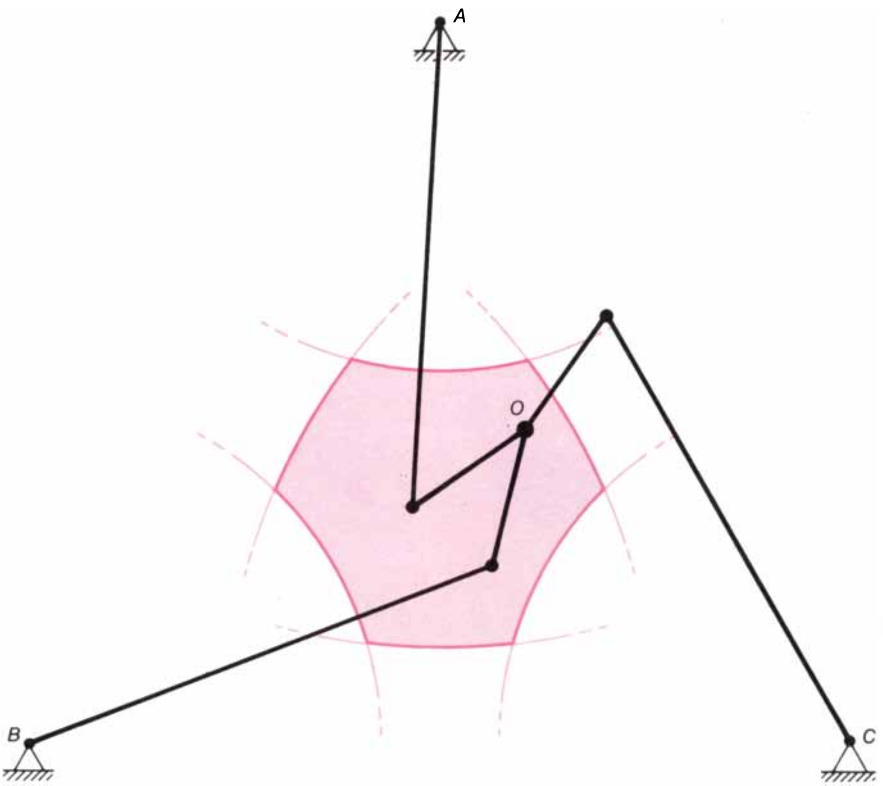

LINKAGE of three double cranks joined by a central pin also gives rise to a two-manifold. The motion is confined to a plane, and the central pin can reach any point inside the hexagon.

pin within a ring centered on the anchor point of that double crank. Hence the central pin can reach any point that lies in the intersection of the three rings corresponding to the three double cranks. The intersecting region is a curvilinear hexagon in the plane.

There is more to the configuration space, however, than one planar hexagon. Remember that for one double crank there are two configurations of the linkage for each point reached by the end of the second crank. Similarly, for each point in the interior of the planar hexagon, the elbow of each double crank can be bent in one of two ways. With three double cranks the total number of configurations for each point in the interior of the planar hexagon is 2^3, or eight configurations. The configuration space for the three double cranks can therefore be assembled by abstractly gluing the edges of eight abstract curvilinear hexagons.

How are the eight abstract hexagons to be glued together in the configuration space? When the central pin of the three double cranks lies on an edge of the hexagon in the plane, one of the double cranks is forced into straight alignment. The alignment can be straight in one of two ways: the bars of the double crank can point in the same direction or in opposite directions. In both cases, however, the configuration of the entire linkage is specified as soon as one specifies the position of the central pin and the sense of the bend of the two unstraightened double cranks.

For each of the two edges of the planar hexagon that can be traced when the first double crank is straight, there correspond only 2^2, or four, distinct edges in the configuration space instead of eight. Along any such edge in the configuration space two hexagons are glued together. The two hexagons represent the two ways in which the first double crank can bend. The opposite edge of the first hexagon in the configuration space is also glued to the opposite edge of the second hexagon, because the first double crank also becomes straight there. The bending of the other two double cranks does not change from the first to the second hexagon. The analysis is identical for all the other edges of the configuration space.

Similarly, when the central pin of the three double cranks lies on a vertex of the hexagon in the plane, two of the double cranks are forced into straight lines. In the configuration space, therefore, there are only two points instead of eight that correspond to each vertex of the planar hexagon, one point for each way in which the third double crank can bend. The four abstract hexagons that correspond to the four ways in which the two straight double cranks can bend when the central pin moves inside the planar hexagon must share a vertex in the configuration space.

The configuration space is a surface that, unlike the planar hexagon, has no corners and no boundary. The surface can be tiled, or covered, with eight hexagons. There are 6×2^2, or 24, edges between the tiles and 6×2, or 12, vertexes where four tiles meet.

The description of the configuration space for the three double cranks that we have given so far is logically complete, because all the abstract gluings have been specified. Nevertheless, it is much more satisfying to carry out the gluings and exhibit the manifold as a closed surface in space. It turns out that such a construction is always possible when the gluing description gives rise to a manifold satisfying a technical condition called orientability. The manifold we have described is orientable and so the gluing can be done, but it is not straightforward.

It was proved in the mid-19th century that every orientable two-manifold is topologically equivalent to the surface of a doughnut with some number of holes. The number is called the genus of the surface. For example, the sphere is a surface of genus 0. The genus of the surface of a one-handle coffee cup, like the genus of the surface of a one-hole doughnut, is 1. The genus of the surface of a pretzel depends on the brand.

For any surface divided into polygonal cells of arbitrary shape the number of polygonal faces minus the number of edges plus the number of vertexes is a numerical constant that depends only on the surface. Remarkably, the number is independent of the way in which the surface is divided into polygonal cells. The constant is called the Euler number, after the Swiss mathematician Leonhard Euler. The Euler number of a surface of genus n is equal to $2 - 2n$. Because the surface can be curved in space, the polygons need not be planar and their edges may curve almost arbitrarily. For example, a sphere can be divided into eight triangles by connecting the north and south poles to four points along the equator. Because there are no holes on the sphere, its genus is 0, and its Euler number should be $2 - 2 \times 0$, or 2. One can easily verify that there are six vertexes and 12 edges on the surface, and so the Euler number of the eight triangular regions on the sphere is indeed 2. It is worth noting that the number of faces, vertexes and edges in the example is also characteristic of the regular octahedron, which is topologically equivalent to the sphere.

Since the configuration space of the three double cranks can be divided into eight hexagonal faces, with 24 edges and 12 vertexes, the Euler number of the configuration space must be $8 - 24 + 12$, or -4. The genus n of the manifold can be calculated by setting the Euler number -4 equal to $2 - 2n$. It follows that n is equal to 3. The eight hexagons in the configuration space for the three double cranks can be depicted in their proper relation to one another on a three-hole doughnut [*see illustration on next page*].

