

CMSC 425: Lecture 12

Procedural Generation: Fractals

Reading: This material comes from classic computer-graphics books.

Procedural Generation: One of the important issues in game development is how to generate the objects that make up your scene: humans and animals, tools and weapons, furniture, structures, natural objects like plants, mountains, rocks, water, and clouds. Usually these objects, or *models*, are generated by a human designer. But, since this is not a course in art, we will not delve into this subject. Alternatively, objects can be generated by the computer. This is called *procedural generation*. In the next few lectures, we will discuss a number of different techniques for generating geometric models procedurally.

Fractals: One of the most important aspects of any graphics system is how objects are modeled. Most man-made (manufactured) objects are fairly simple to describe, largely because the plans for these objects are be designed “manufacturable”. However, objects in nature (e.g. mountainous terrains, plants, and clouds) are often much more complex. These objects are characterized by a nonsmooth, chaotic behavior. The mathematical area of *fractals* was created largely to better understand these complex structures.

One of the early investigations into fractals was a paper written on the length of the coastline of Scotland. The contention was that the coastline was so jagged that its length seemed to increase as the length of the measuring device (mile-stick, yard-stick, etc.) got smaller. This phenomenon was identified mathematically by the concept of the *fractal dimension*. The other phenomenon that characterizes fractals is *self similarity*, which means that features of the object seem to reappear in numerous places but with smaller and smaller size.

In nature, self similarity does not occur exactly, but there is often a type of *statistical* self similarity, where features at different levels exhibit similar statistical characteristics, but at different scales.

Iterated Functions and Attractor Sets: One of the examples of fractals arising in mathematics involves sets called *attractors*. The idea is to consider some function of space and to see where points are mapped under this function. An elegant way to do this in the plane is to consider functions over complex numbers. Each coordinate (a, b) in the real plane is associated with the complex number $a + bi$, where $i^2 = -1$. Adding and multiplying complex numbers follows the familiar rules:

$$(a + bi) + (c + di) = (a + c) + (b + d)i,$$

and

$$(a + bi)(c + di) = ac + adi + bci + bdi^2 = (ac - bd) + (ad + bc)i.$$

Define the *modulus* of a complex number $a + bi$ to be length of the corresponding vector in the complex plane, $\sqrt{a^2 + b^2}$. This is a generalization of the notion of absolute value with reals. Observe that the numbers of given fixed modulus just form a circle centered around the origin in the complex plane.

Now, consider any complex number $z_0 = (a_0 + b_0i) \in \mathbb{C}$. If we repeatedly square this number, $z_i \leftarrow z_{i-1}^2$, for $i = 1, 2, 3, \dots$ then it is easy to verify that with each step the modulus is also squared. If the modulus of z_0 is strictly smaller than 1, then the resulting sequence of complex numbers will become smaller and smaller (in terms of their moduli) and hence will converge to the origin in the limit. If the modulus of z_0 is strictly larger than 1, the moduli will grow to infinity, implying that the sequence will move arbitrarily far from the origin. Finally, if the modulus is equal to 1, it will remain so, and the sequence will spiral around the unit circle.

In general, we can do this using any function $f : \mathbb{C} \rightarrow \mathbb{C}$ on the complex plane. We define the *attractor set* (or *fixed-point set*) to be a subset of nonzero points that remain fixed under the mapping. Note that it is the set as a whole that is fixed, even though the individual points tend to move around (see Fig. 1).

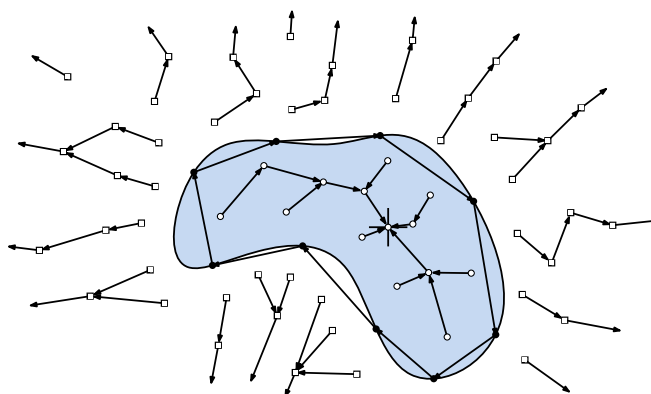


Fig. 1: An iterated function. The subset within the blue region converges to the origin. The attractor set (the black curved boundary) is preserved under the function.

Julia and Mandelbrot Sets: For any complex constant $c \in \mathbb{C}$, consider the iterated function

$$z_i \leftarrow z_{i-1}^2 + c \quad \text{for } i = 1, 2, 3, \dots$$

Now as before, under this function, some points will tend toward ∞ and others towards finite numbers. However there will be a set of points that will tend toward neither. Altogether these latter points form the *attractor* of the function system. This is called the *Julia set* for the point c . An example of such a set is shown in Fig. 2(a).

