DEMENTED DIMENSIONS

by

ANDERSON HASSELL NORTON III

(Under the Direction of Theodore Shifrin)

ABSTRACT

This paper examines some of the historical problems that mathematicians have faced in defining dimension. Several definitions are considered and several mathematical "monsters" are introduced, such as fractals and space-filling curves, which challenge previously developed definitions and suggest new definitions. Space-filling curves introduce particularly interesting challenges because, counter to intuition, they actually fill *too much* of the target space to establish homeomorphisms. The paper picks up with a comparison of two classes of constructions for space-filling curves and the question of one-to-oneness: How much too much do space-filling curves fill? Both types space-filling constructions relate to fractals, which, in turn, lead us to consider fractal dimensions: self-similarity dimension, Hausdorff dimension, box-counting dimension.

INDEX WORDS: Dimension, Fractal, Space-filling

DEMENTED DIMENSIONS

by

ANDERSON H. NORTON III

B.S., Virginia Polytechnic Institute & State University, 1995

M.A., The University of Georgia, 1999

A Thesis Submitted to the Graduate Faculty of The University of Georgia in Partial Fulfillment

of the Requirements for the Degree

MASTERS OF ARTS

ATHENS, GEORGIA

2004

DEMENTED DIMENSIONS

by

ANDERSON HASSELL NORTON III

Thesis Advisor:	Theodore Shifrin
Committee:	Malcolm Adams
	Edward Azoff
	Roy Smith

Electronic Version Approved:

Maureen Grasso

Dean of the Graduate School

The University of Georgia

May 2004

ACKNOWLEDGEMENTS

Many thanks to my teachers over the last several years in the Mathematics Department at The University of Georgia. Their enthusiasm for the mathematical ideas of their students is inspiring to me as a student and a teacher. Among my influential teachers, four also agreed to work with me on this thesis. I would not have written a thesis if I were not already convinced they would help me to make it fun, and each of Roy Smith, Ed Azoff, and Malcolm Adams has contributed uniquely to make the entire process enjoyable. Thanks to my advisor, Ted Shifrin, for years of guidance and friendship. And, finally, I would like to recognize Paul Oser, another wonderful teacher and a colleague from the math department at Georgia's Governor's Honors Program who loaned to me the title of this thesis.

TABLE OF CONTENTS

CHAPTER 1: PROLOGUE

This paper is intended as a set of course notes for a mathematical excursion in dimension. Specifically, the notes challenge the reader to define and redefine dimension while considering the topological and analytical implications of each definition and responding to properties of particularly peculiar mathematical objects. This approach supposes and requires some foreknowledge of calculus, real analysis and topology; a course in Point-Set Topology or Real Analysis should more than adequately serve as a prerequisite.

The debate over appropriate definitions of dimension is a general theme underlying the mathematical development of these notes. In fact, many of the challenges and subsequent solutions (including the introduction of new terms and concepts) reflect historical debates in which mathematicians like Brouwer, Hilbert, Cantor, Cauchy, and Peano were involved.

CHAPTER 2: A BRIEF HISTORY OF THE PROBLEM

This chapter deals with the history of problems mathematicians have faced in defining dimension. Although all students should have developed an intuitive definition of dimension upon entering college, a consistent definition that satisfies our intuitive requirements is difficult to make explicit. In this chapter, we will see examples of "monsters" that challenge naïve definitions while suggesting new paths in developing the comprehensive definition we seek. Ultimately, these paths will lead us to the study of a particularly interesting class of monsters, fractals.

Exercise 0: Assign a dimension to each of the objects indicated below.

Exercise 1: Define dimension (consider it as a first draft).

One obvious definition is inherited from an object's coordinates once it is given in a Euclidean space, $\Re^n$. We could say that an object has dimension n if it can be described by points with n coordinates. We can call this *Euclidean dimension*, but it may be problematic for determining the dimension of objects like that shown in example (1) of Figure 1. Though the object is just a copy of $\Re^1$, we would not know whether to assign it dimension 1 or 2 because it may be tilted relative to the axis depending on how we assign coordinates to the object. We could revise our definition to say that the smallest possible number of Euclidean coordinates gives the dimension, but example (2) may still be troublesome.

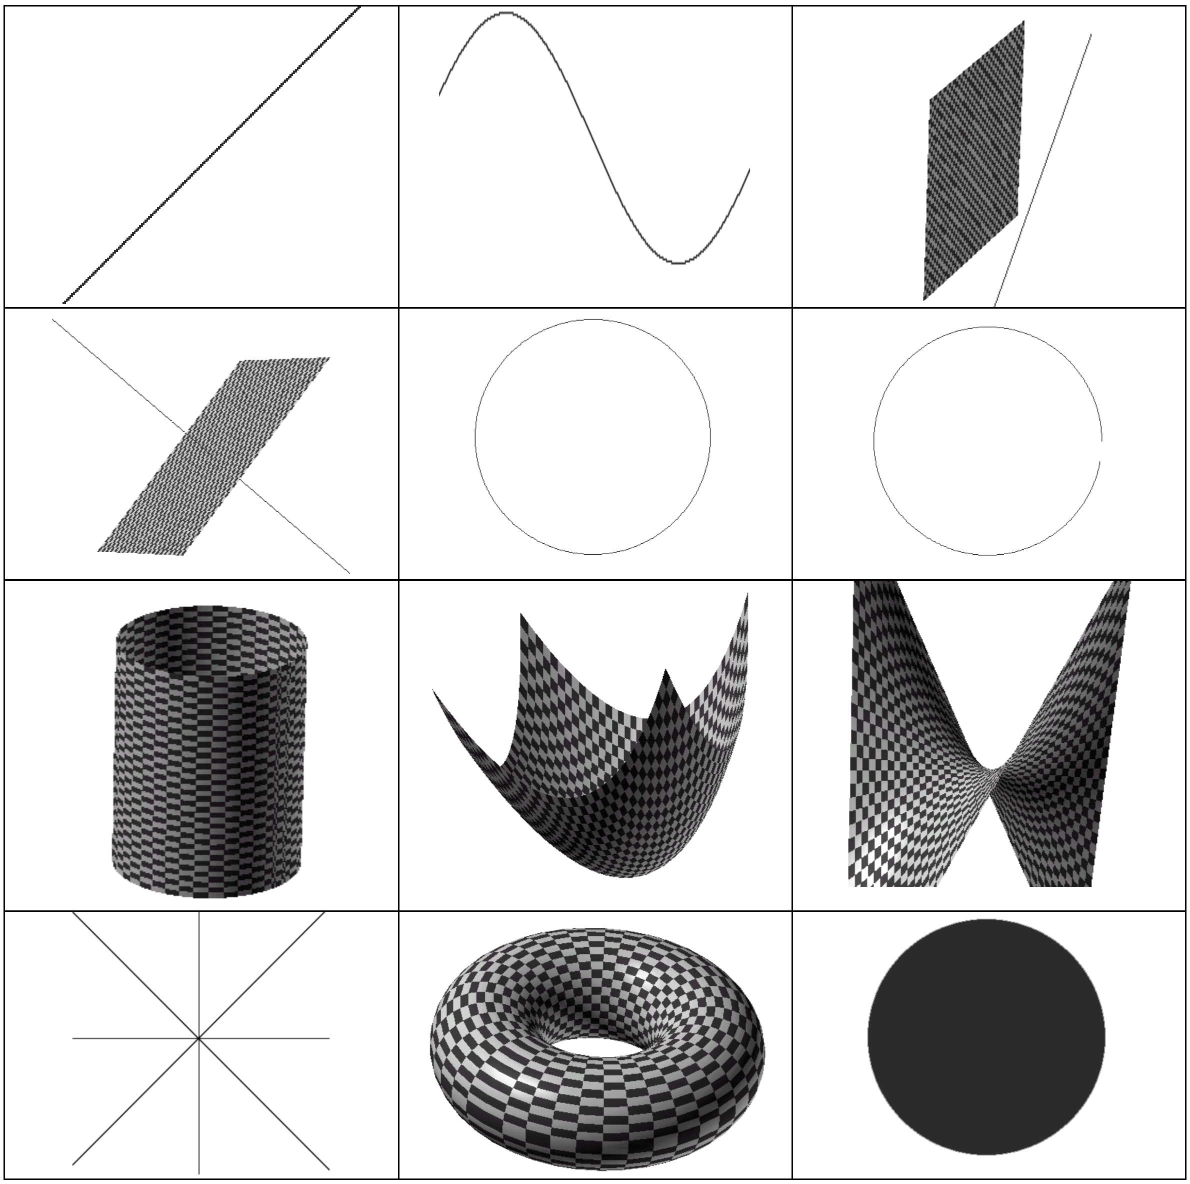

Figure 1. Examples of mathematical objects, numbered from left to right and top to bottom as examples 1 through 12.