The visual representation of a manifold such as the three-hole doughnut is satisfying because it is concrete, but it also has a number of disadvantages. For example, many of the symmetries that are present in the abstract description of the manifold must be given up in order to picture the manifold in ordinary space. In the abstract description we initially gave for the configuration space of the three double cranks every hexagon is congruent to every other hexagon. Moreover, a rotation of any one hexagon by 120 or 240 degrees leaves its shape unchanged. In the visual representation of the configuration space, however, most of the abstract symmetry has been lost. The hexagons on the three-hole doughnut are neither congruent to one another nor rotationally symmetrical: a rotation of one of them by 120 or 240 degrees cannot be done without changing its shape.

Another problem with the three-hole doughnut is that the geometric properties of its surface vary from point to point: properties of the surface around the outer rim are different from the properties of the surface near one of the holes. It must be emphasized that the geometric properties we refer to are intrinsic to the surface. Intrinsic geometry can be determined by measurements made on the surface itself, without reference to the surrounding space where the surface is found. It is to be distinguished from the extrinsic geometry of the surface, which describes how the surface is bent in space. For example, if a flat sheet of paper is bent without distortion to form a cylinder or a cone, both the cylinder and the cone have the same intrinsic geometry as the flat sheet, although their extrinsic geometries are quite different.

The fact that a surface appears to bend in a nonuniform way when it is viewed from above is therefore not a reliable indicator of its intrinsic geometry. What is the intrinsic difference between the inner and outer regions of the surface of the doughnut? Imagine that a small piece is cut out of the convex, outside portion of the doughnut and flat-

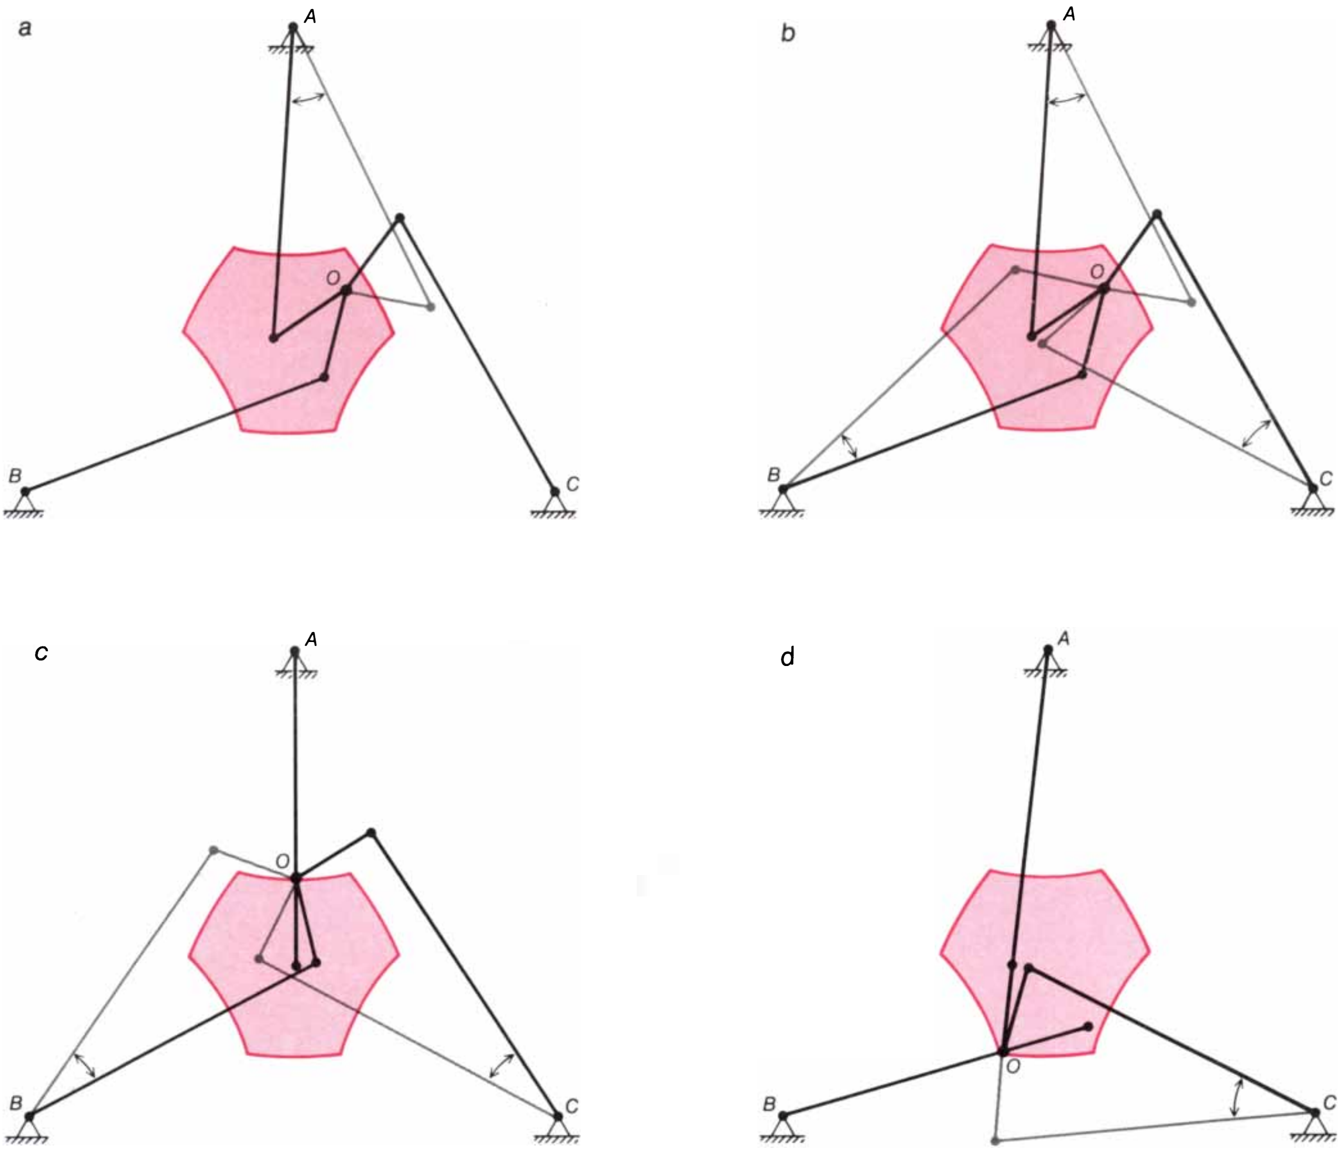

CONFIGURATION SPACE of the three double cranks is a two-manifold on which distinct points represent distinct configurations, or possible arrangements, of the linkage. Every point inside the curvilinear hexagon traced by the central pin of the linkage can be reached when any one of the double cranks is bent in either of two ways (*a*). The bending of each double crank is independent of the bending of the other cranks, and so every point inside the hexagon gives rise to 2^3, or eight, configurations of the linkage (*b*). Whenever the central pin reaches an edge of the hexagon, only two double cranks can be bent; the linkage can assume only four configurations (*c*). When the central pin reaches a vertex of the hexagon, only one of the double cranks can bend; the linkage can assume only two configurations (*d*).

tened on a table. The piece rips open as it is flattened, much like the peel of an orange. The reverse of the process is often exploited by tailors to form a part of a garment intended to fit a convex shape, such as the bust of a dress. A pointed section called a dart is removed from the fabric and the two sides of the gap are sewn together.

