CHAPTER 1

An Introduction to the Shape of the Universe

Topics:

- 2-dimensional universes
- Straightest paths in different spaces
- Gluing diagrams
- The flat torus and the curved torus
- Games on a flat cylinder and a flat torus
- The extended diagram for a flat torus
- The 3-torus
- The 3-sphere

---

1. To Infinity and Beyond

People normally assume that our universe goes off infinitely in all directions, but the true shape and size of our universe is not known. It is hard to believe that we really don't know something as basic as the shape of our universe, and it's hard to imagine how the universe could *not* go on forever in all directions. If the universe were finite, would that mean that the universe has an end that you can't go beyond?

Let's imagine the following scenario, featuring an astronaut whose name is A. 3D-Girl. Like her name, she is extraordinary and super-duper in every way. She's smart, funny, charming, adventurous, brave, and anything else you'd want in an interstellar explorer. She's been sent off deep into space to investigate a potential colonization site that was identified by astronomers through super-duper technologically advanced telescopes. The astronomers were pretty excited when they saw this planet, because it seemed to have many of the same characteristics as Earth. They even named it "Earth 2". A. 3D-Girl's spaceship, aptly named A. Spaceship, has a super-duper hyper-warp drive that allows it to zoom off in any direction, going as fast as she wants and as far as she wants. After months of planning and training for the mission, A. 3D-Girl blasts off in her spaceship, traveling in a straight line past Mars, Jupiter, Saturn, Uranus, Neptune, and Pluto, beyond Alpha Centauri, in fact beyond all known stars and planets, until Earth 2 finally comes into view.

As the spaceship gets close to Earth 2, A. 3D-Girl notices more similarities with our earth. Earth 2 is a blue and green planet with a single moon. It's also the third planet from its sun. After landing, she cautiously descends from her spaceship and looks around. To her great relief, the place feels comfortable and

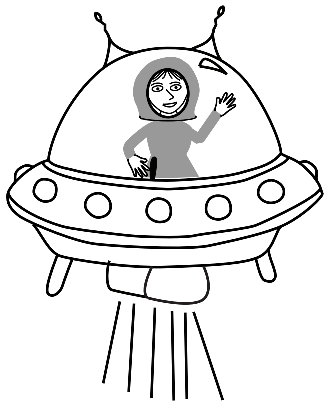

FIGURE 1. A. 3D-Girl waving to the crowd as her spaceship blasts off.

familiar. After looking around a bit more, she does a double take. This place is *too* familiar. There's even a doughnut shop that looks exactly like the one that A. 3D-Girl stopped at before she boarded her spaceship. She checks and rechecks her flight log. Nope, she didn't make a mistake in navigation, and didn't make any U-turns. In fact, she didn't make any turns at all in getting here from Earth. She set A. Spaceship to full throttle and went along the straightest, shortest path from Earth to the exact coordinates where the astronomers had told her to go. As crowds of people gather around her, A. 3D-Girl boldly remarks, "Huh? What happened?"

Although it is hard to imagine ever developing a spaceship as technologically advanced as A. 3D-Girl's, this story is not completely unbelievable. Many elementary schools teach children that people in the middle ages thought the earth was flat, and it was not until 1492 when Christopher Columbus sailed the ocean blue that people found out that the earth is round. In fact, this is a myth that was put forth in 17th-century works of fiction including Washington Irving's fictionalized biography *The Life and Voyages of Christopher Columbus.* But what is noteworthy about this myth is that to this day people accept it as completely reasonable. We can imagine ourselves in the middle ages, looking out in all directions and concluding that either the Earth must have an edge way off in the distance or it must be a plane that goes on forever.

The story about A. 3D-Girl's voyage isn't that different from the myth that Christopher Columbus discovered that the earth is round. When we look out into space in all directions, we can't imagine how space could wrap around and come back to where it started. So we assume that either the universe must end way off in the distance or it must go on forever. But just because we don't know how to visualize a universe that closes back on itself like the Earth, doesn't mean it isn't possible.

Though the actual shape of the universe is not known, in this and the next seven chapters, we'll develop our intuition about some of the possibilities for the shape of our universe. We'll learn how to visualize universes of different dimensions, and how we might distinguish one such universe from another. While you may end up with more questions than you started with about our own universe, we hope you'll find the voyage intriguing.

2. A. Square and His Universe

Sometimes, when you're faced with a hard math problem, it's a good idea to start by solving an easier version of the same problem. So instead of thinking about the shape of our universe, we begin one dimension down, by considering how 2-dimensional creatures living in a 2-dimensional universe might think about possible shapes for their universe.

You might have heard of the short novel titled *Flatland: A Romance of Many Dimensions* which was written in 1884 by the English school teacher Edwin A. Abbott. The book was meant to be a commentary on Victorian society, but it can give us some insight into the world of 2 dimensions. Whether or not you've read *Flatland* does not matter here (though if you would like to read it, you can download a copy from the internet for free). We will simply introduce you to the main character, whose name is A. Square. His portrait is given in Figure 2.

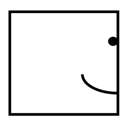

FIGURE 2. A portrait of A. Square.

A. Square lives in a 2-dimensional universe. He and his space have no thickness, like a shadow or an image projected on a screen. In the book *Flatland*, A. Square's universe is a plane. However, there are other possibilities for his 2-dimensional universe. For example, his universe could be a sphere, a torus (that is the surface of a doughnut), a 2-holed torus, a 3-holed torus, etc. (see Figure 3). Continuing in this way, we find that there are infinitely many 2-dimensional universes.

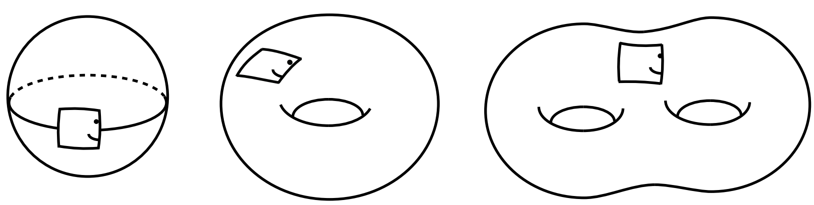

FIGURE 3. A. Square on the surfaces of a sphere, a torus, and a 2-holed torus.

While you might think of these surfaces as 3-dimensional objects in our space, keep in mind that we are only interested in the surfaces of the doughnuts (or the doughnut hole when he's on a sphere) and not the cake inside. This means that a 2-dimensional being like A. Square could live in one of these universes, but a 3-dimensional person could not. Furthermore, if A. Square lived in one of these universes, he could only move back and forth and up and down, staying in the icing of the doughnut with no way to get to (or even known about) the cake underneath.

We say that A. Square's universe is 2-*dimensional* because he can only move in two perpendicular directions. Back and forth is one direction, and up and down is another direction. Of course he can combine these two directions to go in any diagonal direction he wants.

In general, the number of *dimensions* of a space is the number of perpendicular directions in which creatures in the space can move. A line and a circle each have only 1 dimension because a creature living inside one of these spaces can only move backward and forward. This is true no matter how the line or circle is situated in our space. In particular, even though we need 3 dimensions to tie a knot, if we consider a knotted circle as a space by itself, it would only be 1-dimensional. Figure 4 illustrates A. Dot who lives in a circle and her cousin A. Dash who lives in a knotted circle. Each of the cousins is only able to move backwards and forwards in her space. They seem happy, but imagine what boring lives they must have. They talk on the phone on a daily basis to report how many laps around their space they run. However, they have no idea that their spaces are different because the difference can only been seen with our 3-dimensional perspective.

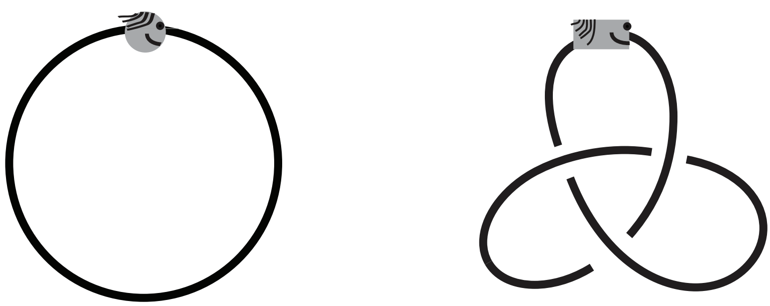

FIGURE 4. A. Dot lives in a circle and A. Dash lives in a knotted circle.