There is a sense in which example (2) may need only one coordinate as well; the curve may be described by an equation with only one independent variable such as $y = \sin(x)$. Even general curves in space may be parametrized by a single independent variable, t [Figure 2], and surfaces such as the torus of example (11) can be parametrized by two independent variables, as we will see later in this section. This idea introduces an *independent variable dimension*, which agrees with the revised Euclidean dimension for graphs of linear functions and disagrees on graphs of most non-linear functions. However, both definitions appeal to algebra in defining dimension and neither will be very useful in more general settings where we have no explicit equations or coordinates for generating the object. Let us then consider a topological approach to dimension.

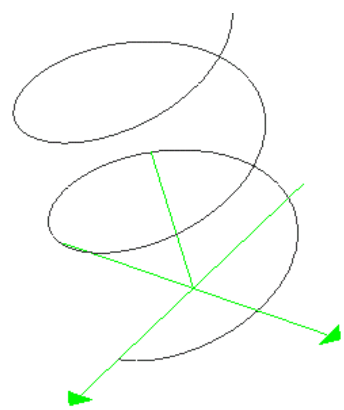

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} \cos t \\ \sin t \\ \frac{t}{4} \end{bmatrix}$$

Figure 2. Parametrized equation and graph of a helix.

While there is plenty of room for disagreement in assigning dimensions to the examples in Figure 1, we could all agree on a few canonical examples: A point ($\Re^0$), for instance, should have dimension 0 because it has no length, area, or volume. Likewise, a line ($\Re^1$) should have dimension 1 because it has length but no area or volume; a plane and space should have dimensions 2 and 3, respectively, for similar reasons. We could go on to spaces of higher dimensions, but we already have enough to be concerned about in dimensions no greater than 3. In fact, all of the objects we will study can be embedded in $\Re^3$.

$\Re^3$ provides us with a relatively tame ambient space in which to study embedded subsets. Most importantly, it is a metric space – we can unambiguously define distances between points in the space. Recall from courses in analysis or topology that a closed set is a subset of a space whose complement is open in that space, and that an open set of a metric space is a subset of that space such that each point in the subset has a neighborhood contained in that subset. Finally, a neighborhood about a given point in a metric space is the set of all points whose distance to the given point is less than some given nonnegative value.

Among other things, we know from real analysis (ref. Rudin) that subsets of $\Re^3$ are compact if and only if they are closed and bounded. The compact analogues of our canonical examples are the point at the origin (I^0), the unit segment (I^1), the unit square (I^2), and the unit cube (I^3); all of these objects are closed and bounded subsets of $\Re^3$. Moreover, they are all connected (not contained in a union of two disjoint, nonempty open sets). Though $\Re^3$ is our present playground for objects, we will not rely on coordinates in $\Re^3$ for developing a general definition of dimension.

We have already introduced a few topological properties and will be required to introduce more as further problems arise. We may require connectedness, for instance, in order to avoid the confusion of assigning dimensions to objects like example (3) shown in Figure 1. Should it have dimension 1 or 2? We would rather assign a dimension to each connected component; however, we might still be puzzled by example (4) in the same figure.

Connectedness also suggests a definition of dimension. Note that the boundary of I^n is a collection of copies of I^{n-1}, for each of our compact canonical examples. Likewise, for the unbounded canonical examples (Euclidean n-spaces), we see that each can be divided into two or more components by removing a subset that is a copy of a canonical example with dimension one less than itself. For instance, a plane can be divided into two components by removing any line in the plane. We can generalize this to define an object as having dimension 2 if it can be divided into two or more components by removing a line, but not by removing any point. Likewise, an object has dimension n if it can be divided into two or more components by a (linearly embedded) copy of $\Re^{n-1}$, but not by one of lower dimension. We will call this attempt at an explicit definition *Poincaré splitting dimension* because it is based on Poincaré's inductive approach to dimension (Poincaré, 1952, pp. 32-33). Alternatively, we could allow for the removal of a finite number of copies and call this definition *Poincaré multiple splitting dimension*.

Exercise 2: Consider the effect of the two alternatives on our assignment of dimensions in Figure 1, especially for examples (5) through (8).

Note that the saddle of example (9) is a ruled surface, which can be separated by the removal of a single line (ref. Oprea). While the multiple-removal definition resolves some potential concerns about the dimension of the cylinder of example (7), it is still unsettling that the paraboloid of example (8) should remain dimension 3. After all, all three examples have area but no volume. In fact, we could stretch the saddle to form the paraboloid; in other words, those objects are topologically equivalent.

Two objects are topologically equivalent (homeomorphic) if there exists a continuous function from one onto the other that has a continuous inverse (a homeomorphism). Throughout the rest of this paper, we will use the term "map" to refer to continuous functions. Intuitively, a function is a map if points that are relatively close in the domain have relatively close images. We can formalize this notion with the $\varepsilon - \delta$ definition of continuity:

> A function $f: X \to Y$ $(X \subset \Re^m, Y \subset \Re^n)$ is continuous (is a map) if given $a \in X$ and given $\varepsilon > 0$ there exists a $\delta > 0$ such that $|f(a) - f(b)| < \varepsilon$ whenever $|a - b| < \delta$, $b \in X$.
>
> (Munkres, 1975, p. 126)

Note that the spaces we are considering are subspaces of an ambient Euclidean space (as indicated in the definition above) so that they inherit the Euclidean metric indicated by the absolute value signs in the definition above. We will later see the definition for uniform continuity, which differs from this definition only by the fact that δ may be chosen independent of *a*.

Exercise 3: Can you find a (continuous) map from $\Re^1$ onto S^1 (the circle in the *xy*-plane centered at the origin with radius 1)? If so, does the map have an inverse? In each case, describe the map explicitly or explain why it doesn't exist.

Exercise 4: Same questions for the *xy*-plane and the paraboloid.

Exercises 3 and 4, above, provide examples of pairs of spaces that are homeomorphic and of ones that are not. In fact, if dimension is to be a topological invariant (i.e., homeomorphic spaces have the same dimension), then it is clear that the paraboloid should have dimension two despite the assignment given by the Poincaré multiple splitting definition. Insisting on topological invariance suggests a second definition of dimension, called *simple topological dimension*. As such, an open, connected object has dimension n if it is homeomorphic to n-dimensional Euclidean space. We know from topology that homeomorphisms preserve openness and compactness (ref. Munkres), so for compact objects, we could look for a map to our compact canonical examples. We should note that it is not possible for an object to be homeomorphic to both compact and unbounded objects, because unbounded objects in Euclidean space are not compact. However, it was a formidable task for mathematicians at the turn of the last century to determine whether an object could not be homeomorphic to both I^n and I^m where $n \neq m$ (likewise, whether $\Re^n$ can be homeomorphic to $\Re^m$ for $n \neq m$).

Even if we do show that simple topological dimension is well-defined, our new definition of dimension faces other serious problems. Consider example (10) in Figure 1: it has Euclidean dimension 2 and Poincaré dimension 1, but it is not topologically equivalent to any of our canonical examples. In fact, we even have this problem with the circle, although in that case the problem is easily resolved by considering dimension as a local property.

Exercise 5: Find a homeomorphism from the open disc of radius ε centered at the origin in $\Re^2$ to $\Re^2$.

Considering dimension locally dispenses with many potential problems. We only need to consider *open neighborhoods* about points in the interior of the object. Recall that an open neighborhood on any object that is a subset of Euclidean space can be defined as the set of all points on the object whose Euclidean distance to a given point is less than some given positive value (e.g. open intervals are open neighborhoods of the real line). In considering open neighborhoods of interior points, we do not have to worry about boundaries of the object and can dismiss the canonical compact dimension *n* objects in favor of the unbounded ones, which are homeomorphic to connected, open neighborhoods within themselves. For instance, consider the example of the torus (11) in Figure 1. A torus with given parameters *a* and *b* is defined to be the set of all points formed from a choice of a point on a circle of radius *a* and a choice of a point on a circle of radius *b*. To visualize how these choices define the points of a torus, consider Figure 3. The first choice, given by angle ϕ, gives the location (along the lower dashed circle of Figure 3) of a circular slice of the torus. The second choice, given by angle θ, gives the location of a point on that slice. Algebraically, we can define the torus as the image of the map from $\Re^2$ to $\Re^3$ given by:

$$f(\theta,\phi) = ((R + r\cos\theta)\cos\phi, (R + r\cos\theta)\sin\phi, r\sin\theta)$$

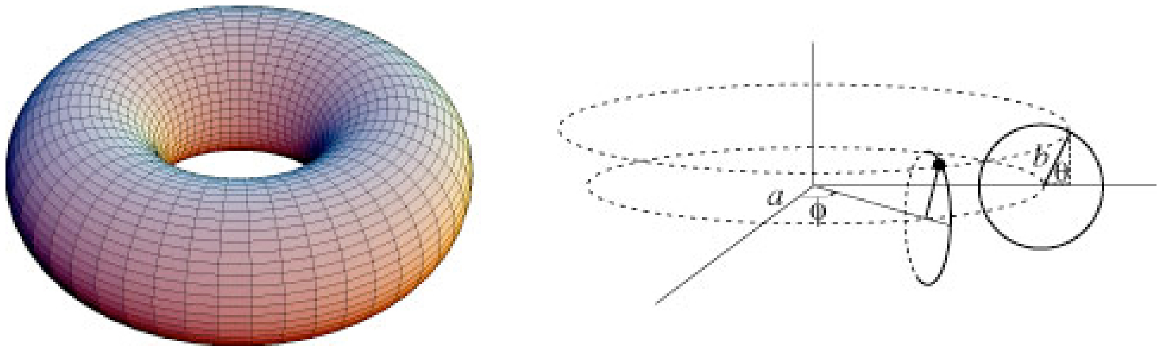

Figure 3. Addressing points on a torus.