On the other hand, when a small piece is cut out of the surface of the doughnut near the hole, the piece wrinkles and overlaps itself when it is flattened on a table. The tailor can reverse the process by slitting the fabric and sewing a godet, or pointed patch, into the slit. The device is often used to make a skirt that is tight below the knees but flares at the bottom. Whether a finished piece of fabric splits, overlaps or conforms to a flat surface when it is spread out on the surface is an important property of its intrinsic geometry.

The intrinsic geometry of the surface of a one-hole doughnut varies in a way quite similar to that of the three-hole doughnut. As we have emphasized, the one-hole doughnut and the square whose opposite edges are identified have the same topology. On the square, however, the intrinsic geometry is much simpler than it is on the one-hole doughnut: the intrinsic geometry in a small region around every point in the square is the same as the geometry in a small region of the plane. The property holds even for points on the edges or the vertexes of the square. In other words, the intrinsic geometry in any small region of the square whose opposite edges are identified is the same as that in any other small region on the square. When the intrinsic geometry of a manifold has uniformity of this kind, the geometry is said to be locally homogeneous.

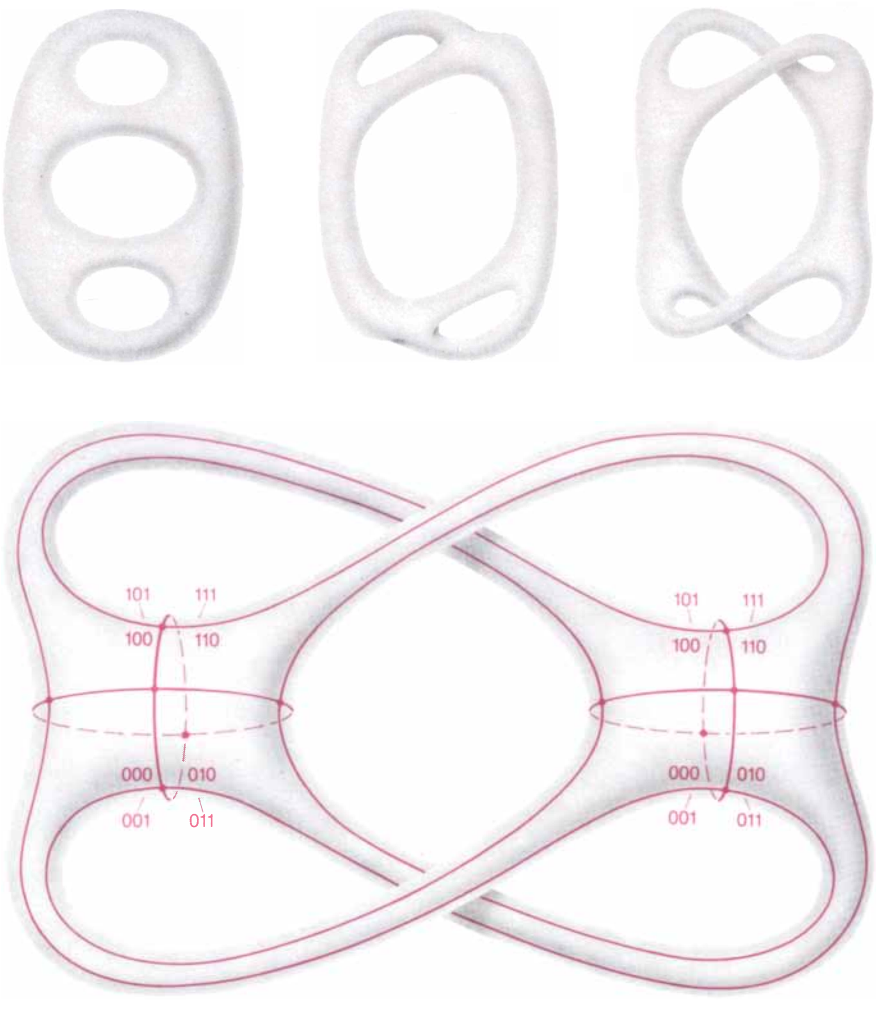

EIGHT ABSTRACT HEXAGONS make up the configuration space of the three double cranks; each hexagon corresponds to one of the eight ways the three double cranks can bend to reach one of the points in the interior of the curvilinear hexagon. If the elbow of a double crank is bent clockwise, its configuration is labeled 0; if it is bent counterclockwise, its configuration is labeled 1. The abstract hexagons are labeled with three binary digits. In order to visualize the eight hexagons and the relations imposed on them in the configuration space the hexagons can be placed on the surface of a three-hole doughnut and distorted as if they were made of rubber. The three-hole doughnut has been deformed to the topologically equivalent manifold at the bottom of the illustration in order to show the eight hexagons more symmetrically. The binary digits assigned to the abstract hexagons reflect the pattern of gluings. If the digits for two hexagons match at two of the three positions, the hexagons are glued along two opposite edges; the two edges correspond to the edges of the curvilinear hexagon where the double crank associated with the nonmatching binary digit is straight. Any four hexagons with one matching binary digit meet at a vertex in the configuration space. The diametrically opposite vertexes of the four hexagons meet at a second point in the space. The two points represent straight configurations of the two double cranks associated with the two nonmatching digit positions.

The introduction of the concept of local homogeneity was an important advance in the understanding of two-manifolds. About 100 years ago it was proved that any surface—not only the one-hole doughnut—can be generated in such a way that its geometry is locally homogeneous. Moreover, no manifold can be given more than one kind of locally homogeneous geometry.

For a surface there are only three kinds of intrinsic geometry that are locally homogeneous. The first kind is the simple Euclidean geometry of the plane. On the plane the circumference of a circle is equal to pi times its diameter, and the sum of the interior angles of a triangle is 180 degrees. The plane is said to have zero Gaussian curvature, which is a measure of the intrinsic shape of a surface first developed by Carl Friedrich Gauss.