Notice that even though the surfaces in Figure 3 have no boundaries, they all have finite area. Figure 5 illustrates some 2-dimensional universes with infinite area. Since there is only a finite amount of space on a page, we use dotted lines to mean that the spaces go on forever. The illustration on the left is an infinite cylinder. The illustration next to it has two planes joined with a tube and a handle on the lower plane. The illustration on the bottom has a tube and two such handles. Keep in mind that just like the surfaces in Figure 3, these universes consist only of surfaces. In particular, A. Square and his friends can't get to the 3-dimensional space inside of the tubes or handles.

Observe that we could create a 2-dimensional universe with any number of planes joined together with any number of tubes, and each plane could have any number of handles stuck to it. Pushing this idea even further, we could have infinitely many planes joined with infinitely many tubes, and each plane could have infinitely many handles. This gives us infinitely many 2-dimensional universes all with infinite area.

Of course, A. Square can't picture any of these spaces. He thinks there is only one possible universe, and it's a plane that goes on forever in all directions. We denote such an infinite plane by $\mathbb{R}^2$, since we can think of it as the coordinate space with an x-axis and a y-axis. Similarly, we use $\mathbb{R}^3$ to denote the coordinate space

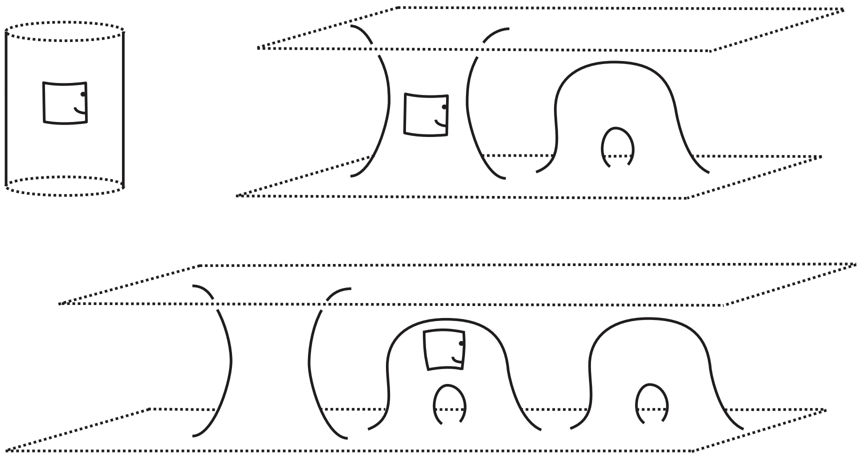

FIGURE 5. Some 2-dimensional universes with infinite area.

with x, y, and z axes. Knowing that there are infinitely many finite and infinite 2-dimensional universes gives us the idea that perhaps our own 3-dimensional universe could be something besides $\mathbb{R}^3$. But we will have to wait until later in the chapter to see some possibilities.

Maybe we should consider the disk illustrated in Figure 6 as a possible 2-dimensional universe. The problem is that a disk has a boundary. If A. Square reached the boundary, he would hit a wall and could go no further. He could not fall off the edge of the universe, because the universe is all there is. He can't step outside his universe. He also couldn't go around the edge of the disk and come back on the bottom because, since the disk is only 2-dimensional, it doesn't have a top and a bottom. Imagine if there was a similar wall in our 3-dimensional universe. Suppose we sent out a spaceship and it smashed into a wall at the edge of our universe. This seems neither pleasant nor likely.

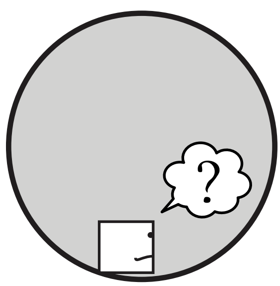

FIGURE 6. If A. Square lived in a disk, when he reached the edge he could go no further.

As we think about possible universes for A. Square and for ourselves, we do not want to consider the possibility that universes might have walls that you could smash into. This just seems too implausible to be true. Thus, from now on whenever we use the word *universe*, we will mean a space that has no boundary. Note that this does not mean that we require every universe to be infinite. Keep in mind that there are universes like those in Figure 3 which have finite area but no boundary. By contrast, when we use the word *space*, we won't make any assumptions about whether or not the space has a boundary. For example, we say that a disk like the one in Figure 6 is a 2-dimensional *space* but not a 2-dimensional *universe*.

Now that we know that there are so many possibilities for 2-dimensional universes, we want to help A. Square learn about his own universe. Ideally, he would like to create a list of all possible 2-dimensional universes. Even though such a list would be infinite, if it was set up in a systematic way, it would give him an idea what the possibilities were. Even if his list isn't complete, he would like to determine some *property* of his universe that would enable him to distinguish it from other spaces on the list or at least to reduce the number of possibilities. For example, if he could determine that his universe had finite area, he could eliminate all of the spaces with planes and tubes and handles.

You can think about properties of a space the way we think about characteristics of people. For example, suppose you had a list of all of the students in your math class, and you wanted to find out the name of the girl who sits next to you. Knowing that she's a girl might help you to reduce the number of possibilities, but you still wouldn't know which one was her. If the list included major and year of graduation, and you were able to find out this information about her, you could reduce the possibilities even further.

We'd like to do the same thing with properties of a universe. For example, suppose that A. Square somehow determines that his universe is one of the infinite surfaces illustrated in Figure 5. The number of planes, tubes, and handles that the surface has is a property which would enable us to distinguish them. However, if A. Square were living in one of these universes, he would not be able to see the planes, tubes, and handles like we can. So we need to think of properties that A. Square could use to distinguish among these universes. If we can succeed at that, perhaps it will give us some insight about properties we could use to learn about our own universe.

3. Straightest Paths in Flatland

One of the tools that A. Square can use in investigating his universe is the notion of a "straight line" between two points. You've probably heard people say "the shortest path between two points is a straight line". This is true if the two points are in a plane. However, if A. Square is on a sphere, what does "straight line" even mean? There are no straight lines or even straight line segments on a sphere. A path which appears to A. Square to be the shortest path between two points will appear to us as a curve since every path on a sphere is a curve.

One way to think about what "straight" means to A. Square is to imagine that he has a 2-dimensional flashlight that he is shining in front of him. We could picture his flashlight like the lamp that's attached to a miner's helmet. A 2-dimensional ray of light from his flashlight stays within his universe, illuminating the shortest path to get to a nearby point. So if he follows the beam of his flashlight, his path will appear to him to be a straight line. We will call the path of a light ray within his universe a *straightest path*, even if it doesn't actually follow a straight line in our 3-dimensional universe.

Another approach that A. Square could take to finding the straightest path between two points would be to attach a piece of string to one of the points in his space and then go to the other point and pull the string tight. Then if he walks along the string, he knows he will be going as straight as he can.

If A. Square lives in a plane, then the straightest path between two points is a line segment. But if he lives in a sphere, then the straightest path between two points is an arc on what's known as a *great circle*. The great circles on a sphere are circles whose radius is the same as the radius of the sphere. An equator is an example of a great circle. Great circles can be thought of as the intersection of a sphere with a plane that divides the sphere into two equal halves. Note that if the plane does not divide the sphere exactly in half, then the circle of intersection is not a great circle. Paths which follow great circles curve as little as possible while staying within the surface of the sphere. Airplanes and ships follow great circles around the earth because these are the shortest paths between two points on earth. If you have looked at flight paths in airline magazines, you may have noticed that the flight paths are drawn as curves rather than as straight lines. These flight paths appear to be longer than necessary, but that's because a flat map of our spherical earth will necessarily be distorted.

Let's consider how A. Square might use the idea of straightest paths to distinguish between a plane and a sphere. Suppose his universe is a plane, and he leaves home to take a trip following a straightest path (making use of his headlamp to go straight). Much to the disappointment of his family, if he continues following this straightest path, he will never return home. However, if A. Square lived on a sphere and left on a similar voyage following a straightest path, after some time he would happily return to his awaiting loved ones (see Figure 7). At this point he would know that his universe could not be a plane. It's important in this experiment that A. Square use a taut string or a light ray to follow a straightest path. Otherwise, he could just walk in a circle on the plane and return to his starting point.