A common method for approximately rendering Julia sets is to iterate the function until the modulus of the number exceeds some prespecified threshold. If the number diverges, then we display one color, and otherwise we display another color. How many iterations? It really depends on the desired precision. Points that are far from the boundary of the attractor will diverge quickly. Points that very close, but just outside the boundary may take much longer to diverge. Consequently, the longer you iterate, the more accurate your image will be.

For some complex numbers c the associated Julia set forms a connected set of points in the complex plane. For others it is not. For each point c in the complex plane, if we color it black if $\text{Julia}(c)$ is connected, and color it white otherwise, we will a picture like the one shown below. This set is called the *Mandelbrot set* (see Fig. 2(b)).




Fig. 2: (a) a Julia set and (b) the Mandelbrot set.

Fractal Dimension: One of the important elements that characterizes fractals is the notion of *fractal dimension*. Fractal sets behave strangely in the sense that they do not seem to be 1-, 2-, or 3-dimensional sets, but seem to have noninteger dimensionality.

What do we mean by the *dimension* of a set of points in space? Intuitively, we know that a point is zero-dimensional, a line (or generally a curve) is one-dimensional, and plane (or generally a surface) is two-dimensional, and so on. If you put the object into a higher dimensional space (e.g., a line in 5-space) it does not change the dimensionality of the object, it is still a 1-dimensional set. If you continuously deform an object (e.g. deform a line into a circle or a plane into a sphere) it does not change its dimensionality.

How do you define the dimension of a set in general? There are various methods. Here is one, which is called *fractal dimension*. Suppose we have a set in d -dimensional space. Define a d -dimensional ε -ball to be the interior of a d -dimensional sphere of radius ε . An ε -ball is an open set (it does not contain its boundary) but for the purposes of defining fractal dimension this will not matter much. In fact it will simplify matters (without changing the definitions below) if we think of an ε -ball to be a solid d -dimensional hypercube whose side length is 2ε (an ε -square).

The dimension of an object depends intuitively on how the number of balls it takes to cover the object varies with ε . First consider the case of a line segment. Suppose that we have covered the line segment with n ε -balls. If we decrease the size of the covering balls exactly by $1/2$, it is easy to see that it takes roughly twice as many, that is, $2n$, to cover the same segment (see Fig. 3(a)).

Next, let's consider a square that has been covered with n ε -balls. If we decrease the radius by $1/2$, we see that it now takes roughly four times, that is, $4n$, as many balls (see Fig. 3(b)). Similarly, one can see that with a 3-dimensional cube, reducing the radius by a factor of $1/2$ increasing the number of balls needed to cover by a factor of $2^3 = 8$. While this is easiest to see for cubes, it generally holds (in the limit) for any compact "solid" object.

This suggests that the nature of a d -dimensional object is that the number of balls of radius ε that are needed to cover this object grows as $(1/\varepsilon)^d$. To make this formal, given an object

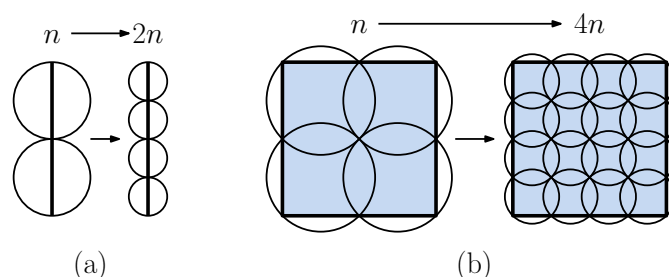


Fig. 3: The growth rate of covering numbers and fractal dimension.

X in d -dimensional space, define

$$N(X, \varepsilon) = \text{smallest number of } \varepsilon\text{-balls needed to cover } X.$$

(It will not be necessary to the absolute minimum number, as long as we do not use more than a *constant factor* times the minimum number.) We claim that an object X has dimension d if $N(X, \varepsilon)$ grows as $c(1/\varepsilon)^d$, for some constant c . This applies in the limit, as ε tends to 0. How do we extract this value of d ? Observe that if we compute $\ln N(X, \varepsilon)$ (any base logarithm will work) we get $\ln c + d \ln(1/\varepsilon)$. As ε tends to zero, the constant term c remains the same, and the $d \ln(1/\varepsilon)$ becomes dominant. If we divide this expression by $\ln(1/\varepsilon)$ we will extract the d .

Thus we define the *fractal dimension* of a set X to be

$$d = \lim_{\varepsilon \rightarrow 0} \frac{\ln N(X, \varepsilon)}{\ln(1/\varepsilon)}.$$

Formally, a set is said to be a *fractal* if:

- (i) it is *self-similar* (at different scales)
- (ii) it has a *noninteger fractal dimension*

The Sierpinski Triangle: Let's try to apply this to a more interesting object. Consider the triangular set X_0 shown in the upper right of Fig. 4. To form X_1 , we scale X_0 by $1/2$, and place three copies of it within the outline of the original set. To form X_2 , we scale X_1 by $1/2$, and place three copies of it as before. Let $X^* = \lim_{i \rightarrow \infty} X_i$. This limit is called the *Sierpinski triangle*.