The second locally homogeneous geometry is the geometry on the surface of a sphere. A circular cap cut from the surface of a sphere rips open when it is flattened on a plane much like the piece from the convex region of the three-hole doughnut. Hence the circumference of a circle on the sphere is less than the circumference of a circle having the same radius on the plane. The missing circumference suggests the standard name for the locally homogeneous geometry of the sphere: elliptic geometry, from the Greek word for falling short. The interior angles of a triangle constructed on the sphere add up to more than 180 degrees, and the greater the ratio of the area of the triangle to the area of the surface of the sphere, the greater the sum of the angles [*see top illustration on opposite page*]. The sphere has constant positive Gaussian curvature.

As one might expect, a circle cut from a surface having the third kind of locally homogeneous geometry overlaps when it is flattened like the piece from the region near a hole of the three-hole doughnut. The circumference of such a circle is greater than the circumference of the corresponding circle on the plane. The geometry is therefore called hyperbolic geometry, from the Greek word for excess. It is impossible to define a complete hyperbolic surface with any analytic formula, but one can make approximate models of large pieces of such a

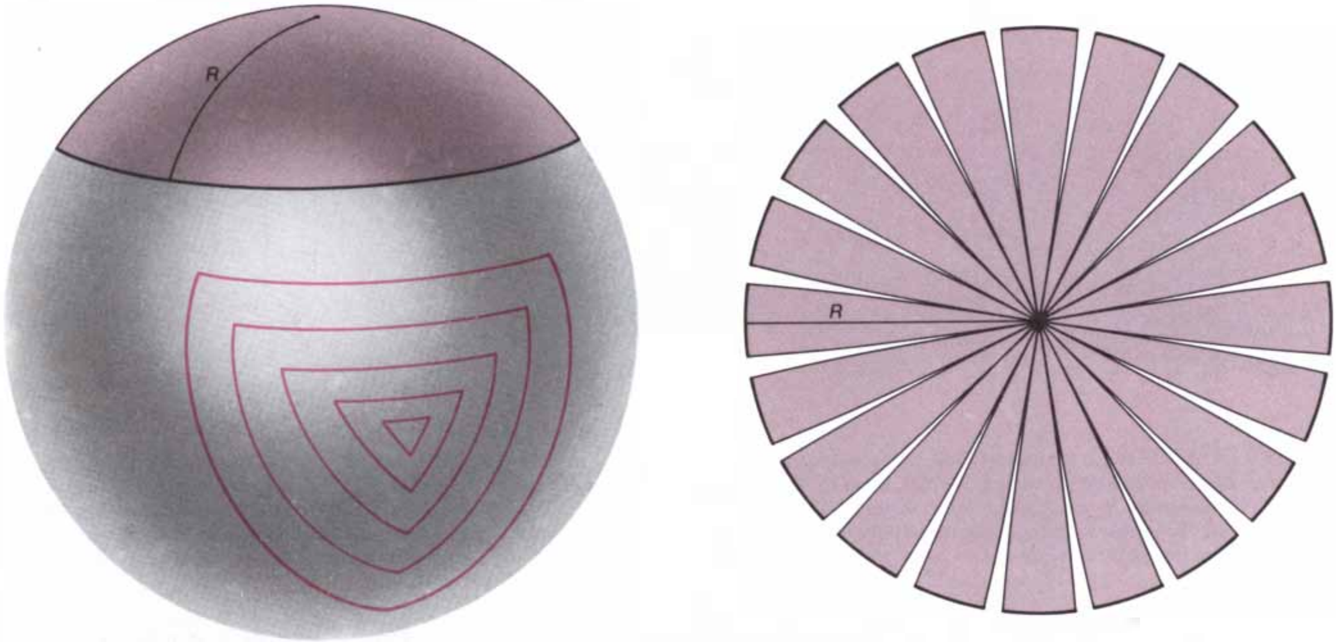

GEOMETRY ON A SPHERE, which is called elliptic geometry, differs from the ordinary Euclidean geometry on the plane. The interior angles of a triangle do not add up to 180 degrees on the sphere as they do on the plane; instead the sum increases with the area of the spherical triangle (*left*). A circular cap cut out of a sphere and flattened on the plane would crack and split as shown at the right. The area of a circle on the sphere is less than the area of a circle that has the same radius on the plane. The sphere has a constant positive curvature.

surface [*see illustration below*]. The interior angles of a triangle constructed on the surface add up to less than 180 degrees, and the greater the area of the triangle, the smaller the sum of the angles. The hyperbolic surface has constant negative Gaussian curvature.

It is instructive to show how the three-hole doughnut can be given a locally homogeneous geometry. Remember that each of the eight hexagons from which the manifold was originally constructed can be bent and deformed in any way, as long as no hexagon is cut or torn. The method is then to deform each hexagon in such a way that it has a locally homogeneous geometry and still fits together with the other hexagonal pieces of the manifold as it does on the three-hole doughnut.