Although it is not homeomorphic to any of the canonical objects, if we take the set of all points within a distance of $\varepsilon > 0$ (for ε small enough) of a given point in the image, we get an open disc that is homeomorphic to $\Re^2$. There is a subtlety here that we should not overlook. No matter how small we choose an ε-neighborhood around a point (say $f(\theta,\phi)$) on the torus, we never actually get a *flat* disc. Instead, we can consider the flat disc that is tangent to the torus at that point, meaning it is contained in the image of the best linear approximation to the *differentiable* map defining the torus.

The approximation for a differentiable map is so good that if we zoom in on smaller and smaller ε-neighborhoods (open sets that consist of all points on the object within a distance of ε of a given point on the object) about a point in the image of the differentiable map (the torus), we will not be able to distinguish between the image of $\Re^2$ under the linear approximation and the torus itself. This phenomenon is demonstrated in Figure 4.

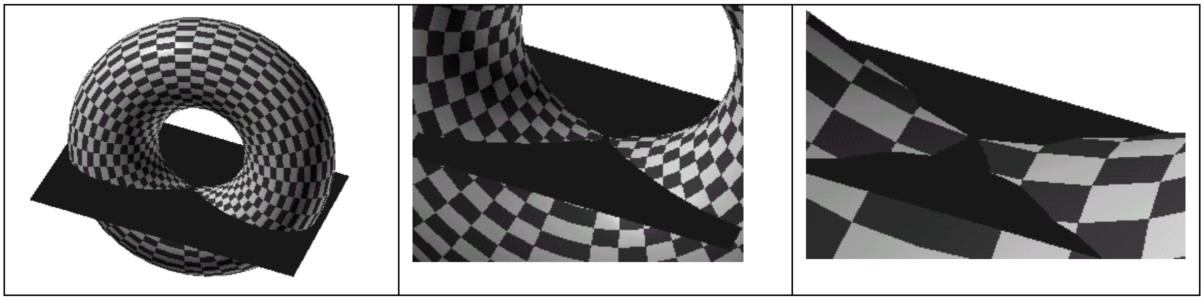

Figure 4. Zooms on a saddle point on the torus and the best linear approximation to the torus at this point.

The description above characterizes a class of objects called differentiable manifolds. We will consider the larger class of continuous manifolds: subsets of an ambient real space $\Re^N$ with the property that every point in the subset has an ε-neighborhood homeomorphic to $\Re^n$ for some nonnegative integer $n \le N$. So, we see that assigning dimension to any manifold is a simple matter so long as we consider dimension locally (i.e. we can assign dimension n to manifolds since the manifold inherently provides us with a local homeomorphism to a canonical object of dimension n). Examples (3), (4) and (10) in Figure 1, however, are not manifolds.

Suppose we were to loosen our criteria of finding local homeomorphisms in order to assign dimension and instead assign dimension n to an object if there exist onto (continuous) *maps* in each direction between (neighborhoods of each point within) the object and (neighborhoods of) $\Re^n$. This would resolve issues surrounding the object in example (10) because we can find onto maps in both directions between it and $\Re^1$. In fact there is an onto map in each direction between the unit segment and example (10). One of them stretches the unit segment as it traces back and forth across the arms in example (10); the other simply flattens example (10) onto the unit segment.

Exercise 6: Why are the objects indicated by examples (3), (4) and (10) not manifolds?

This approach would have seemed reasonable around the turn of the last century, until Cantor, Peano, and Hilbert introduced a class of functions that challenged intuition. These "monsters" demonstrate the need for the full strength of homeomorphisms in guaranteeing topological invariance between objects. In fact, mathematicians even questioned whether those maps would be strong enough, because Brouwer had not yet proved his theorem on Invariance of Domain, which is described at the end of the next section. Also in the next section, we will review some of the space-filling curve constructions and find some similarities between those objects and fractals, which emerged later in the nineteenth century to present new problems in assigning dimensions. For now, however, let us be content to resolve the problem exemplified by objects that are not manifolds by introducing a new definition of dimension and a new topological concept called covering.

We have already discussed ε -neighborhoods about points on an object in attempting to find homeomorphisms to Euclidean space. Although it is impossible to find a neighborhood about the vertex in example (10) that is homeomorphic to any such $\mathfrak{R}^n$, examining the intersection of neighborhoods covering the object is useful. Let us then create a *covering* of the object in example (10) –a collection of (open) neighborhoods in the ambient space about points on the object whose union is the entire object. We could possibly *refine* this covering by finding a new collection of neighborhoods that covers the object, such that every neighborhood in the new covering is contained within at least one neighborhood in the old covering (Munkres, 1975, p. 302). Figure 5 illustrates a covering of the object in example (10) from Figure 1 and a refinement of it.

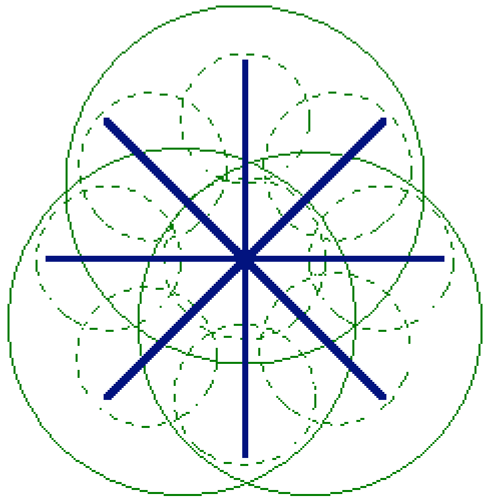

Figure 5. A covering (solid circles) and a refinement of it (dashed circles).

Notice that each point on the object in Figure 5 is contained within only one or two neighborhoods of the refinement, even though the original cover covered the vertex three times. It is a property of the unit interval, as well, that every cover can be refined so that no point on the interval is contained in more than two neighborhoods of the refinement.

Exercise 7: Determine the analogous properties for the unit square and unit cube. In other words, what is the smallest integer n such that every cover of the unit square (cube) can be refined until no point lies in more than n neighborhoods?

Such properties would provide an appropriate foundation for a new definition of dimension so long as we can show that the property is topologically invariant. Given that homeomorphisms take open neighborhoods of one space to open neighborhoods of another, topological invariance should not be hard to prove.

Exercise 8: Show that the properties referred to in Exercise 7 are topologically invariant.

Now, we can define the *topological dimension* of an object by way of the order of refinements of a cover: A refinement has order $m+1$ if some point on the object lies in $m+1$ neighborhoods in the refinement, and no point of the object lies in more than $m+1$ neighborhoods in the refinement (Munkres, 1975, p. 302). Suppose that, given any cover of an object, we can find a refinement of order $m+1$, but we cannot find a refinement of order m or less for some cover, then we say that the object has topological dimension m. Consider this definition applied to the objects in Figure 1. It seems to satisfy all of our concerns so far.

Exercise 9: What is the topological dimension of the rational numbers?

Before we go on to examine more potential problems, it is important to note that the concept of covering may also be used to redefine compactness. So far, we have assumed that a compact set is simply one that is closed and bounded. While this is true for compact sets in $\Re^3$, it will sometimes be useful to use the more general definition: a set is compact if every open cover of the set has a finite subcover (a subset of the cover). This definition will be especially helpful in the next section. Its equivalence (in real spaces) to the closed and bounded definition of compactness can be found in Munkres' *Topology* (1975) on page 174.

CHAPTER 3: SPACE-FILLING CURVES

A curve is the image of an interval under a (continuous) map. We are concerned with curves that are mapped onto the unit square. Such constructions are potentially dangerous to our efforts in establishing various concepts (particularly those of simple and local topological dimensions) as topological invariants.

Challenge: Visualize a map from the unit interval onto the unit square. Is your map continuous and onto: Is it also one-to-one? If your initial map does not satisfy one or more of these criteria, try to imagine a new map that satisfies more of them. (Note: This exercise may require several revisions, which you should make as you work through this section.)