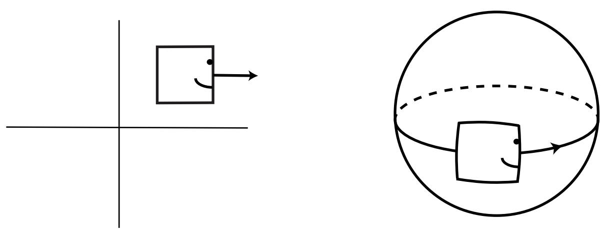

FIGURE 7. Straightest paths on a plane and a sphere.

In order to see if a similar method would enable A. Square to differentiate between a sphere and a torus, consider the path illustrated in Figure 8. A string could be pulled taut along this path so that A. Square would know that it's indeed a straightest path. If A. Square travels along this path, he will eventually return to his starting point as he did on the sphere. Thus finding a straightest path that returns to its starting point will enable A. Square to distinguish between a torus and a plane, but will not give him enough information to distinguish between a torus and a sphere.

In search of some other property that a torus has and a sphere doesn't have, A. Square tries to find a different type of straightest path on the torus. The problem

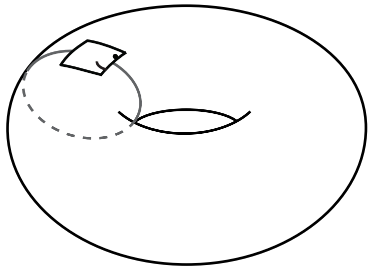

FIGURE 8. A straightest path on a torus.

is that finding such paths on the torus is not as straightforward as it seems. For example, take a look at the grey circle on the left of Figure 9. At first glance, we might think that it is a straightest path. It's nice and round and rests just at the top of the surface. However, if A. Square marks points p and q on the grey circle and pulls a string tight between these points, the string will not lie on the grey path. The black string between points p and q on the left in Figure 9 is just a bit shorter than the grey path going between the points. So the grey path on the left of Figure 9 isn't a straightest path on the torus after all. While it's hard to identify a straightest path on the top of a torus, an inner circle going around the hole of the torus is a straightest path because the circle cannot be made any shorter. Such a shortest path is illustrated on the right of Figure 9

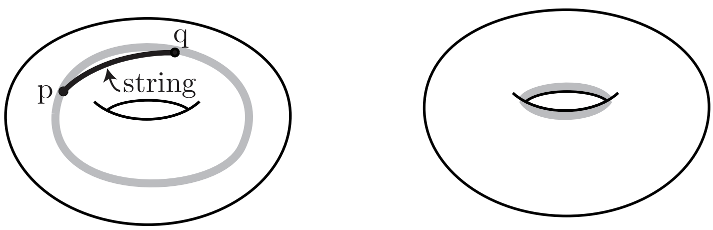

FIGURE 9. The grey circle on the left is not a straightest path, but the grey circle on the right is.

When looking at the illustration on the left of Figure 9, it's important to keep in mind that to us the black string is curved, though not as curved as the grey string. A truly straight segment would cut through the inside of the torus, but A. Square's string could not leave the surface of the torus to follow such a path.

Now let's imagine that A. Square lives on a torus and has two balls of string, one black and the other one grey. He pins down the end of the black string, then follows the straightest path illustrated on the left in Figure 10 pulling the string taut as he goes. After returning to his starting point, he ties the two ends of the black string together. Then he repeats this process with the ball of grey string going around the inner circle of the torus, again as illustrated on the left in Figure 10. He notices that the black and grey circles are both straightest paths, and they intersect at precisely one point (indicated by a dark grey dot).

On the other hand, if A. Square lives on a sphere, any two straightest paths will intersect in two points as illustrated by the black and grey circles on the right

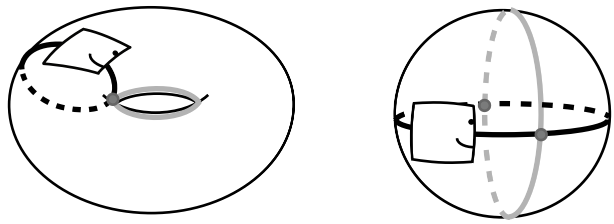

Figure 10. Two straightest paths can meet in exactly one point on a torus, but not on a sphere.

in Figure 10. Thus this experiment with colored string will allow him to distinguish between a torus and a sphere.

Could A. Square use additional balls of colored string to distinguish between tori with different numbers of holes? He would have to begin by finding some straightest paths on a torus with two holes. You may want to experiment with this yourself. Get a bagel, a giant pretzel, some pins, and a few skeins of colored yarn, and use them to try to distinguish the bagel from the pretzel. (How many holes does a pretzel have?) Afterwards, when you eat the bagel and pretzel make sure you first remove the yarn and the pins.

4. Exploring the Shape of a Cave

Having watched A. Square use straightest paths to compare different universes, we would like to use similar ideas to compare 3-dimensional universes. But before considering such a hard problem, let's check on A. 3D-Girl, who has been suffering from nightmares and insomnia ever since her confusing voyage to Earth 2.

During a restless night of sleep, A. 3D-Girl thinks she's woken up in a deep, dim, dank cave. Feeling her way along the wall, she discovers that the cave has no exit and no entrance to any other passageway. While exploring the entire cave, she starts thinking about a book she read a long time ago about a square trying to figure out the shape of his universe, and suddenly a ball of grey string appears in her hand. She places one end of the string under a big rock so it can't move, and trails the string behind her as she walks. Eventually she returns to the big rock where she started without ever retracing her steps. From this exploration she concludes that the cave is like the inside of a doughnut. The cave has some nooks and crannies as all caves do, but the overall shape is as illustrated in Figure 11. Depressed at the thought of a life walking around in circles trailing grey string, she decides to go back to sleep hoping to wake up in her own comfortable bed.

Unfortunately, the next thing she knows she's in a different cave, again with no exit and again with a ball of grey string in her hand. After exploring for a while in the dark, she discovers that this cave has a single four-way intersection. She winds the string around both circular passageways making it into a figure-eight. She pictures this cave as illustrated in Figure 12.

She becomes agitated as she dreams that she is waking up in more and more complicated caves. What if the cave had two three-way intersections? The cave could look like two circular loops connected by a passageway (as illustrated on the left of Figure 13), or it could be one circular loop with a passageway cutting across

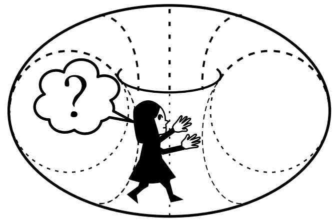

FIGURE 11. A cave which looks like the inside of a doughnut.

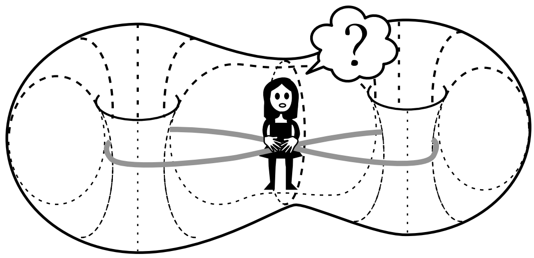

FIGURE 12. A cave with one four-way intersection.

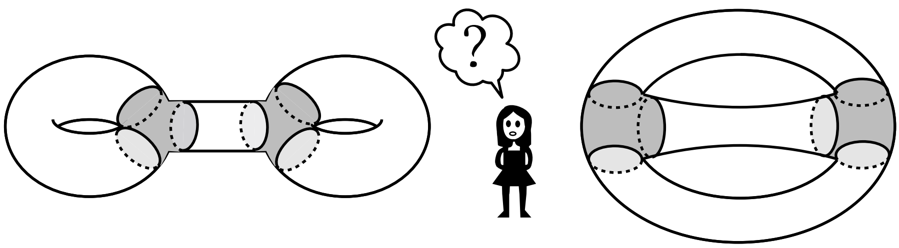

FIGURE 13. These caves each have two three-way intersections.

the middle (as illustrated on the right of Figure 13). How would she know which one it was? Could she use a ball of string to help figure it out?