All eight hexagons on the three-hole

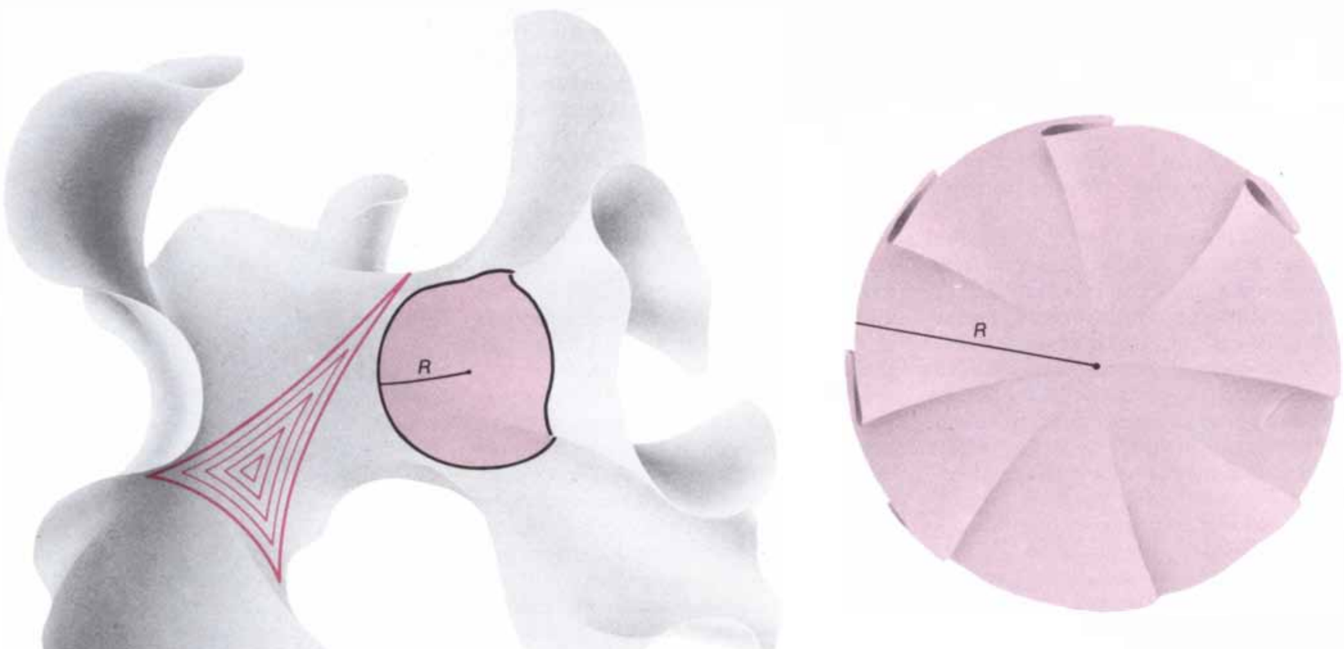

GEOMETRY ON A SURFACE of constant negative curvature (*left*) is called hyperbolic geometry. The sum of the interior angles of a triangle is less than 180 degrees, and the sum decreases as the triangle grows. A circle cut out of a hyperbolic surface would wrinkle and overlap as shown at the right; its area is greater than the area of a circle that has the same radius on the plane. A paper model of a hyperbolic surface can be made by gluing many equilateral triangles along their edges in such a way that seven triangles meet at each vertex.

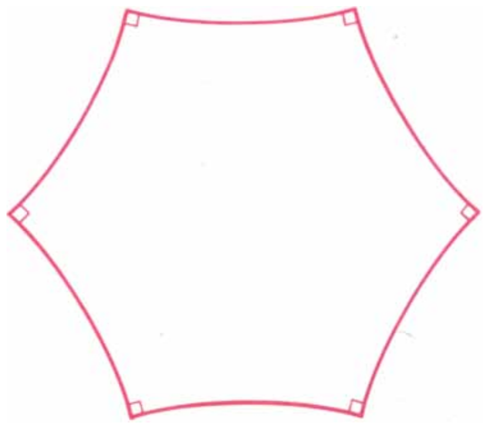

ANGLES OF A REGULAR HEXAGON can be reduced in size if the hexagon is allowed to grow on a hyperbolic surface. In the illustration the hexagon grows until each interior angle is equal to a right angle. The eight hexagons in the configuration space of the three double cranks must meet four hexagons at a vertex, and so all six interior angles of each hexagon must be right angles. If the eight hexagons are topologically deformed into right-angled hyperbolic hexagons and abstractly glued as before, the resulting two-manifold has a constant curvature; its geometry is said to be locally homogeneous. The geometry of the configuration space is thereby greatly simplified.

doughnut meet four at a vertex. If the hexagons were Euclidean, the angles at any vertex would add up to 480 degrees, which is impossible. If the hexagons were spherical, the sum of the four angles at a vertex would be even greater than 480 degrees, which is also impossible. In the hyperbolic plane, however, the larger the polygon, the smaller the interior angles. A sufficiently large hexagon in the hyperbolic plane must have interior angles of 90 degrees; four such hexagons would fit snugly together at a vertex. Hence if the eight hexagons are placed on the surface of a hyperbolic plane and inflated on that plane until each interior angle shrinks to 90 degrees, the manifold constructed by gluing the eight hexagons will have a locally homogeneous, hyperbolic geometry. The manifold cannot be directly visualized in its new form, but the geometric properties of the manifold are much simpler.

The reader may enjoy verifying that the surface of any doughnut with two or more holes can be cut into hexagons that meet four at a vertex. The resulting manifold can be given a hyperbolic geometry by constructing it from right-angled, hyperbolic hexagons. A more traditional procedure is to cut the surface open into a polygon whose vertexes all meet on the surface at one point. The three-hole doughnut, for example, can be split open into a dodecagon, or 12-sided polygon, as well as into eight hexagons. If the surface is sufficiently complex, the polygon derived from the split must have at least six vertexes. If all six vertexes are to fit together properly, the interior angles of the polygon must be reduced. The reduction is done by allowing the polygon to grow in the hyperbolic plane. When the edges of the polygon are glued in pairs, the new surface is topologically identical with the original surface, but it has the locally homogeneous geometry of the hyperbolic plane.

There are only four finite surfaces for which the locally homogeneous geometry is not hyperbolic, because the polygons that arise from the cuts in the surfaces have fewer than six sides. The one-hole doughnut gives rise to a square, and all four corners of the square can be abstractly glued without changing the interior angles of the square. Since no further deformation is needed, the locally homogeneous geometry given to the two-torus is Euclidean. Similarly, the sphere and a nonorientable two-manifold called the projective plane are given an elliptic geometry, and a nonorientable two-manifold called the Klein bottle is given a Euclidean geometry.

A three-manifold can be curved in much the same way as a surface can: every two-dimensional slice of a positively curved three-manifold would split open if it were placed in ordinary Euclidean space, and every two-dimensional slice of a negatively curved three-manifold would wrinkle and overlap. Elliptic geometry, Euclidean geometry and hyperbolic geometry all have their three-dimensional counterparts.

In 1976 one of us (Thurston) began to suggest that locally homogeneous, hyperbolic geometry is the key to the understanding of almost all three-manifolds. The development has come as a surprise to many topologists, because three-manifolds are so much more complicated than two-manifolds. Whereas any orientable two-manifold can be specified and listed according to its genus, every three-manifold, like a tangled loop of string, seems to have its own distinct properties and resists fitting into any larger pattern. On closer scrutiny, however, larger patterns have begun to emerge. The patterns depend on the fact that many three-manifolds can be given a locally homogeneous geometry.

How can such a simple geometric structure be imposed? A procedure quite similar to the one we have described for two-manifolds works for a large number of cases. The manifold is cut apart into a polyhedron, and one must determine how many vertexes of the polyhedron are to be fit in place when it is abstractly glued back together. For example, in the Seifert-Weber space all 20 corners of the dodecahedron that generates the space are abstractly glued together. The solid angle formed at each vertex of a Euclidean dodecahedron is much too large for 20 such vertexes to fit together at a point. If the dodecahedron is placed in three-dimensional hyperbolic space, however, it can be expanded until the solid angle at each vertex is small enough to pack 20 equal vertexes around a point [*see illustration on the cover of this issue*]. When opposite faces of the hyperbolic dodecahedron are abstractly glued together after a rotation of three-tenths of a turn, the resulting manifold is a Seifert-Weber space with a locally homogeneous, hyperbolic geometry.

The Poincaré dodecahedral space is also derived by gluing the faces of a dodecahedron, but the vertexes are glued together in five groups of four. The solid angle at the vertex of an ordinary dodecahedron is slightly too small to pack tightly around a point in groups of four, but a suitably large dodecahedron in a positively curved space has corners that are just the right size. The enlarged dodecahedron makes it possible to construct a Poincaré dodecahedral space with a locally homogeneous, elliptic geometry [*see illustration at right on opposite page*]. In this context it is worth mentioning how a locally homogeneous geometry is given to the three-torus. In the construction of the manifold the eight vertexes of a cube are abstractly glued. Because the eight corners can fit together at a point without distortion, the locally homogeneous geometry of the three-torus is Euclidean.