We first must check that our constructions are indeed maps (i.e., they are continuous) and that they are onto. We may also want to know whether the map is one-to-one. An inverse function of the map exists if and only if the map is one-to-one, although we have not yet concluded that the inverse to a map would have to be continuous. If there is an inverse map and it *is* continuous, then the map is a homeomorphism. We may also find maps in both directions that are one-to-one and onto but not inverses of each other.

As an initial attempt to satisfy the various criteria laid out in Exercise 9, we might imagine a geometric construction that fills the square by spiraling around it or zigzagging through it. However, we cannot hope to do this in any one step because we will always leave gaps in the square that are not hit by the curve. Instead, we could define the curve as the limit of a process or a sequence of maps. For the zigzag, we can imagine a sequence of maps with more and more peaks (as illustrated in Figure 6), so that we get arbitrarily close to every point in the square.

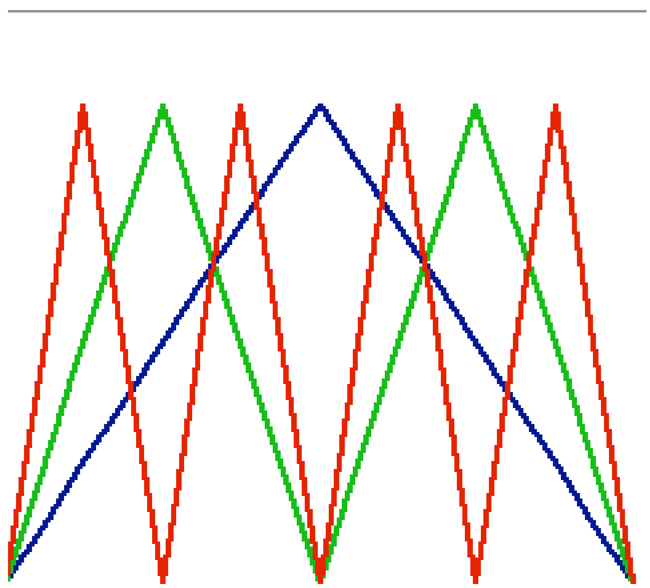

Figure 6. The first three images in a sequence of zigzag functions.

However, this sequence does not have a limit function at all (i.e. it does not even converge pointwise). Even if we had a continuous sequence of functions that did converge, there is no guarantee that the limit would be continuous. As an example, consider the sequence of maps $f_n(x) = x^n$ on the interval $[0,1]$ (graphs of the first, second, third and tenth function in this infinite sequence are displayed in Figure 7). This also provides an example of a sequence that is one-to-one at each stage, but not in the limit.

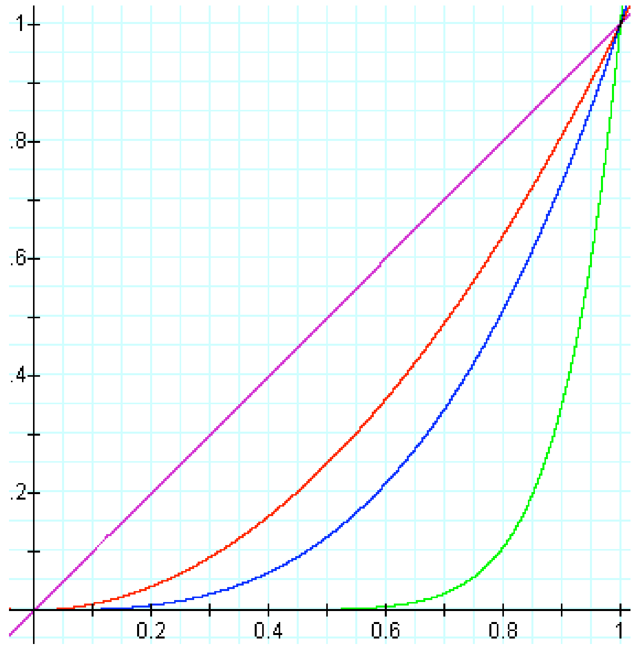

Figure 7. A sequence of maps whose limit is not a map.

We require a sequence of maps that will yield a (continuous) map in its limit such that its image (the curve) gets as close as we please to each point in the unit square. In fact, if we are able to show that, for every $\varepsilon > 0$, the curve passes within ε of each point in the square, then we can prove that the curve actually hits every point in the square.

In order to formalize the claim above, we will need the following topological theorem:

Theorem 1: The continuous image of a compact object is compact.

Since we are working in Euclidean space, where compact means closed and bounded, we could prove that the continuous image of a closed and bounded object is closed and bounded. However, the theorem is easily proven in general (non-metric) spaces by using concepts already introduced in the discussion of covering dimension, namely the definition of compactness.

Before going on, it may be useful to introduce (recall) a more topological definition of continuity, which is equivalent to the one given above so long as we are considering metric spaces. A function is continuous if and only if the preimages of open sets in the range are open sets in the domain. A proof of the equivalence of the two definitions for continuity given so far can be found in Munkres (1975) on page 127. The proof is straight-forward and the reader might attempt its proof before referencing the published one. Even then, the published proof might only be used for hints in its reconstruction.

Exercise 10: Use the definition of compactness given at the end of the last chapter in order to prove Theorem 1.

In addition to a metric, subsets of $\Re^3$ also inherit the Hausdorff property. A Hausdorff space is one in which two points can always be separated by disjoint open sets. This implies that points in $\Re^3$ are closed sets. In turn, this property implies the following theorem.

Theorem 2**:** Let $\{f_n\}$ be a convergent sequence of maps from I^1 to I^2 with a limit, f, that is also a map from I^1 to I^2. If for all ε >0 there exists an N such that each point in the unit square lies within ε of $f_n(I^1)$ for all n>N, then f is onto. (That is, f is a space-filling curve).

Exercise 11: Use Theorem 1 to prove Theorem 2.

Theorem 2 should guide future attempts at constructing space-filling curves: we have to construct a sequence of maps whose limit is a map and whose images get arbitrarily close to each point in the unit square. We can satisfy these requirements by recursively partitioning the unit square into four squares. The squares created at the first iteration would each be divided into four more squares at the second iteration, and so on. Let us call the 2^{2k} squares of length (width) $\frac{1}{2^k}$ (created by the k^{th} iteration) k-squares. For the k^{th} map in the sequence of maps, we will require that its image intersect each of the k-squares. Moreover, if the interval $[0, \frac{1}{2^{2k}}]$ is mapped to the 1st k-square and each successive interval of length $\frac{1}{2^{2k}}$ is mapped to a unique k-square, then the sequence will converge uniformly (as defined below). Figure 8 illustrates the square criterion for $k = 1$, with the image of the unit segment under first element, f_1, of a sequence of maps. Notice that the second fourth of the unit segment is mapped to the second square, where its image is displayed by a thick curve.

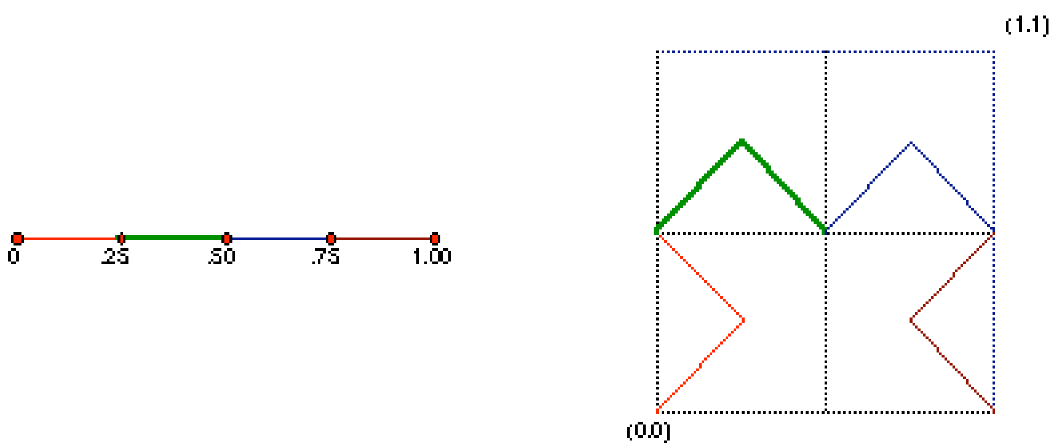

Figure 8. The second fourth of the unit segment and its image under f_1.

We could show that a sequence of functions satisfying the square criterion described above converges to a limit value for each point in its domain. This convergence defines a limit function, but we want more than that. In order to show that the limit function is continuous, we need uniform convergence (examples of the inadequacy of pointwise convergence are left for the reader in Exercise 13).