After some thought, A. 3D-Girl decides that she could distinguish these two caves by seeing if there is a circular path from one three-way intersection back to itself that doesn't pass near the other three-way intersection, as there would be in the cave on the left in Figure 14. There would be no such path if she were in the cave on the right.

But what if she were in a cave that had many different types of intersections, and included some branches which seemed to go on forever and others which looped back on themselves? The thought of listing all the possibilities and coming up with a plan to distinguish them makes her head spin so much that she passes out. Luckily, the next time she wakes up she is actually back in her own bed. But her night adventures leave her wondering what could help her answer these sorts of questions. Perhaps if she checks back with A. Square, she can get some ideas from him.

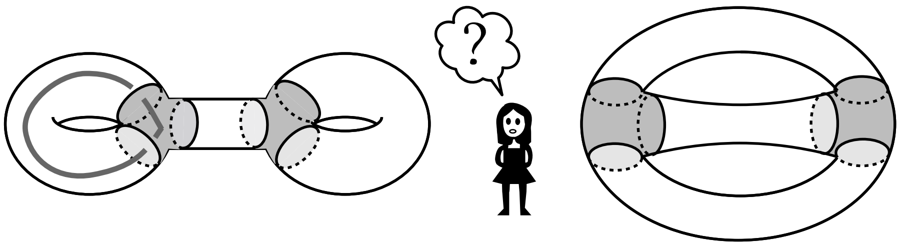

FIGURE 14. On the left there's a path from a three-way intersection back to itself that doesn't go near the other three-way intersection.

5. Creating Universes by Gluing

A. Square was able to use the notion of straightest paths to distinguish between a plane (denoted by $\mathbb{R}^2$), a sphere (denoted by S^2), and a torus (denoted by T^2). Note that we use the exponent 2 whenever we want to indicate that a space is 2-dimensional. This is in contrast with ordinary 3-dimensional space which is denoted by $\mathbb{R}^3$. The terms "ball" and "sphere", and "disk" and "circle" are often used interchangeably in normal life. But when talking about spaces of different dimensions, we want to use these terms more precisely. A sphere S^2 is the 2-dimensional surface that is the boundary of a 3-dimensional ball B^3. A circle S^1 is the 1-dimensional boundary of a 2-dimensional disk D^2.

We will now introduce the idea of *gluings* and *instant transport* to give A. Square a way to visualize some of the possibilities for his universe. Later, we will use similar techniques to enable us to visualize some analogous 3-dimensional universes. We begin by trying to help A. Square visualize a sphere. Drawing the usual picture of a sphere (like the one on the right side of Figure 10) won't help A. Square at all, since he doesn't have the 3-dimensional perspective necessary to understand that the dotted arcs are behind the rest of the picture. Instead we want to break up the sphere into pieces that A. Square can understand.

As 3-dimensional creatures we know that a sphere can be obtained by gluing two hemispheres together along their boundaries. To draw a picture of two hemispheres being glued together, we put arrows on the boundaries of each hemisphere to indicate how they should be matched up (as illustrated in Figure 15). In particular, the head of the arrow on one hemisphere is glued to the head of the arrow on the other hemisphere. A diagram which uses arrows to indicate how the boundaries of shapes are glued together is called a *gluing diagram.*

We can think of each hemisphere as being a disk that has been curved, like a disk made of play dough that you mold around half of a ball or even around your fist. Of course, A. Square cannot visualize disks being curved around a 3-dimensional object, because that would require him to have 3-dimensional vision. However, he can draw two flat disks with arrows on their boundaries. In order to help him understand that the boundaries of the two disks are glued together, we explain to him that when he arrives at the boundary of one disk, he will be instantly transported to the boundary of the other disk (see Figure 16). We show him the gluing diagram in Figure 16 to help him understand. A. Square reads a lot

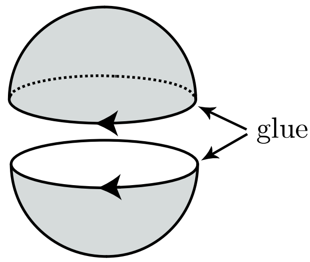

FIGURE 15. A sphere consists of two hemispheres that are glued together along their boundaries.

of science fiction so he accepts the idea that instant transport can take him from one disk to the other.

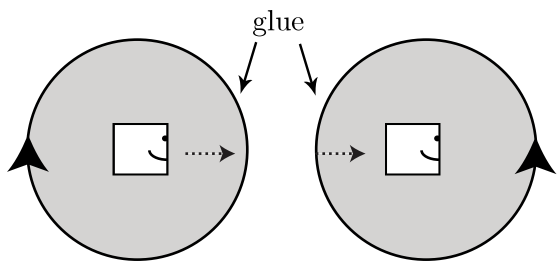

FIGURE 16. When A. Square arrives at the boundary of one disk, he is instantly transported to the boundary of the other.

Another way to think about gluing diagrams is as pages of a road atlas. When you drive off the right side of the map on one page of an atlas, there is a number or a letter to indicate that you drive onto the left side of a map on another page. If there were a road that went all the way around the earth, you could follow the pages of the atlas going from the map on one page to a map on another page and another and another, until eventually you would come back to the page where you started.

Now that A. Square can visualize a sphere as two disks glued together, we would like to use a gluing diagram to help him visualize a torus. Let's consider a 2-dimensional universe consisting of a square where opposite sides of the square are glued together. We put arrows on the sides of the square to indicate how they are attached. The head of the single arrow is glued to the head of the other single arrow, and the head of the double arrow is glued to the head of the other double arrow, as illustrated on the left side of Figure 17.

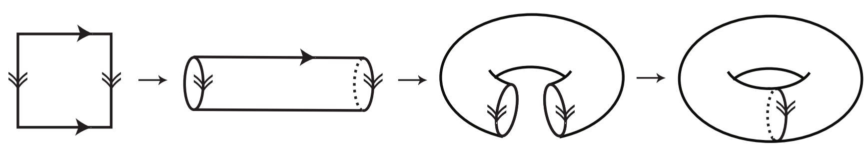

FIGURE 17. Gluing up opposite sides of a square.

As 3-dimensional people, we can physically glue up opposite sides of the square (if the square is made of a sufficiently flexible material). After gluing together one pair of sides we get a cylinder. Then if we glue together the double arrows on the ends of the cylinder, we get a torus as illustrated on the right side of Figure 17. You might want to check for yourself that it doesn't matter which pair of sides we glue together first—we'll still end up with a torus after both pairs of sides are glued.

A. Square can't visualize gluing up the sides of the square. So we have to explain to A. Square that if he goes through the top edge of the square, he is instantly transported to the bottom edge, and if he goes through one side, he is instantly transported to the other side (see Figure 18). It's just like a video game where a character who goes off the screen on one side reappears on the opposite side. When A. Square is not reading science fiction, he likes to play video games. So this explanation also makes sense to him.

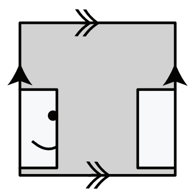

FIGURE 18. When A. Square exits from the right edge, he returns on the left edge.

If we want to physically glue up opposite sides of a square, the square needs to be made out of a flexible material. In particular, if the square is made of paper, we can only glue together one pair of opposite sides at a time (you should try this to see for yourself what happens). In fact, there is no way to make a nice round torus out of paper. For this reason, it is more convenient for us to think of the torus the way A. Square does, as an *abstractly* glued up square. We use the word "abstractly" here to mean that we're thinking about it in terms of instant transport rather than as a square with sides that are physically glued together. When we think of a torus this way, we call it a *flat torus*; whereas when we imagine a torus as the surface of a doughnut or a bagel, we call it a *curved torus*.

A flat torus is a "better" universe than a curved torus because a flat torus is the same everywhere. For example, if A. Square watches his friend B. Triangle move around on a flat torus, she will always have the same area even if she goes out one edge and comes in another (as illustrated on the left side of Figure 19). This is not the case if she moves around on a curved torus. For example, on the right side of Figure 19 we see that B. Triangle might get a stomach ache if she eats a big dinner at a restaurant near the outer circle of a torus and then walks towards the inner circle. It wouldn't be very pleasant to live in a place where you would get a stomach ache every time you walked in a certain neighborhood.