Lest the reader be misled, we must point out that the preceding examples are not really typical because they are highly symmetrical. When a three-manifold is defined by gluing the faces of an irregular polyhedron, more care must be taken to give the polyhedron a shape that leads to a locally homogeneous geometry for the three-manifold. The shapes of the polyhedral faces that are glued together must match, and the angles between the faces that surround any edge must add up to 360 degrees.

There are at least two major differences between the geometry of two-manifolds and the geometry of three-manifolds. First, there are five more kinds of locally homogeneous geometry that can be given to three-manifolds, in addition to the three we have mentioned. The additional geometries come about because in three or more dimensions an intrinsic curvature is defined for each two-dimensional slice that passes through a point. A locally homogeneous geometry need not have the same curvature in all the two-dimensional slices. Nevertheless, an understanding of the intrinsic curvature of all eight geometries can be based on the simpler geometries of two-manifolds.

The second difference between the two- and three-manifolds might seem to present insurmountable complications. It is possible to combine three-manifolds in such a way as to yield new three-manifolds that cannot be given a locally homogeneous geometry. Fortunately, topologists know how to split a three-manifold into primitive pieces by purely topological methods.

One of us (Thurston) has proposed that after a three-manifold has been cut up into its primitive pieces, each of the resulting pieces does in fact admit a locally homogeneous geometry of one of the eight possible types. The conjecture has been proved for wide classes of manifolds, and it has been empirically tested for many other examples either by hand or with the aid of a computer, and it has never been found wanting. It now seems unlikely a counterexample will be found.

The empirical studies also suggest that for most three-manifolds the complications of three-manifold geometry do not come into play. Indeed, it has been proved that, in a certain technical sense of the word, "most" three-manifolds can be given a locally hyperbolic geometry. The finding is fortunate, because hyperbolic three-manifolds have many beautiful properties. For example, in 1971 G. D. Mostow of Yale University proved that if a three-manifold can be given a locally hyperbolic geometry, the geometry is completely determined by the topology. A consequence of Mostow's theorem is that all manifolds having a locally homogeneous geometry can in principle be classified. Moreover, for hyperbolic three-manifolds the theorem gives a rough-and-ready test of identity. When a manifold is deformed into its geometrically tractable form, its volume can be measured, and the theorem guarantees that the volume depends only on the topological type of the manifold. It can often be quite difficult to distinguish two manifolds in their arbitrary topological form, and so the volume turns out to be a handy signature for each manifold.

With these results in mind it is worth returning to our initial speculations about the topological structure of the universe. Observational evidence suggests the universe is homogeneous everywhere and has either elliptic, hyperbolic or Euclidean geometry. There is also strong support for the theory that the universe is currently in a stage of expansion that has continued since the beginning of the big bang. It is interesting to speculate on what the distant future holds in store, but there are essentially only two possibilities. One is that the mutual gravitational attraction of the matter in the universe will finally

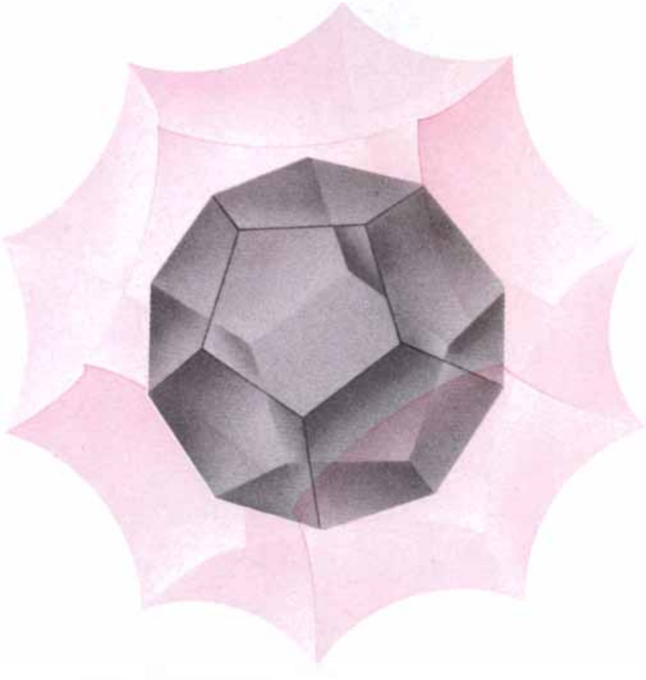

SEIFERT-WEBER SPACE can be given a locally hyperbolic geometry if the dodecahedron that generates the space is allowed to grow in hyperbolic space. The growth of a polyhedron in hyperbolic space is similar to the growth of a polygon on a hyperbolic surface. When the dodecahedron grows, the solid angle at each vertex shrinks, and so each vertex becomes progressively sharper. The abstract gluings that lead to the Seifert-Weber space specify that all 20 vertexes of the dodecahedron must meet at a point. The solid angle at each vertex must therefore be shrunk in the hyperbolic space until all 20 of the solid angles are small enough to fit together around a single point.

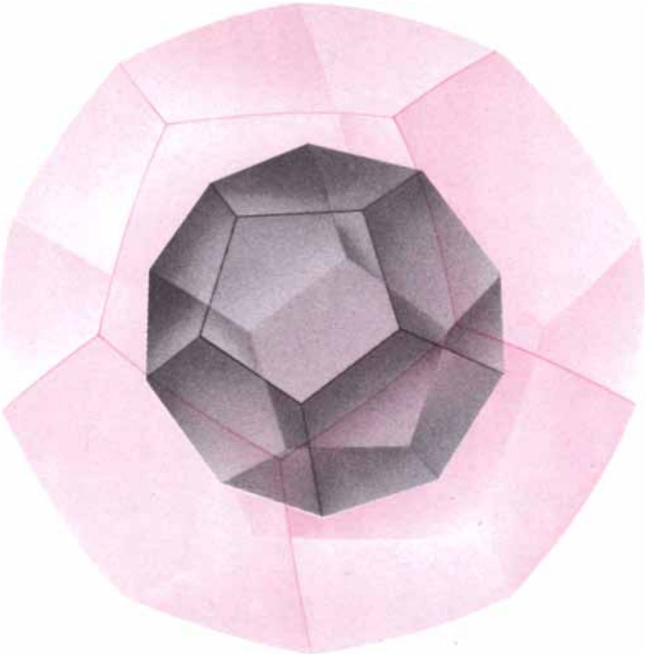

POINCARÉ DODECAHEDRAL SPACE is also generated by gluing pairs of opposite faces of a dodecahedron, but one member of each pair is matched to its counterpart after a rotation of one-tenth of a turn instead of the three-tenths of a turn required for the Seifert-Weber space. The abstract gluings lead to the identification of four vertexes of the dodecahedron at every vertex in the manifold. The solid angle at each vertex of an ordinary dodecahedron is slightly too small for four dodecahedrons to fit snugly around a point, but the solid angles can be increased if the dodecahedron is inflated in elliptic space. The effect is the reverse of the inflation in hyperbolic space.

halt the expansion and cause the universe to recollapse in a "big crunch." A second possibility is that the gravitational attraction is not strong enough to halt the expansion, and the universe will expand forever.