> Definition: Consider a sequence of functions $f_n: X \to Y$ from the set X to the metric space Y. Let d be the metric for Y. We say that the sequence f_n converges uniformly to the function $f: X \to Y$ if given $\varepsilon > 0$, there exists an integer N such that $d(f_n(x), f(x)) < \varepsilon$ for all $n \geq N$ and for all x in X.
> (Munkres, 1975, p. 129)

The term *uniform* in uniform convergence refers to the existence of a single integer N for which *all* values of f_n (with $n > N$) fall within a given $\varepsilon > 0$ of the values of the limit function. In other words, N depends on ε, but it is independent of points in the domain. A second characterization of a uniformly convergent sequence can be found in Rudin's *Principles of Mathematical Analysis* (1964, p. 147), along with a justification for the following criterion that is equivalent to the definition, at least for maps between Euclidean spaces (and actually for all complete metric spaces).

Cauchy criterion: Let $\{f_n\}$ be a sequence of functions from one Euclidean (complete metric) space D to another. The sequence converges uniformly on D if and only if for every $\varepsilon > 0$ there exists an integer N such that for all $m, n \geq N$ and for all $x \in D$, $|f_n(x) - f_m(x)| < \varepsilon$

(Note: The brackets used in the equation above and throughout this paper indicate a measure of distance between points.)

Exercise 12: Prove that any sequence of maps satisfying the criterion of k-squares described above also satisfies the Cauchy criterion and thus is uniformly continuous.

Exercise 13: Construct a sequence of maps that converges pointwise to a function that is not continuous.

Theorem 3: If a sequence $\{f_n\}$ converges uniformly to *f* and all f_n are continuous then *f* is continuous (Rudin, 1953, p. 149).

Proof of Theorem 3: Under the hypotheses, the sequence $\{f_n\}$ is uniformly convergent, and so it is also convergent to some function f. Now, given an $\varepsilon > 0$, we must find $\delta > 0$ so that $|f(x) - f(y) < \varepsilon|$ whenever *x* and *y* are in the domain of *f* and $|x - y| < \delta$. The following (triangle) inequality guarantees that we can do this. We need only argue that, for some positive integer m, we can make each of the absolute values on the right less than $\frac{\varepsilon}{3}$ whenever $|x - y| < \delta$ (where δ is a value that depends on ε and is guaranteed by the continuity of f_m) .

$$|f(x) - f(y)| \leq |f(x) - f_m(x)| + |f_m(x) - f_m(y)| + |f_m(y) - f(y)|$$

Exercise 14: Check the details for the suggested proof above. Where does the proof use uniform convergence? Specifically, why would the inequality fail to establish continuity in the case of the limit of the sequence of functions $f_n(x) = x^n$ on the interval $[0,1]$?

Theorems 2 and 3 and Exercise 12 demonstrate that any sequence of continuous functions from the unit interval to the unit square satisfying the k-square condition described above produce a space-filling curve in their limit.

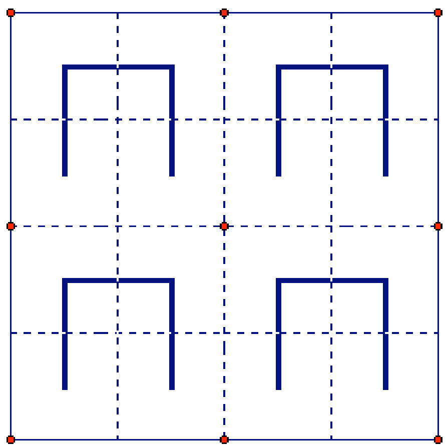

Figure 9a. Ordering the sub-squares.

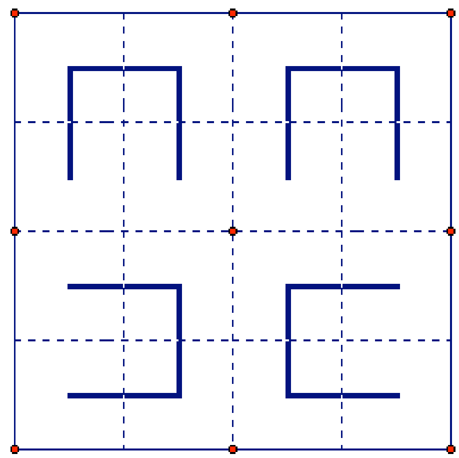

Figure 9b. Transforming the order.

Recall that constructing a space-filling curve under the prescribed guidelines amounts to ordering the k-squares in such a way that they respect the order of the (k-1)-squares. Let us try this using a base-4 addressing system where every point in the unit square is given by some expansion in base-4. We can order each of the four new k-squares within each (k-1)-square 0,1,2,3 starting from the lower left square clockwise to the lower right square. Our constructed curve would then respect this order. But this first try would yield a disconnected set after the first stage, as indicated by Figure 9a. We could correct this by reflecting the lower left and lower right squares as indicated by the direction in which the "U" shapes open in Figure 9b. Now we can extend the construction at each stage to a (connected) curve by joining endpoints. This completes our construction satisfying all of our conditions for a space-filling curve called the Hilbert curve, named after David Hilbert who first wrote about it in 1890 (Peitgen et al, I, 1992, pp. 110-111). The first few stages of Hilbert's construction are displayed in Figure 10.

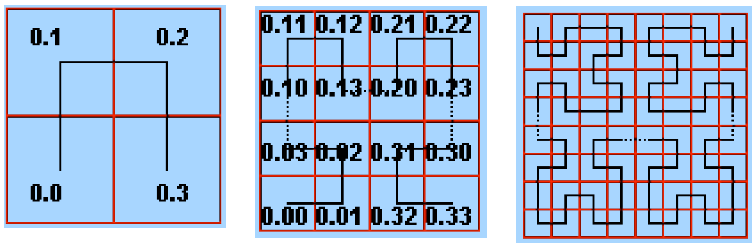

Figure 10. Hilbert's curve at stages 1, 2 and 3.

Exercise 15: A Peano curve follows the same pattern through the squares except that it begins with a "^" shape with endpoints in the lower corners and vertex at the center of the unit square. Follow the pattern of Hilbert's construction to create the first few stages of Peano's curve.

Peano created other space-filling curves as well, including one that divided each square (and subsequent sub-squares) into nine sub-squares instead of four. We could use a base-nine addressing system on these sub-squares to suggest new space-filling curves. In fact, we could partition the unit square into any composite number of parts, so long as the length and width of the sub-squares shrink exponentially (Why is it important for *both* to shrink?). Shrinking the squares exponentially by a factor of n allows us to use a base-n addressing system. However, self-similarity in the two examples we have constructed thus far suggests a new approach. In particular, notice that each stage of the Peano curve contains four scaled-down copies of the previous change, although they may be transformed by rotations or reflections. In fact the openings of each "U" in Figure 9b indicate the transformations used.

Using transformations, we can develop space-filling curves through iterations of transformations, rather than through a sequence of maps. To describe the recursive transformations, we use the metaphor of a multiple reduction copy machine (MRCM), wherein an image is shrink by some factor and copied a specified number of times (Peitgen et al, 1992, p. 28). The copies are then arranged to form a new image, which is used as the new image for the same recursive process. Note that Figure 8b represents the results through one iteration of an MRCM and that we can recursively apply the MRCM to each of the four resulting squares. The image after each iteration is identical to the image of the corresponding map in the sequence of maps. In the next section, we will further examine the application of MRCM's – in generating fractals. For now, it will suffice to note that their application creates self-similarity, as can be seen in the Peano curve.

The question remains whether there exists a one-to-one space-filling map. In the case of Peano's curve, it is obvious after a few stages that the map is not one-to-one. However, the Hilbert curve does not appear to intersect itself. In fact, upon further examination we can conclude that at each stage the map is one-to-one. But the following theorem implies that if the limit of these maps is indeed one-to-one, the limit function would also be a homeomorphism! On the other hand, we might consider examples of sequences of one-to-one functions whose limits are not one-to-one (such as $f_n(x) = \frac{x}{n}$).

Theorem 4: A bijective (one-to-one and onto) map from a compact space to a Hausdorff space has a continuous inverse (i.e. the map is a homeomorphism).

Exercise 16**:** Use Theorem 1 to prove the following special case of Theorem 4: A bijective map from the interval to a metric space is a homeomorphism. (Hint: A function is continuous if and only if the preimages of *closed* sets in the range are *closed* sets in the domain.)

The proof of the special case is good enough to determine that bijective maps from the unit interval to the unit square must have continuous inverses and are thus homeomorphisms. The proof of Theorem 4 seems to require us to know that compact subsets of Hausdorff spaces are closed; proofs of this and Theorem 4 can be found in Munkres (1975, pp. 166-167). However, we should not take for granted that all bijective maps are homeomorphisms. Consider, for instance, the function

$$f(t) = (\cos(2\pi t), \sin(2\pi t))$$

taking the half open interval [0,1) to the unit circle S^1 in $\Re^2$.

Exercise 17**:** Prove that *f* is a bijective map, but that its inverse is not continuous.