6. Games on a Gluing Diagram for a Cylinder and a Torus

One way to check your understanding of abstract gluings is to play games on a gluing diagram. For example, you can play Connect Four on a gluing diagram for a cylinder (see Figure 20). The normal game of Connect Four is like tic-tac-toe but

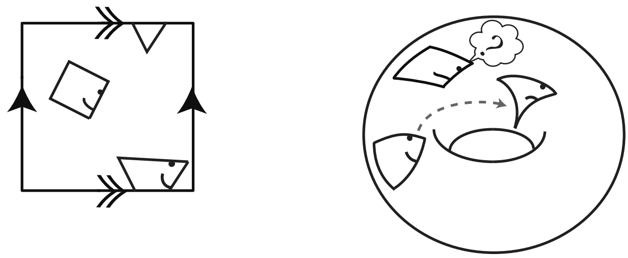

Figure 19. B. Triangle changes her shape when she moves around on a curved torus.

with gravity. One player has black tokens and the other player has grey tokens. The players take turns dropping their tokens into a vertical plastic grid which has seven columns and six rows. As you drop in a token, it falls down due to gravity. The game ends when one player has four in a row either horizontally, vertically, or diagonally.

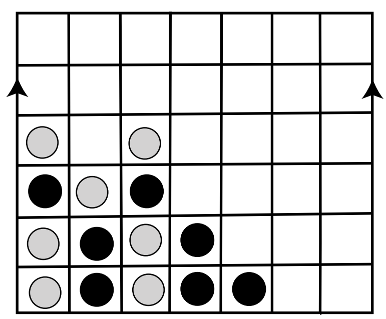

Figure 20. The game of Connect Four on a cylinder.

While normal Connect Four is played in a vertical rectangular grid, we can also play Connect Four on such a grid where the pair of vertical sides have been glued together as indicated by the arrows in Figure 20. This gives us a cylindrical grid. There still is gravity which causes the tokens to drop down to the bottom of the cylindrical grid as illustrated in Figure 20. As with normal Connect Four, you win if you are the first player to get four in a row horizontally, vertically, or diagonally. For example, consider the game in progress in Figure 20. It's Grey's turn, and Grey can win in one move. Can you see how?

You should find a partner and play Connect Four a few times on the gluing diagram for a cylinder. In 1988, Victor Allis "solved" Connect Four in the sense that he determined a strategy wherein the first player can always win. How might your strategy for Connect Four on a cylinder be different from your strategy in a normal Connect Four game?

We can also play games on a flat torus. In Figure 21, we see a game of tic-tac-toe on a torus where **X** has won. This game would be much harder to draw (and to think about) on a curved torus. You should stop reading now to play a game of torus tic-tac-toe with a friend.

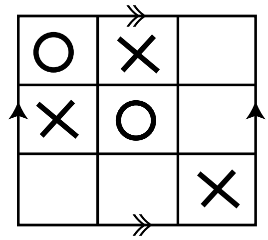

FIGURE 21. **X** has won this game of torus tic-tac-toe.

There are many other games that you can play on a flat torus. Could you play Connect Four on a flat torus? You should check out the website

`http://www.geometrygames.org/TorusGames/`

where you can play tic-tac-toe, mazes, crossword puzzles, word search puzzles, jigsaw puzzles, chess, pool, and gomoku all on a flat torus.

7. Extended Diagrams

One way to plan your strategy when playing games on a gluing diagram is to draw an *extended diagram* to help you see where you go when you go off one side of the board. In order to construct an extended diagram for the flat torus, let's first imagine that A. Square lives all by himself in a flat torus. To keep track of his experience, let's imagine that curtains are hung at the edges of the gluing diagram. This gives him the impression that his neighborhood is a square, but he can walk through the curtains. He puts on his hat to go out for a walk. When he walks through the curtains on the right side of the gluing diagram he is instantly transported to the left side. However, he thinks he has entered a second square neighborhood where there is another guy with a hat just like his, who is also taking a walk at the same speed that he is, and that guy is going through the curtain on the right just as A. Square enters through the curtain on the left (see Figure 22). In fact, the square that he sees taking a walk ahead of him is really just A. Square himself.

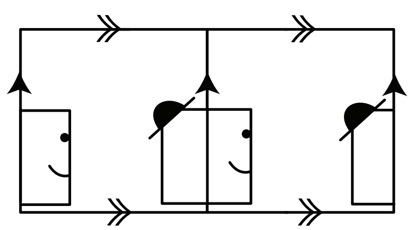

FIGURE 22. A. Square sees a copy of himself up ahead.

A. Square continues his walk in the same direction, believing he is leaving the second square and entering a third square, and so on. Continuing in this fashion, he believes that there are infinitely many squares in front of him, one after another. Similarly, he believes that there are infinitely many squares on top of him stacked one above the other, and infinitely many squares going off on a diagonal. Altogether,

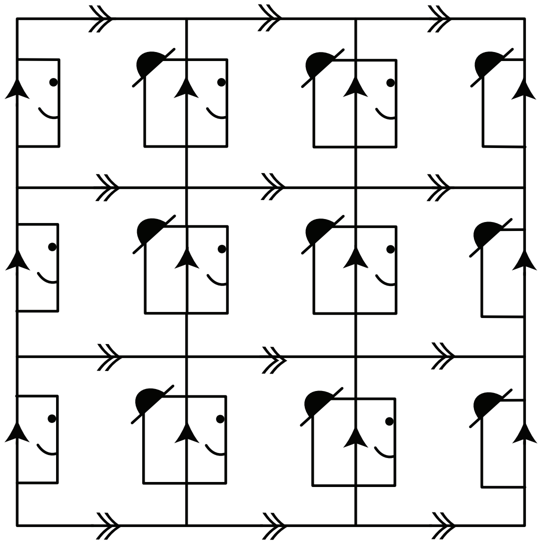

Figure 23. A portion of an extended diagram of a flat torus with A. Square in it.

he imagines there is an entire plane full of squares glued together left to right and top to bottom, containing other creatures who look just like him, wearing hats like his, and moving at exactly the same rate as he is.

The space he imagines he lives in is called the *extended diagram* of the flat torus. We see a portion of the extended diagram of the flat torus in Figure 23. Keep in mind that all of the square gluing diagrams in the extended diagram are really just a single flat torus, and all of the guys with hats on are really just A. Square himself.

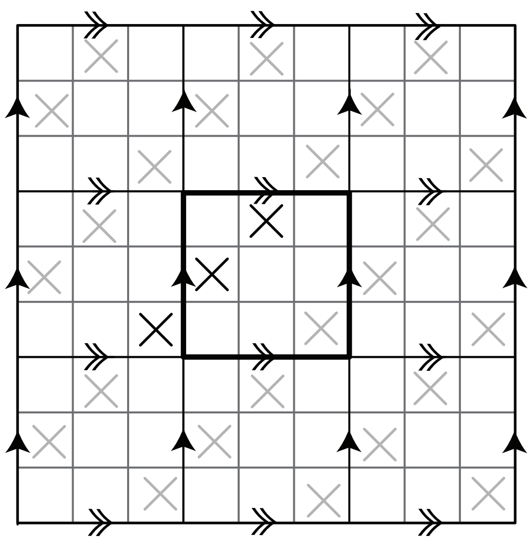

Figure 24. The three black X's in an extended diagram represent three in a row in flat torus tic-tac-toe.

It's a bit like when you are in a dressing room with opposing mirrors and it seems like there are infinitely many copies of you, though one is blocking another so you can't really see them all. You know that really you are the only you, and the copies of you in the mirror are just your images.

If we draw an extended diagram of a flat torus tic-tac-toe board, then it is easier to see all the possible ways you can get three in a row. For example, the row of X's that we saw in Figure 21 looks like an ordinary three in a row in the extended diagram (see the three black X's in Figure 24). In fact, there are infinitely many copies of these three X's in the extended diagram and we could choose any three adjacent ones to represent the 3 in a row on the tic-tac-toe board from Figure 21.