One consequence of the general theory of relativity is that the ultimate fate of the universe depends on its geometry. If the universe has an elliptic geometry, it will eventually recollapse. If it has a hyperbolic geometry, it will expand forever. If it has a Euclidean geometry, it will also expand forever, but the rate of the expansion will approach zero. In principle it would be possible to determine the geometry of the universe by laying out a huge triangle and accurately measuring its interior angles. If the sum of the angles were greater than 180 degrees, the geometry of space would be elliptic; if the sum of the angles were equal to 180 degrees, the geometry would be Euclidean and if the sum were less than 180 degrees, the geometry would be hyperbolic. In practice cosmologists try to estimate the density of the matter in the universe and the rate of the expansion, because the geometry of the universe can be correlated with the two measurements. If the density is high enough for a given rate of expansion, the universe will recollapse.

There is a widespread misconception, however, that the curvature of the universe determines whether the universe is finite or infinite in extent. It is often asserted that if the universe is finite, its geometry must be elliptic and, conversely, that if the geometry of the universe is hyperbolic, the universe must be infinite. The Seifert-Weber space, which is a finite three-manifold with a locally hyperbolic geometry, shows that neither of these beliefs is true. Indeed, most finite topological models of space are three-manifolds like the Seifert-Weber space with a locally hyperbolic geometry. Such manifolds yield models of a finite universe that expands forever.