In order to compute the fractal dimension of X^* , let's see how many ε -balls does it take to cover this figure. It takes one 1-ball to cover X_0 , three $(1/2)$ -balls to cover X_1 , nine $(1/4)$ -balls to cover X_2 , and in general 3^k , $(1/2^k)$ -balls to cover X_k . Letting $\varepsilon = 1/2^k$, it follows that $N(X_k, 1/2^k) = 3^k$. Thus, the fractal dimension of the Sierpinski triangle is

$$\begin{aligned} d &= \lim_{\varepsilon \rightarrow 0} \frac{\ln N(X, \varepsilon)}{\ln(1/\varepsilon)} = \lim_{k \rightarrow \infty} \frac{\ln N(X_k, (1/2^k))}{\ln(1/(1/2^k))} \\ &= \lim_{k \rightarrow \infty} \frac{\ln 3^k}{\ln 2^k} = \lim_{k \rightarrow \infty} \frac{k \ln 3}{k \ln 2} = \lim_{k \rightarrow \infty} \frac{\ln 3}{\ln 2} = \frac{\ln 3}{\ln 2} \approx 1.58496 \dots \end{aligned}$$

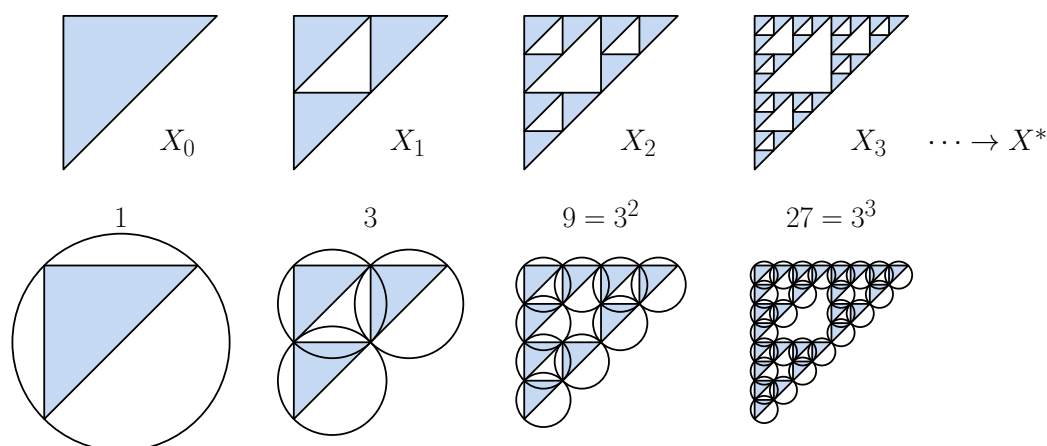


Fig. 4: The Sierpinski triangle.

Thus although the Sierpinski triangle resides (or has been embedded) in 2-dimensional space, it is essentially a $1.58\dots$ dimensional object, with respect to fractal dimension.

Although the above derivation is general, it is often easier to apply the following formula for fractals made through repeated subdivision. Suppose we form an object by repeatedly replacing each “piece” of size x by b nonoverlapping pieces each of size x/a each. Then the fractal dimension will be

$$d = \frac{\ln b}{\ln a}.$$

The Koch Island: As another example, consider the limit of the shapes K_0, K_1, \dots shown in Fig. 4. We start with a square, and with each iteration we replace each line segment of length x by a chain of 8 subsegments each of length $x/4$. The limiting shape is called the *Koch Island*. Note that the area does not change with each iteration, since for each “outward bump” there is a matching “inward bump.” However, the perimeter doubles with each iteration, and hence tends to infinity in the limit. The object itself is of fractal dimension 2, but the object’s *boundary* has fractal dimension

$$\frac{\ln 8}{\ln 4} = 1.5.$$

Since this is not an integer, the boundary of the Koch Island is a fractal.

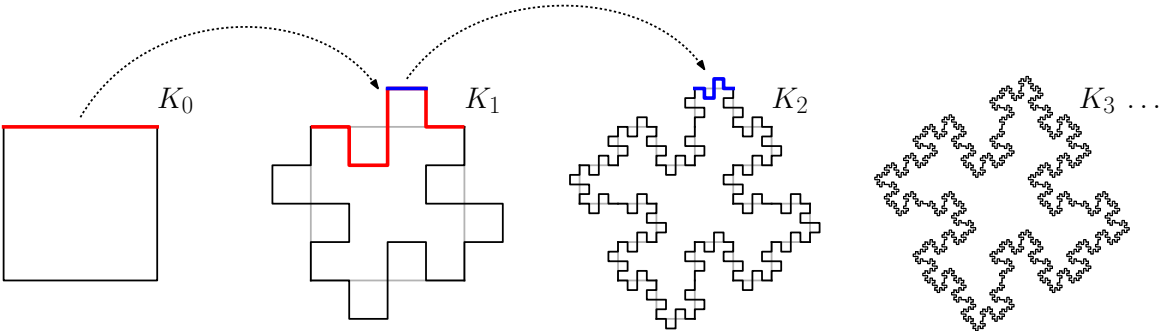


Fig. 5: The Koch Island.