You should try to see if an extended diagram can help you find a winning move for grey in the game of Connect Four illustrated in Figure 20. Note that the extended diagram for a cylinder will look like a strip of paper that goes on indefinitely to the left and the right but not on the top and the bottom.

8. Introducing the 3-Dimensional Torus and Sphere

Now that we have some experience helping A. Square understand some possibilities for his 2-dimensional universe, let's consider some analogous possibilities for our 3-dimensional universe. It is natural for us to assume that our universe goes off infinitely in all directions, just as it was natural for A. Square to assume that his universe was an infinite plane. We saw that while it is easy for us to think of many different possibilities for A. Square's universe, it is hard for him to visualize these possibilities. We used the idea of gluing up disks and squares in order to help him understand what it would be like if his universe were a sphere or a torus. This gives us the idea that our own 3-dimensional universe might be an analogously glued up 3-dimensional shape.

We begin by exploring a 3-dimensional version of the flat torus. Suppose that you live in a large cubical room with pairs of opposite walls glued together and the floor and the ceiling glued together. We imagine this space "abstractly" just as A. Square imagines a flat torus. We don't assume that the room is flexible and exists in a higher dimension where opposite walls curve around so that they can be glued together. Rather, we imagine that a person is instantly transported from one wall to the same point on the opposite wall. This 3-dimensional universe is called a 3-*torus* or a 3-dimensional torus and is denoted by T^3.

In order to help us visualize this, let's imagine that A. 3D-Girl lives in a 3-torus. If she floats through the ceiling, she comes back through the floor (see Figure 25). This is similar to what we saw in Figure 18 when A. Square travelled through the right edge of the square and reappeared on the left edge.

Even though we visualize the 3-torus as a glued up cube, there is no boundary to this space. We never have to worry about walking "off the edge" or bumping into a wall. Figure 26 illustrates a gluing diagram for a 3-torus where pairs of identical arrows are used to indicate that opposite walls are glued together. Notice that the head of an arrow is glued to the head of the similar arrow with the colored part glued to the colored part and the white part glued to the white part. We use three different types of arrows because there are three pairs of opposite walls that are glued together.

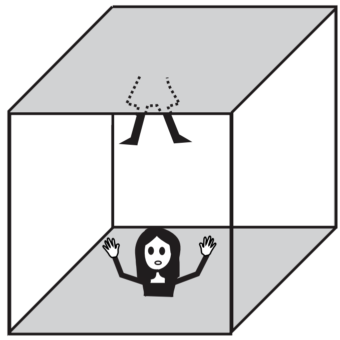

FIGURE 25. If we glue up opposite faces of a cube, we get a 3-torus.

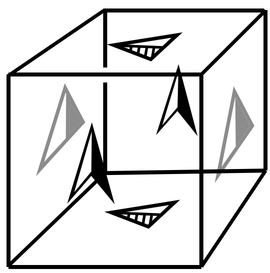

FIGURE 26. The gluing diagram of a 3-torus.

A very, very large version of the 3-torus could actually be the shape of our universe. If this were the case and we had a spaceship with a super-duper hyper-warp drive that allows it to zoom off as fast as we want and as far as we want, then we could send it off in one direction and it would eventually come back to Earth from the opposite direction. In reality, though, such a spaceship would probably take billions of years to return to Earth. In fact, if our universe were a 3-torus and we had a really powerful telescope, we could see a copy of our own Earth (though the image we would see could be billions of years old since it would take the light billions of years to traverse the entire universe to get back to us).

Another possibility for the shape of our universe is a 3-*dimensional sphere* (also known as a 3-*sphere*), which is denoted by S^3. It is very important to remember that S^3 is NOT a solid 3-dimensional ball (which we denote by B^3). Rather S^3 is the 3-dimensional analog of a 2-dimensional sphere S^2. The ball B^3 cannot be a universe because it has a boundary, and we are assuming that no universe can have a boundary.

There are several ways to visualize S^3, each of which is analogous to an approach that A. Square could use to visualize S^2. The first way to think about S^3 is to formally define it in terms of coordinates using an equation. With this approach, we would define

$$S^2 = \{(x, y, z) | x^2 + y^2 + z^2 = 1\}$$

as a subset of $\mathbb{R}^3$. Keep in mind that S^2 is 2-dimensional even though it's a subset of $\mathbb{R}^3$. This definition of S^2 makes it easy for 3-dimensional creatures to picture S^2 as the outside of a round ball with radius 1. By analogy, we could define S^3 in terms of coordinates as

$$S^3 = \{(x, y, z, w)|x^2 + y^2 + z^2 + w^2 = 1\}$$

as a subset of $\mathbb{R}^4$. However, this doesn't give us much intuition about how to visualize S^3 since it's hard enough to picture $\mathbb{R}^4$. A. Square has the same problem if he tries to picture $S^2 = \{(x, y, z)|x^2 + y^2 + z^2 = 1\}$ as a subset of $\mathbb{R}^3$.

For a more visual approach, recall that we helped A. Square visualize a sphere by telling him that a sphere is made of two disks which have been glued together along their boundaries (see Figure 16). In order to construct S^3 in an analogous way, we need to first find the 3-dimensional analogue of a disk. One way of defining a *disk* is as the set of points in $\mathbb{R}^2$ whose distance from the origin is 1. We would express this symbolically as

$$D^2 = \{(x, y)|x^2 + y^2 \leq 1\}$$

as a subset of $\mathbb{R}^2$. Thus the 3-dimensional analogue of a disk is the set of points in $\mathbb{R}^3$ whose distance from the origin is 1. We express this symbolically as

$$B^3 = \{(x, y, z)|x^2 + y^2 + z^2 \leq 1\}$$

as a subset of $\mathbb{R}^3$. Hence the analogue of a disk is just a solid *ball.* Now the 3-dimensional analogue of constructing S^2 by gluing two disks together along their boundaries is constructing S^3 by gluing two solid balls together along their boundaries. If A. 3D-Girl lives in a 3-sphere, then she can travel through the boundary of one of the solid balls into the other solid ball as illustrated in Figure 27.

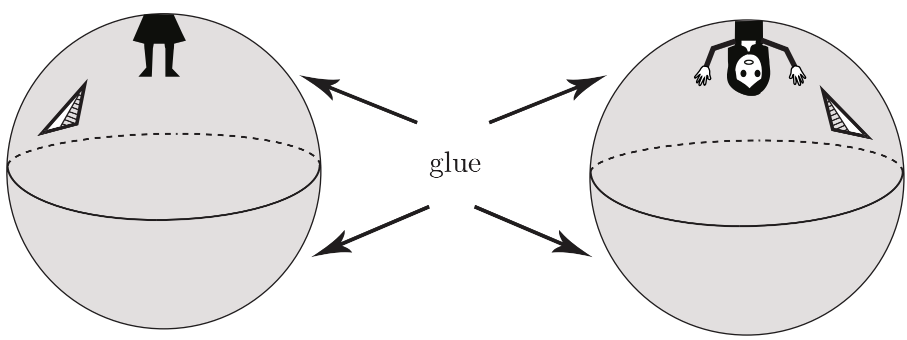

FIGURE 27. S^3 consists of two balls that are glued together along their boundaries.

No matter how we imagine contorting or smashing the balls in Figure 27, it is impossible for us to visualize the boundaries of the balls actually being glued together. This is because we would need four dimensions in order to do the gluing. This is the same problem that A. Square would have if he tried to visualize gluing together the boundaries of the two disks in Figure 16. The best way for us to imagine the gluing is to think about it point by point, so that each point on the boundary of one ball is glued to the corresponding point on the boundary of the other ball. The gluing arrows on the spheres can help us remember how one ball is glued to the other.

Now we know that in addition to $\mathbb{R}^3$, our universe could be a 3-torus T^3 or a 3-sphere S^3. As soon as we discover new possibilities for our universe, we want to know how we could determine whether our universe has one of these forms. If we could somehow map out every cubic inch of our universe, we would know that our universe was finite. So we could conclude that it wasn't $\mathbb{R}^3$. But how could we distinguish between T^3 and S^3? As usual, it's a good idea to start by answering the analogous question for 2-dimensional universes before we attempt to answer it for our own universe.