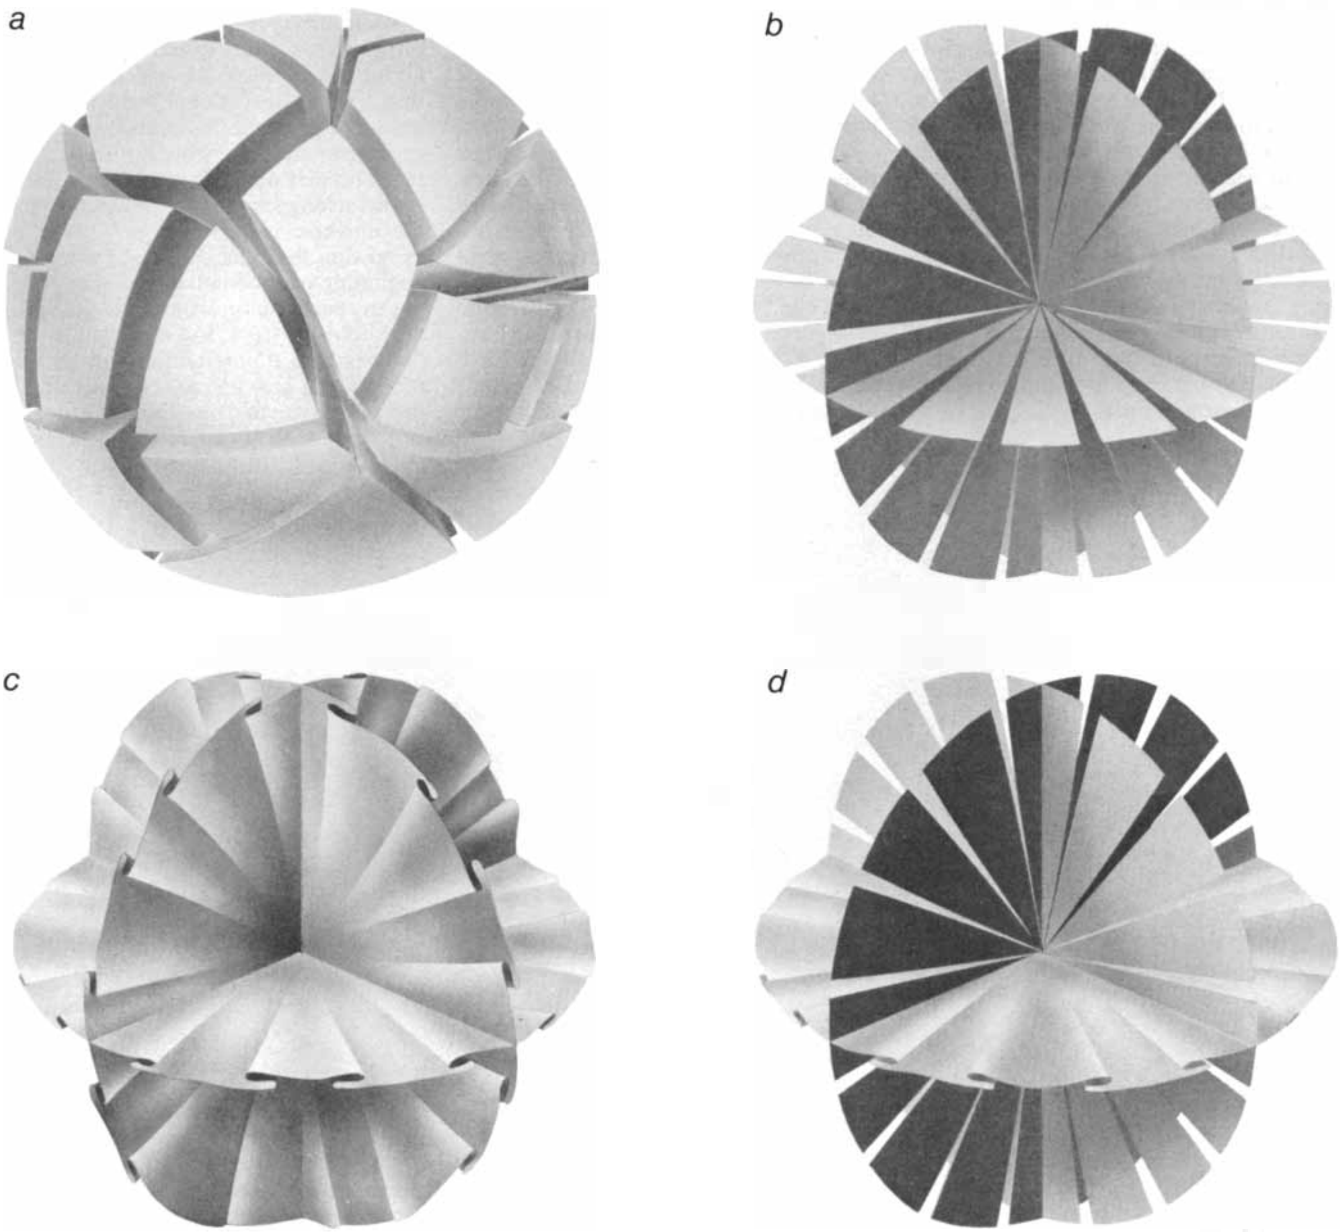

SPHERICAL "SLICE" of a curved three-manifold, analogous to the circular pieces of the two-manifolds shown in the illustrations on page 117, cannot be fitted into ordinary Euclidean space without deformation. A positively curved three-manifold would split open everywhere (*a*); every two-dimensional slice of the manifold has the curvature of an ordinary sphere (*b*). Similarly, every two-dimensional slice of a negatively curved three-manifold would wrinkle and overlap as if it were cut out of a hyperbolic surface (*c*). A three-manifold whose curvature varies with direction can still have a locally homogeneous geometry, as long as the pattern is the same at every point. For example, one slice of the manifold could wrinkle and overlap in ordinary space, whereas two other slices could split apart (*d*).

8087 SUPPORT and FORTRAN for your PC

87FORTRAN/RTOS™ - our adaptation of the Intel Fortran-86 Compiler generates in line 8087 code using all 8087 data types including 80-bit reals and 64-bit intergers. The compiler uses the Intel large memory model, allowing code/data structures of a full megabyte, and supports overlays. Includes RTOS and support for one year **$1350**

87PASCAL/RTOS™ is Intel's ISO-Standard Pascal with 8087-8088 exceptions. These make it possible to use all the 8087 data types directly, while generating modules in one of the three Intel Memory Models. Includes RTOS and support for one year **$1350**

RTOS DEVELOPMENT PACKAGE includes 87FORTRAN, 87PASCAL, PL/M-86, Utilities, TX Screen Editor and RTOS... **$2500**

RTOS—
REAL TIME MULTI-TASKING/ MULTI-USER EXECUTIVE

RTOS is a Micro Ware configured version of iRMX-86. Includes ASM-86, LINK-86, LOC-86, LIB-86, and the ROM Hex Loader **$600**

87BASIC/INLINE™ generates inline 8087 code! Converts the IBM Basic Compiler output into an assembly language source listing which allows the user to make additional refinements to his program. Real expression evaluations run five times faster than in 87BASIC **$200**

MWS-286™ Configured to your specifications, our computer runs RTOS-286 or XENIX. Includes one Intel compiler, seven slot multibus chassis, hard disk, streaming tape backup and Intel Service Contract. Six to twenty times faster than your PC .... **CALL**

87BASIC™ includes patches to the IBM Basic Compiler and both runtime libraries for USER TRANSPARENT and COMPLETE 8087 support. Provides super fast performance for all numeric operations including trigonometrics, transcendentals, addition, subtraction, multiplication, and division **$150**

87MACRO™ - our complete 8087 software development package. It contains a "Pre-processor," source code for a set of 8087 macros, and a library of numeric functions including transcendentals, trigonometrics, hyperbolics, encoding, decoding and conversions... **$150**

87DEBUG™ - a professional debugger with 8087 support, a sophisticated screen-oriented macro command processor, and trace features which include the ability to skip tracing through branches to calls and software and hardware interrupts. Breakpoints can be set in code or on guarded addresses in RAM **$150**

FOR→BAS™ - a library of interface routines which allow MS Fortran programs to call the IBM Basic Compiler library and access features such as the RANDOM NUMBER GENERATOR, SOUND, PLAY, DRAW and SCREEN commands **$150**

8087-3 CHIP $175
including DIAGNOSTICS and 180-day warranty

64K RAM Upgrade $47.50

87/88GUIDE - an excellent tutorial on writing 8087 code and interfacing it with compilers. Full of code that runs! **$30**

MATRIXPAK™ manages a **MEGABYTE!** Written in assembly language, our runtime package accurately manipulates large matrices at very fast speeds. Includes matrix inversion and the solution of simultaneous linear equations. Callable from MS Fortran 3.2, 87MACRO, 87BASIC, and RTOS .... each **$150**

MICROSOFT FORTRAN 3.2 or PASCAL 3.2 - These IEEE compatible compilers support double precision and the 8087 .... each **$259**

MICROSOFT C COMPILER includes Lattice C and the MS Librarian **$350**

SCO XENIX. ™MICROSOFT	**$595**
Unisource VENIX/86 ™VenturCom, Inc.	**800**
FLOAT87 for MS or LATTICE C	**125**
SuperSoft Fortran 66	**299**
Computer Innovations C86	**345**
STSC APL★PLUS/PC	**500**
TURBO PASCAL	**45**
TURBO PASCAL with 8087 Support	**85**
HALO GRAPHICS	**CALL**
GRAPHMATIC	**125**
ENERGRAPHICS	**295**
Professional BASIC	**295**
Kidger Optical Design Program	**3000**
COSMOS REVELATION	**850**
SuperCalc III with 8087 Support	**CALL**
MAYNARD WS1 HARD DISK	**995**
MAYNARD WS2 HARD DISK	**1170**
256K RAM CHIPS	**CALL**
MAYNARD ELECTRONICS Boards	**CALL**

ALL ITEMS IN STOCK

Micro Ware
P.O. Box 79
Kingston, MA
02364
(617) 746-7341

You Can Talk To Us!

HOW TO REACH THE PEOPLE WHO MAKE THE FUTURE HAPPEN IN SPAIN

INVESTIGACION Y CIENCIA is the Spanish edition of SCIENTIFIC AMERICAN which attracts a young audience of affluent professionals who have come to the top because of their technical expertise. Almost half of them hold top management job titles. The fortunes of major corporations in Spain increasingly depend on technically sophisticated people. INVESTIGACION Y CIENCIA keeps these people in touch with the advances in science that drive the growth of Spain's industry.

SCIENTIFIC AMERICAN speaks the languages of more than half the world's population. We are the one publication in the world today providing efficient coverage of technology-based management. We reach men and women in industry, whose qualification to make technical decisions places them in key positions in their country's government and industry.

SCIENTIFIC AMERICAN in eight languages: English, Spanish, Italian, French, German, Japanese, Chinese and now Russian has gathered in its audience three million people who make the future happen around the world.

For more information on our Spanish-language edition contact:

Elena Sanchez-Fabres
Prensa Scientifica, s.a.
Calabria 235-239
Barcelona 29, SPAIN
Telephone (011) 343-322-0551

or in New York
John Kirby,
V.P./Advertising Director
SCIENTIFIC AMERICAN
415 Madison Avenue
New York, New York 10017
Telephone 212-754-0262