At this point, we know that we can construct space-filling curves by satisfying the *n*-square condition. Later, we will explore another method, Cantor's construction. So far, we have one candidate (the Hilbert curve) that appears to be one-to-one and thus a homeomorphism from the unit interval to the unit square. We already know to be skeptical of such appearances (Figure 10 illustrates a sequence of curves that appear to have a one-to-one limit, but consider the sequence of curves on the open interval illustrated in Figure 6.), but knowing the Hilbert curve would provide a homeomorphism from the unit interval to the unit square, destroying the topological invariance of many of our conceptualizations of dimension, gives us even more reason to be careful. Let us examine just one suspicious point on this curve.

Since the Hilbert map is onto, we know that the center of the unit square must get hit, but from which direction? The vertical line of symmetry in the construction implies that it gets hit from at least two directions. We will examine this more closely later using addressing systems, but we can already generalize this problem of symmetry and one-to-oneness in our constructions. This gives us reason to believe none of our space-filling constructions can be one-to-one, so we might better spend our energy trying to prove this assertion rather than trying to construct one-to-one space-filling curves. In fact, we could compare one of several topologically-invariant properties that differ between the unit segment and the unit square. Finding just one topologically-invariant property of one space that is not shared by the other would demonstrate that there can be no homeomorphism between them. For example, when we delete an interior point from the unit segment (any point but an endpoint), it becomes disconnected; this is not so for the unit square. The larger task would be to prove that dimension itself is a topologically-invariant property; this is just what Brouwer did in 1911.

The following formulation of Brouwer's theorem on invariance of domain is found in texts on algebraic topology. It is sometimes presented as a generalization of the Jordan curve theorem and as an application of homology.

> Theorem 5: If U and V are homeomorphic subsets of S^n (the sphere with covering dimension *n*) and U is open in S^n, then V is also open in S^n.
>
> (Spanier, 1966, p. 199)

Exercise 18: Show that Theorem 5 implies that $\Re^n$ is homeomorphic to $\Re^m$ if and only if $m = n$.

Instead of proving Theorem 5, we will prove a modified version of its most important implication. In particular, we will show that $\Re^n$ is diffeomorphic to $\Re^m$ if and only if $m = n$. A diffeomorphism $f : X \rightarrow Y$ is a map between subsets of Euclidean spaces that is smooth (i.e. around each point in the domain, there is an open set of which *f* can be extended by a map with continuous partial derivatives of al orders) and has a smooth inverse (ref. Guillemin and Pollack). Let ϕ be such a map between $\Re^n$ and $\Re^m$, and, without loss of generality, let $n \geq m$. Its derivative is a linear function between the Euclidean spaces. In fact, multivariable calculus tells us that the linear function is given by the Jacobian $m \times n$ matrix of partial derivatives. Moreover, applying the chain rule to the identity map

$$i : \Re^n \xrightarrow{\phi} \Re^m \xrightarrow{\phi^{-1}} \Re^n$$

we get that $d\phi_x (d\phi_{\phi(x)}^{-1})$ is the $n \times n$ identity matrix. Thus, each matrix in the product must have maximal rank, giving $n = m$. Having proved invariance of domain for diffeomorphisms between Euclidean spaces, we will now assume it for homeomorphisms.

Exercise 19: Use invariance of domain to prove that the unit square and the unit cube are not homeomorphic.

Now we know that space-filling curves cannot be one-to-one, but which points get hit more than once? Can we remove them? Do any points get hit infinitely often? We return to our *k*-square addressing system in order to answer such questions.

We have said that our construction of space filling curves amounts to an ordering of the *k*-squares at each stage. So, in following the *k*-square condition for our constructions, the base-4 expansion of each point in the unit interval gives the address of its image in the unit square! For example, for the Hilbert curve whose first three stages are mapped out in Figure 9, a point with a base-4 expansion beginning with $.032$ would go somewhere in the shaded square shown in Figure 11. To see this, just examine the three illustrations in Figure 10. The first one illustrates the ordering of the 1-squares with the lower-left square being the first one (corresponding to points with a 0 in their tenths place) and the other squares following the order in which they are hit by the curve. The other two illustrations likewise demonstrate the orderings of the 2-squares within each 1-square and the 3-squares within each 2-square.

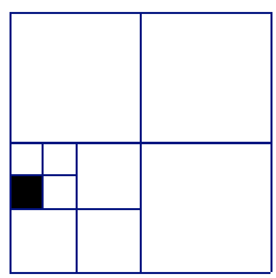

Figure 11. Example of addressing of a point in the unit square.

Now, we face a potential problem with this system. Since numbers terminating in 0's can also be written as numbers terminating in 3's (eg. $.032\overline{0} = .031\overline{3}$), we must consider whether the addressing system provides a well-defined description of our space-filling curve. It may or may not, but we can alleviate this problem by choosing representations of the domain given by quaternary expansions with no infinitely repeating 3's. So, we can return to our question of finding many-to-one points.

Exercise 20**:** Using the addressing system described above, answer the following question concerning any given *k*-square construction of a space filling curve: Which points in the unit square get hit more than once? Argue that there are only a countably infinite number of points that get hit more than once and that none gets hit more than four times.

CHAPTER 4: FRACTALS AND HAUSDORFF DIMENSION

In this final section, we will examine the construction of a space-filling curve that does not follow the *k*-square construction. It is based on the construction of a simple fractal called Cantor's set. This fractal has important implications and brings up apparent paradoxes in set theory, especially concerning countable and uncountable sets. The strange properties that this and other fractals hold yield a new conceptualization of dimension, embodied by two new definitions. The first is based on the self-similarity that fractals possess, and the second is named for Hausdorff, whom we have already mentioned.

In short, fractals are "fragmented, fractured, self-similar shapes of fractional dimension" (Mandelbrot, 1977, p. 4). They are constructed through infinite recursion of an algorithm or function; thus, they can never be truly constructed outside of the imagination. We have already seen two fractals, in Hilbert's and Peano's curves. We will continue by describing simple fractals such as Cantor's Set, consider their implications, and then move into the introduction of a complex class of fractals called Julia sets.

Cantor's set, like all of the fractals we will consider, is the result of applying a simple process indefinitely. Beginning with a unit segment, we remove the open middle third of the segment. Then, we remove the open middle thirds of each of the remaining segments and so on (see Figure 12). At least one natural question arises from this process: what is left in Cantor's set? Will there be nothing but a discrete set of endpoints? Will there be some collection of tiny intervals?

Figure 12. The first few iterations in the construction of Cantor's set.

Exercise 21: Demonstrate that the number of endpoints in the Cantor set is countably infinite.

Hint: Consider how many new endpoints are created at each stage.

Exercise 22: Show that there is an uncountable number of points in the Cantor set fractal.

Hint 1: Consider the ternary (base 3) expansion of points in the unit interval.

Hint 2: Show that the remaining set is in one-to-one correspondence with the unit segment.

The Cantor set is precisely those points on the unit interval with no 1's in their ternary expansion. Since there is an uncountable number of points in Cantor's set, it must include more than endpoints. In fact, we know precisely the ternary expansions of these extra points. Although endpoints are points that eventually terminate in repeating 0's or repeating 2's, there are uncountably many points in Cantor's set whose ternary expansions do not terminate but alternate arbitrarily between 0's and 2's indefinitely. As such, the Cantor's set is isomorphic to the set X of infinite sequences of 0's and 1's.

We already noted that Peano's curve is a fractal, and so it should not surprise us that Cantor's set can also be used to construct space-filling curves, albeit ones of a different nature. We will begin by demonstrating that the two uncountable sets mentioned above (Cantor's set and the set of binary free choices, X) are homeomorphic. Until now, in discussing homeomorphisms, we have not explicitly mentioned the topologies of the spaces we compare. In fact, two topological spaces are only homeomorphic relative to their defined topologies because the topology defines what are considered open sets in the space. However, we have been dealing with metric spaces, for which there are standard topologies based on neighborhoods of various radii (and for which our stated definition of open sets applies). In light of Exercise 22, the following function would seem a natural choice for the homeomorphism.

$$\Psi(x) = \sum_{n=1}^{\infty} \frac{2x_n}{3^n}, \text{ where } x = (x_1, x_2, \ldots) \text{ is an element of } X.$$

The image of X under Ψ, which we will call C for Cantor's set, inherits a subspace topology from the standard topology of the real line. Now, to complete the argument that C and X are homeomorphic with respect to the topologies that we described, we need only show that Ψ and Ψ^{-1} take open sets to open sets. If we can do this, we will have that Ψ and Ψ^{-1} are continuous, i.e. Ψ is a homeomorphism. First, we must decide how open sets of X look. The most natural choice is to apply the product topology to the infinite product $\prod_{n=1}^{\infty}\{0,1\}$, where $\{0,1\}$ is given the discrete topology. The reader may wish to review the product topology in another text (Munkres, 1975, p. 114), but we note here that the basis for this topology on X consists of all points (ordered sequences of 0's or 1's) in X having the same values in some fixed finite set of positions. Recall that the topology on X determined by its basis is simply "the collection of all unions of elements" of the basis (Munkres, 1975, p.78). Points are not open in this topology because the fixed set of positions for the basis element would have to be infinite.