9. Distinguishing a 2-Sphere from a 2-Torus

In Figure 10, we saw that A. Square could figure out whether he was living in S^2 or T^2 by considering different straightest paths. In particular, he observed that two straightest paths can meet in exactly one point on T^2, but not on S^2.

Another way to distinguish between S^2 and T^2 is to consider what would happen if a wall were built along a straightest path. For example, on the left side of Figure 28 we see that A. Square is in the northern hemisphere of S^2 and B. Triangle is in the southern hemisphere. The Great Circular Wall which was built along the equator divides S^2 into two pieces preventing A. Square from visiting B. Triangle. On the other hand, as we see on the right side of Figure 28, a wall along a straightest path on T^2 might make it inconvenient for A. Square to visit B. Triangle, but if he's willing to walk around the long way he will eventually get there.

Of course, walls can be built on a torus which separate it into two regions as we see in Figure 29. However, such a wall does not follow a straightest path.

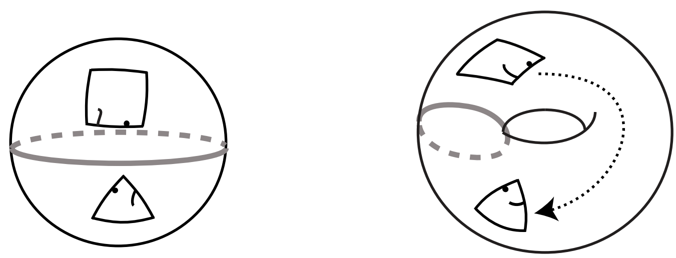

FIGURE 28. A wall along a straightest path in S^2 divides the space into two regions, but it doesn't in T^2.

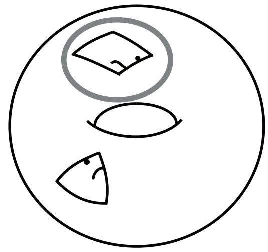

FIGURE 29. This wall is not along a straightest path.

In order to explain to A. Square how he can use walls along straightest paths to distinguish between a torus and a sphere, we need to use gluing diagrams rather than the illustrations in Figure 28, since he can't understand 3-dimensional illustrations. Also, it's easier to draw straightest paths in the gluing diagram than on the curved torus, since straightest paths in the gluing diagram just consist of line segments. To show him a wall along a straightest path in T^2, we draw Figure 30. It's then easy for him to check that the wall doesn't divide T^2 into two separate regions.

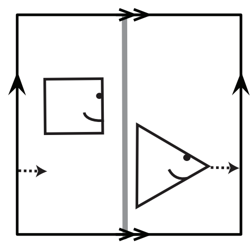

FIGURE 30. This wall doesn't separate T^2 into two regions.

Next we want to show A. Square that any wall along a straightest path in S^2 will separate it into two regions. The problem is it's not so easy to draw a straightest path in the gluing diagram for S^2, because the gluing diagram for S^2 flattens the hemispheres into disks which distorts distances. This means that if we draw a line segment in one of the disks, it won't necessarily correspond to a great circle on S^2. Instead, let's start out with a great circle which goes between A. Square and B. Triangle (illustrated as a thick black circle in the left image of Figure 31). This great circle will be our wall. To make it easier to draw this circle in the gluing diagram, we divide the sphere into two hemispheres so that half of the black circle is in each hemisphere, as illustrated in the second image of Figure 31. Then we flatten out the hemispheres into disks so that the black semi-circles are now diameters of the disks, as illustrated in the right image of Figure 31. We only show A. Square the picture of two disks on the right, since the other pictures in the figure wouldn't make sense to him anyway. We then explain to A. Square that even if he goes from one disk to the other (as indicated by the dotted lines and arrows) he will still be on the opposite side of the wall from B. Triangle.

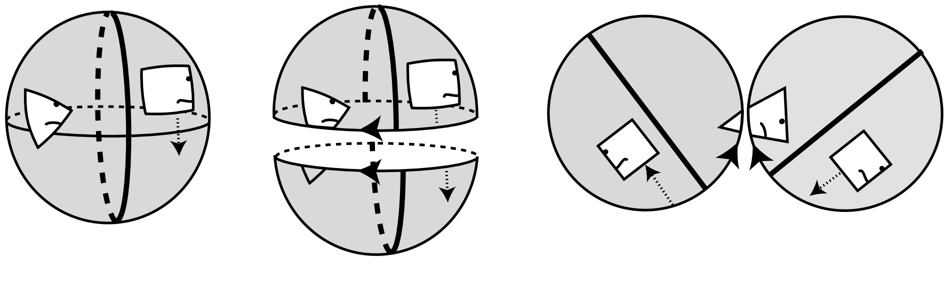

FIGURE 31. This wall separates S^2 into two regions.

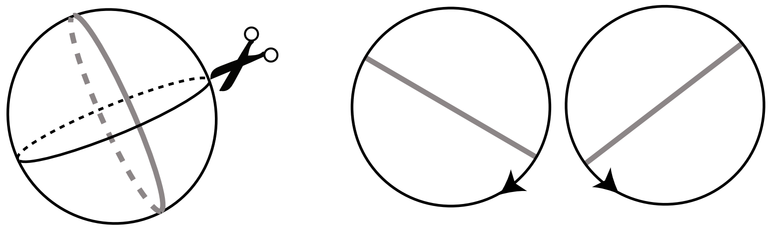

FIGURE 32. The grey circle will become a diameter in each disk if we cut along this great circle.

You may be concerned that we have only shown A. Square that this particular wall separates a sphere into two regions. But given any great circle, we can cut the sphere into hemispheres so that half of the great circle is in each hemisphere, and hence each half of a great circle becomes a diameter of one of the disks. For example, if we want to build a wall along the great circle which is grey in Figure 32, we can cut the sphere into hemispheres along the black circle, giving us the same picture of disks that we had in Figure 31. So we have in fact shown that any wall along a straightest path on the sphere separates it into two regions.

10. Distinguishing a 3-Sphere from a 3-Torus

Now let's see if we can use walls along straightest paths to distinguish between T^3 and S^3. Since we are inhabitants of our space, we have to use gluing diagrams, just as we did for A. Square. To build the wall, we lay a row of bricks along a straightest path, expanding the wall by laying more bricks up, down, left, and right, just as you would if you were building a brick wall to keep deer out of your vegetable garden. Recall that for the walls we built in T^2 and S^2, one side of the wall met the other side so the wall was circular. Since we now have one more dimension, not only do the sides of the wall meet, but the top and bottom of the wall also meet.

Figure 33 illustrates such a wall in T^3. We didn't put the arrows on opposite faces of the cube, but you should remember that opposite sides of the cube are glued together. This means that the top of the wall is indeed glued to the bottom of the wall and the left side of the wall is glued to the right side of the wall. If A. 3D-Girl walks with her back to the wall, she can go visit her friend on the other side of the wall. Thus this wall does not separate T^3 into two regions, just as the wall in Figure 30 didn't separate T^2 into two regions.

We have to be a little more careful with the gluing diagram for S^3, since the two balls have been "flattened" to fit in three dimensions just as the two hemispheres of S^2 had to be flattened to fit in two dimensions. Recall that we saw in Figures 31 and 32 that given any wall along a straightest path in S^2, we can divide S^2 into two disks so that the wall cuts each disk exactly in half. By analogy we now see that given any wall along a straightest path in S^3, we can divide S^3 into two balls so that the wall splits each ball exactly in half. In Figure 34 we illustrate a wall along a straightest path in S^3 which splits each ball exactly in half. As we can see in the figure, even if A. 3D-Girl walks away from the wall into the next ball, she

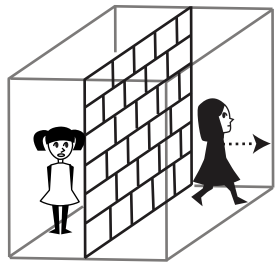

FIGURE 33. A. 3D-Girl is going to visit her friend on the other side of this wall in T^3.

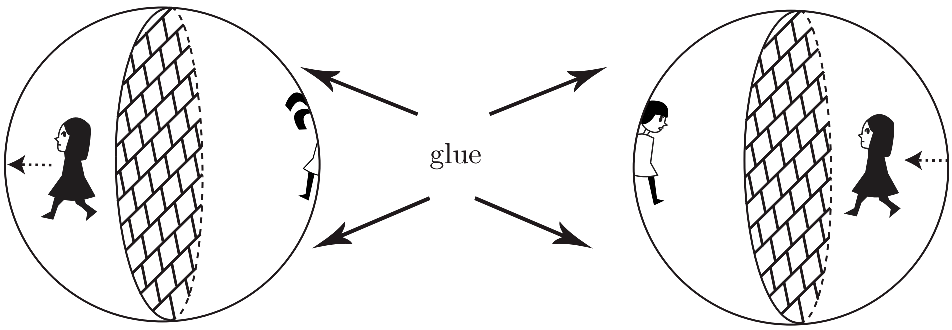

FIGURE 34. No matter where A. 3D-Girl goes, she can't visit her friend on the other side of the wall.

will remain on the same side of the wall. In fact, no matter where she goes in S^3, her friend will always be on the other side of the wall.

Thinking about 3-dimensional universes like T^3 and S^3 reminds A. 3D-Girl about her nightmare with the caves, and she starts to worry that she'll have the same nightmare when she goes to sleep later. She has a momentary panic about losing her ball of string, and forces herself to take a few deep breaths. Once she's a bit more calm, she starts thinking about whether she could build walls along straightest paths to distinguish different caves.

Recall that in Section 4, A. 3D-Girl was faced with the problem of distinguishing between the two caves in Figure 35. To solve this problem, she had figured out that for the cave on the left she could lay out a loop of string from one three-way intersection back to itself that didn't go near the other three-way intersection. This wouldn't be possible for the cave on the right of Figure 35.

Now she realizes that she could also distinguish between the two caves in Figure 35 by building walls along straightest paths. In particular, for the cave on the left in Figure 36, she could build a wall that separates it into two doughnut-shaped regions which she couldn't do for the cave on the right.

A. 3D-Girl reasons that even if she didn't want to go to the trouble of building a complete wall, she could use rocks to mark where the wall should go. Knowing two different techniques to distinguish the caves reassures A. 3D-Girl enough that

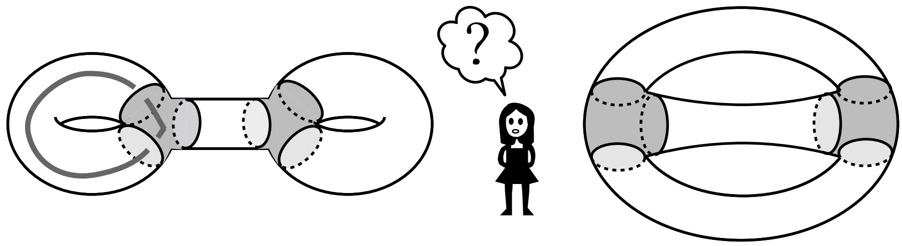

FIGURE 35. On the left, there's a path from a three-way intersection back to itself that doesn't go near the other three-way intersection.

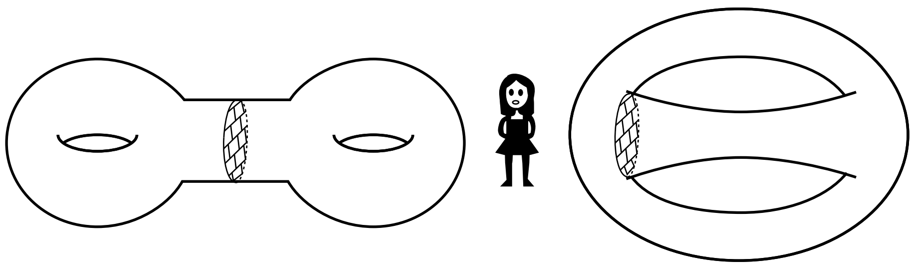

FIGURE 36. The wall on the left separates the cave into two doughnut-shaped regions.

she can go to sleep, during which has the most wonderful, peaceful, blissful sleep she's had since her voyage to Earth 2.

11. Exercises

1. Suppose A. Square lives alone in a small flat torus and each of his sides is a different color. Imagine that he can see in all directions at once. Draw an *extended diagram* representing how A. Square might imagine his universe. Explain your drawing.
2. Suppose A. Cube lives alone in a small 3-torus and each of her sides is a different color. Imagine that she can see in all directions at once and has depth perception. Draw an extended diagram of A. Cube's universe. Explain your drawing.
3. Suppose A. Square lives alone in a small sphere. What direction would he look if he wanted to see the top of his head?
4. How could A. Square detect the difference between a plane, an infinite cylinder, and an infinite cone? Is an infinite cone or an infinite cylinder a reasonable 2-dimensional universe for A. Square to live in? Why or why not?
5. How do we create an infinite cylinder and an infinite cone by gluing up a planar surface? What could be the 3-dimensional analogue of an infinite cylinder? What about an infinite 3-dimensional cone?

6. Suppose we tell A. Square that his universe is either a flat torus, a sphere, or a plane. What information would help him to determine which it really is?
7. Suppose we were told that our universe was either a 3-torus or $\mathbb{R}^3$. What information would we look for to determine which it really was?
8. Draw a circular path on a torus which intersects both the thick black and the thick grey circle in Figure 10 in just one point. Draw a second circular path also with this property, but which can't be deformed on the torus to the first one you drew. Your circular paths do not have to be straightest paths.
9. Draw a circular path on the torus in Figure 10 which intersects the thick black curve in two points and intersects the thick grey curve in three points.
10. In the Connect Four game on a cylinder illustrated in Figure 37, can either player win in one move if that player goes next?

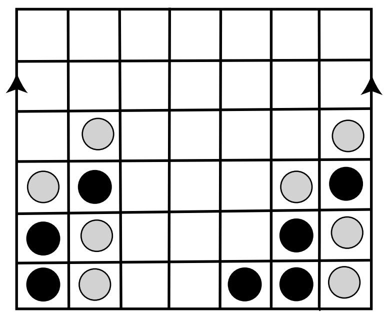

FIGURE 37. Illustration for Exercise 10.

11. In tic-tac-toe, we say two moves are *equivalent* if the strategy for the rest of the game after one of the moves is analogous to the strategy for the rest of the game after the other move. For example, in regular tic-tac-toe, a first move in one corner is equivalent to a first move in any of the other three corners; and a first move in the middle of one side is equivalent to a first move in the middle of any side. Thus in regular tic-tac-toe, there are three inequivalent moves for the first player: the center, a corner, or the middle of a side.

    How many inequivalent first moves are there for tic-tac-toe on a cylinder (see Figure 38).

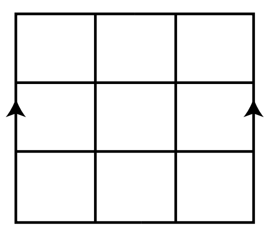

FIGURE 38. Illustration for Exercise 11.

12. In torus tic-tac-toe how many inequivalent moves are there for the first player?

13. In torus tic-tac-toe, how many inequivalent moves are there for the second player?
14. Is either player guaranteed to win torus tic-tac-toe if both players use optimal strategies? Justify your conclusion.
15. Is it possible to have a tie in torus tic-tac-toe if both players wanted to? Justify your conclusion.
16. If you play torus chess, with the pieces set in the usual starting positions, explain how the first player can win in just one move. Now, design a better starting position that does not lead to either checkmate or a stalemate in just one move.
17. The wall built along a straightest path in Figure 30 does not divide T^2 into two regions. Suppose you built two walls in T^2 both along straightest paths. How many regions would the walls divide T^2 into? (Hint: Your answer may depend on whether the walls intersect or not.)
18. It is possible to build three intersecting walls along straightest paths in T^3 that still do not separate T^3 into different regions. Explain or draw a picture showing how this can be done.
19. Explain how A. 3D-Girl can use walls to distinguish between a doughnut-shaped cave (see Figure 11) and a cave with one four-way intersection (see Figure 12).
20. In Figure 32, we started with a wall along the thick black circle on the sphere, then cut along a perpendicular equator to get the gluing diagram on the right. What would the thick black circle look like in a gluing diagram if we cut along an equator that was *not* perpendicular to it? Draw a sketch and explain.