Exercise 23: Name a basis for C under the subspace topology. Use this to complete the argument that C and X are homeomorphic.

We will use this homeomorphism to construct a space-filling curve in a manner that is quite different from our approach in the last section. First, we can find a surjection (onto map) from X onto the unit square by mapping terms in the sequence defining a point $x \in X$ to the binary expansions of coordinates in the unit square. We shall use the ordered odd terms to define the x-coordinate and the even terms to define the y-coordinate of the image point. For example, (0,1,1,0,0,0,0,…) would go to (0.01, 0.1).

Exercise 24: Clearly, the map defined above gives a surjection, but is it truly a map? That is, is it continuous?

Since X and C are homeomorphic, we have a surjection, f, from C to the unit square. Actually, f is the Cartesian product of two maps, f_1 and f_2, from C to I (giving the x- and y-coordinates, respectively, of image points in the unit square). We could generalize this construction to surjections onto objects of higher dimension as well (such as the unit cube) just as we could have done for the space-filling curves of the last section, but let us settle for the completion of a new space-filling curve, filling the unit square. We need only extend our maps, f_1 and f_2, of C to I to ones from I to I. The Tietze extension theorem guarantees that we can so extend.

> Theorem 6 (Tietze): Let Y be a normal topological space, and let A be any closed subset of Y. Then any (continuous) map from A to I may be extended to a (continuous) map from Y to I.
>
> (Munkres, 1975, p. 212)

The proof of Tietze's theorem can be found in Munkres. Here, we will only demonstrate that f_1 and f_2 satisfy the hypotheses. First of all, a normal space is one in which a pair of disjoint closed sets can be separated by a pair of disjoint open sets, each containing one of the open sets. Y=I meets this condition (as do all metric spaces), as described by Munkres (1975) on page 198. The second condition, that C is closed in I, is left as an exercise.

Exercise 25: Use the geometric construction of C to argue that it is closed in I.

So we have a new breed of space-filling curve constructed from Cantor's fractal. Actually, we could construct an extension and thus a space-filling curve without Tietze's theorem. Consider the segments deleted at each stage of the geometric construction of C. We can extend f_1 and f_2 linearly on each of these intervals to define extensions, g_1 and g_2, respectively giving the x- and y-coordinates of image points in the unit square. Instead of settling for an existence proof, we can try to visualize the curve and consider its properties. To begin, consider Figure 13, depicting the linear extension along the open interval removed in the first stage of the construction for Cantor's set. The left and right endpoints of that interval get mapped to (.05,1) and (0.5,0), respectively, in the manner described above. The linear extension simply maps the interval linearly to the open interval between the image points.

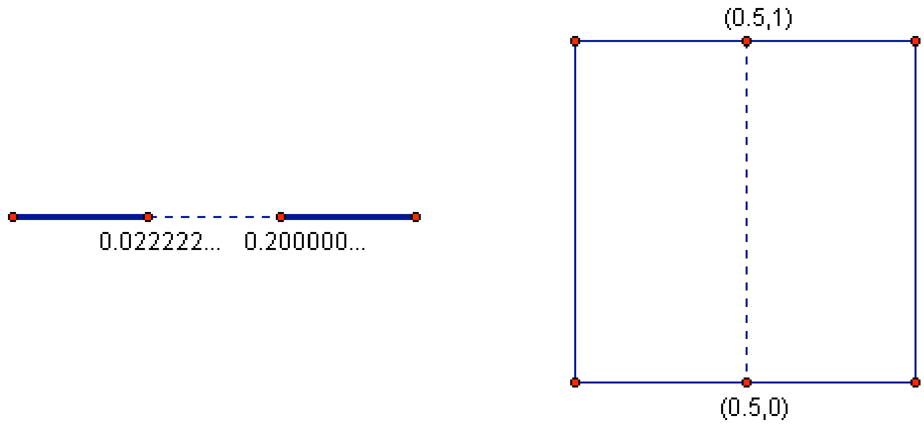

Figure 13. A linear extension on Cantor's set.

To see how our new curve is different from space-filling curves we constructed in the last section, let us return to the question of one-to-one-ness. Previous constructions yielded no points in the range at which the map was worse than four-to-one and that these points occurred at vertices in the grid formed in the unit square by its sub-squares. Now there are potential bad points (points in the range at which the map was not one-to-one) along our linear extensions, but we will consider them in a moment. Points whose *x*- or *y*-coordinates – given respectively by maps f_1 and f_2 from C – can be represented in two ways will also be bad points. These consist of points whose x- or y-coordinates end in repeating 1's in their binary expansions. In other words, these are points in C whose sequence consists of alternating 2's (where every odd (even) position in the sequence are 2's and the other positions are arbitrary).

Exercise 26: Show that the bad points from C form a countable collection of horizontal and vertical line segments as their images, forming an orthogonal grid in the unit square.

Returning to the linear extensions, we see that these are also line segments in the unit square, connecting the images of points that terminate in repeating 0's (the images of right endpoints of C) to those that terminate in repeating 2's (the images of left endpoints of C), as illustrated in Figure 13.

Exercise 27: Show that all of the linear extensions described above connect endpoints of the grid described in exercise 26 and form horizontal or vertical line segments overlapping sections of the grid.

Already we can see that this Cantor construction for a space-filling curve is very different from the sub-square constructions of space-filling curves. Whereas bad points in the latter were restricted to vertices of the sub-squares, in the former there are whole segments' worth of points in the unit square that are hit at least three times and all of the grid lines are hit at least twice. The construction relies on some interesting properties of Cantor's set that have implications for dimension.

We have already noted that Cantor's set is uncountable, and yet a simple computation involving its geometric construction shows that its length is zero. What should this imply about its dimension? We could show that the topological dimension of Cantor's set is zero, as we did for the set of rational numbers in Exercise 9. On the other hand, the construction of Cantor's set appears to produce more than a discrete set of points (whereas the rational numbers are a discrete set). By returning to the MRCM metaphor and examining the geometric construction of Cantor's set, we might note that it is self-similar. In other words, we could magnify the part of the fractal contained in any of the closed subintervals created at any stage in order to reconstruct the whole;

since the fractal is the result of infinite recursion, we do not lose any detail in the magnification. Specifically, we could put the left half of the Cantor set fractal into a metaphorical copy machine and magnify it to 300% to reproduce the whole. Alternatively, we could paste two copies of the left half side-by-side to reproduce the whole. We could use either method for reproducing any self-similar object: scaling a part by some factor (3 in this case) or pasting together some number of copies of that part (2 in this case).

Exercise 28: Consider the values of the scale factor, *s*, and pasting number, *n*, for some familiar self-similar objects, such as a line segment, a filled triangle, or a filled cube. What relation can you produce between the two numbers and the dimension, *d*, of the objects in each case?

What if we were to insist that the relation $s^d = n$ from Exercise 28 hold for all self-similar objects, such as Cantor's set or other fractals? Consider, for instance, Sierpinski's triangle, formed by connecting the midpoints of a triangle with segments and recursively applying this process to all but the midpoint triangles formed at each stage. The result of this infinite process is represented in Figure 14a. We could construct another fractal from a line segment and recursively removing the middle third (just as we do for the geometric construction of Cantor's set) and forming the other two legs of an equilateral triangle in its place. Continuing recursively on each new segment, we generate a fractal called Koch's curve. If we apply this process to each line segment of a triangle, the resulting object is called Koch's snowflake and is depicted by Figure 14b. Though Koch's snowflake is commonly referred to as a fractal, it does not satisfy the self-similarity requirement, although it is self-similar locally.

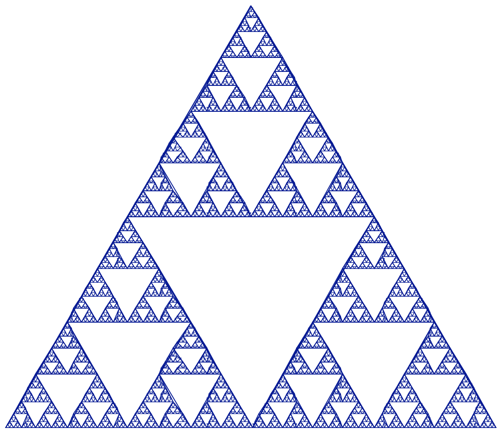

Figure 14a. Sierpinski's Triangle

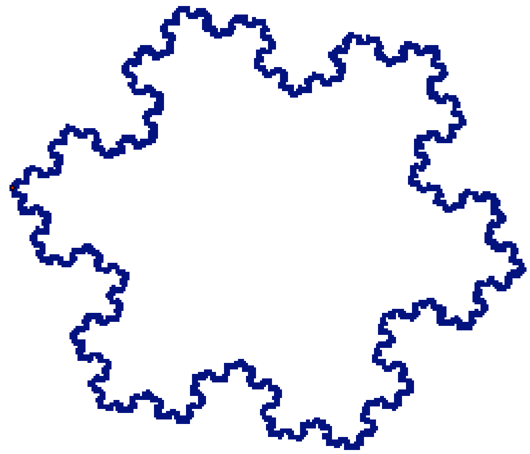

Figure 14b. Koch Snowflake

Exercise 29: What is the self-similar dimension (defined by the relation $s^d = n$) of the Cantor set and the Siepinski triangle? What is the length of each fractal?

As with the space-filling curves, Sierpinski's triangle and Koch's snowflake have infinite length confined in a bounded region. Their dimension suggests that they "almost" fill area, but they actually fill none. Something similar happens with Cantor's set. So, we see that fractals introduce the possibility of non-integer dimensions. We could design our own fractals by defining a simple process and applying infinite recursion. The self-similarity dimension of the fractal is then determined by our scaling factor and the number of parts pasted. Keep in mind that our process should result in a bounded object that increases exponentially in number of self-similar parts while decreasing exponentially in the size of each part. Also, overlap can destroy the self-similarity of the object.

Not only do fractals help us to redefine dimension, they also introduce counterexamples to some of our previous conceptions of dimension and continuity. For instance, consider the Koch curve formed by any of the three "sides" of the snowflake from Figure 14b. We can see that it is homeomorphic to the unit segment and yet the two objects have different self-similarity dimensions! The Koch curve also provides an example of the image of a function that is continuous everywhere and differentiable nowhere; a second example of such an object is developed more formally in Munkres (p. 297).

Since there are many fractal-like objects (such as the Koch snowflake) that are not self-similar, we must employ other means for determining the dimensions of such objects. Furthermore, for many objects (again, such as the Koch snowflake), the measures of their dimensions seem to be dependent on scale. To illustrate, consider the fractal-famed problem of determining the length of the coastline of England. Measuring it in straight mile intervals would yield a smaller length than measuring it a ruler-length at a time because we would miss many of the crevices along the coast. In fact, if we use finer and finer scales, the length of the coastline would appear to increase without bound. To solve such problems, we could introduce a new definition of dimension that measures the n-dimensional volume of objects with an infinitely fine scale.

Like self-similarity dimension, Hausdorff dimension allows for non-integer dimensions, but it also applies to non-self-similar objects. Like topological dimension, it employs a countable collection of small "neighborhoods" to cover the object, but the focus is on their measure and not on their order of intersection. Also, for simplicity's sake, in this case the "neighborhoods" are cubical instead of spherical. The idea is to consider all countable covers of the object by n-dimensional *cubes*. We will consider the most efficient cover (least overlap of the n-cubes) for a given side length ε. Summing up the n-dimensional volume of these n-cubes for the most efficient cover and then taking the limit of these sums as $\varepsilon \to 0$ gives the Hausdorff n-dimensional measure of an object, J. Formally, this can be expressed and generalized by

$$H^n(J) = \lim_{\varepsilon \to 0^+} \inf_{A_\varepsilon} \sum_{A \in A_\varepsilon} \left[sidelength(A) \right]^n$$

where A_ε is any cover of *J* by cubes of sidelength less than or equal to ε.

Exercise 30: Calculate the Hausdorff 1-dimensional measure of the unit segment.

Now, to calculate the Hausdorff 2-dimensional measure of the unit segment, we might imagine a cover A_ε made up of adjacent 2-squares of equal side length. Then there would be $\frac{1}{\varepsilon}$ of them and each would have measure ε^2. But the limit as $\varepsilon \to 0$ for the total area of such covers would be zero. Thus, the Hausdorff 2-dimensional measure has zero as an upper bound and so must be 0. We can be happy with this because one-dimensional objects, such as the unit segment, should have zero 2-dimensional measure. Moreover, we can use Hausdorff dimensions to determine the dimension of an object; we will say that the *Hausdorff dimension* of an object is the largest value of *n* for which the Hausdorff *n*-dimension of the object is non-zero. To check that this definition is workable, we need to prove the following exercise.

Exercise 31: Show that if the Hausdorff *n*-dimensional measure of an object is zero, so is its Hausdorff *m*-dimensional measure for *m*>*n*.

So, what happens when we calculate Hausdorff *n*-dimensional measures of fractals?

Exercise 32: Find upper and lower bounds for the Hausdorff dimension of Sierpinski's triangle (Figure 11a).

Notice that if an object has Hausdorff dimension *d*, then for *0<n<d* the Hausdorff *n*-dimensional measure is infinite. Trying to determine the precise Hausdorff dimension of an object can get confusing because the infimum may be difficult to determine. However, we can make the situation more workable if we consider covers by grid-like patterns of *n*-cubes all of side length ε, as illustrated in Figure 15. Then, the sum in the definition of Hausdorff *n*-dimensional measure could be expressed by $N_\varepsilon(J)\cdot\varepsilon^d$, where $N_\varepsilon(J)$ is the required number of *n*-cubes in the cover.

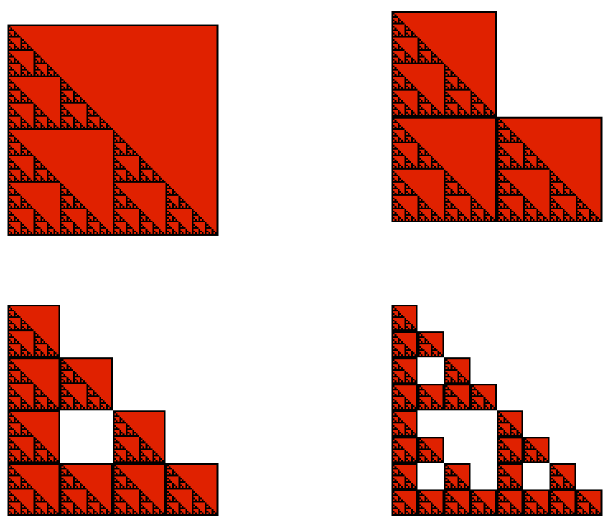

Figure 15. Grid-like covers of Sierpinski's triangle.

Now, we are looking for the unique value d for which $\lim_{\varepsilon \to 0^+}\left[N_\varepsilon(J) \cdot \varepsilon^d\right]$ is non-zero and non-infinite. In other words, $c \cdot \lim_{\varepsilon \to 0^+}\left[N_\varepsilon(J) \cdot \varepsilon^d\right] = 1$ for some positive real value c. But, since log is a uniformly continuous function on $\left[\frac{1}{c}, \infty\right)$, this means $\log c + \lim_{\varepsilon \to 0^+}\left[\log N_\varepsilon(J) + d \log \varepsilon\right] = 0$. Finally, we can show that $d = \lim_{\varepsilon \to 0^+} \frac{\log N_\varepsilon(J)}{\log\left(\frac{1}{\varepsilon}\right)}$. This is referred to as *box-counting dimension*, because it amounts to counting the n-cubes (boxes) that intersect the object for a given grid size, comparing its log to the log of the reciprocal of the grid size, and taking the limit of this ratio as the grid size goes to 0.

Box-counting dimension can be determined within any error of tolerance for any visual object using empirical data that can be generated with a computer program; so it is a very practical measure of dimension. Although I will introduce no new exercises here, the pragmatism of this technique may be noted in its use in various fields of science, including geography, materials science, engineering and biology. A quick internet search for "box-counting" will yield numerous examples and illustrations, such as that of using box-counting dimension for texture image detection in image processing.

Noting the additional criteria placed on covers used to compute Hausdorff dimension, it should not surprise us that box-counting dimension is always greater than or equal to Hausdorff dimension. Likewise, Hausdorff dimension is always greater than or equal to topological dimension. These three definitions constitute a mature conception of dimension, each having evolved from more primitive conceptions and each contributing where the other two are limited: Topological dimension has the desirable attribute of topological invariance but is limited to integer values; Hausdorff dimension allows for non-integer values, which accommodate the need to measure the complexity of fractals, but it is impractical in computation; box-counting dimension provides practicality by imposing natural but unnecessary restrictions on covers. Such exchanges of affordances and limitations are prevalent in the development of mathematics; we can find similar patterns in the development of complex numbers from the counting numbers. They point to logical but subjective nature of the field of mathematics.

REFERENCES

Mandelbrot, B. (1977). *The fractal geometry of nature*. New York: W. H. Freeman and Company.

Munkres, J. (1975). *Topology: A first course*. Englewood Cliffs, NJ: Prentice Hall.

Oprea, J. (1997). *Differential geometry and its applications*. Upper Saddle River, NJ: Prentice-Hall.

Peitgen, H., Jürgens, H., and Saupe, D. (1992). *Fractals for the classroom, Part one*. New York: Springer-Verlag.

Poincaré, H. (1952). *Science and hypothesis*. New York: Dover Publications,Inc..

Rudin, W. (1953). *Principles of mathematical analysis*. New York: McGraw-Hill Inc..