Fractals

Koch Curves as Attractors and Repellers

Przemyslaw Prusinkiewicz and Glen Sandness
University of Regina

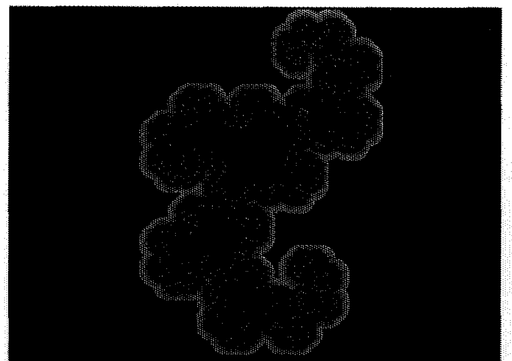

This article presents two methods for generating Koch curves, analogous to the commonly used iterative methods for producing images of Julia sets. The attracting method is based on a characterization of Koch curves as the smallest nonempty sets closed with respect to a union of similarities on the plane. This characterization was first studied by Hutchinson.[1] The repelling method is in principle dual to the attracting one, but involves a nontrivial problem of selecting the appropriate transformation to be applied at each iteration step. Both methods are illustrated with a number of computer-generated images. The mathematical presentation emphasizes the relationship between Koch construction and formal languages theory.

An earlier version of this article appeared in *Graphics Interface 88*.

In recent years the beauty of fractals has attracted wide interest among mathematicians, computer scientists, and artists. Several techniques for generating fractal shapes were developed and used to produce fascinating images. Two techniques, popularized by Mandelbrot's book,[2] have gained a particular popularity. These are the Koch construction, and function iteration in the complex domain. According to Mandelbrot's generalization, the Koch construction consists of recursively replacing edges of an arbitrary polygon (called the initiator) by an open polygon (the generator), reduced and displaced so as to have the same endpoints as those of the interval being replaced. (The original construction[3] was limited to the definition of the now famous "snowflake" curve.) As pointed out by A.R. Smith,[4] this is a language-theoretic approach: the fractal is generated by a rewriting system (a "grammar") defined in the domain of geometric shapes. In contrast, the method of function iteration refers to notions of complex analysis. The main idea is to analyze sequences of numbers $\{x_n\}$ generated by the formula $x_{n+1} = f(x_n)$, where f is a complex function. The fractal, called a Julia set, is a set invariant with respect to f. Sequences of points originating outside the fractal may gradually approach it—in which case the Julia set is said to be an attractor of the process f—or they may diverge from the fractal, and the set is then called a repeller of f. A discussion of fractal generation techniques using attracting and repelling processes has been presented by Peitgen and Richter,[5] among others.

According to the above descriptions, the methods for generating Koch curves and Julia sets appear totally unrelated to each other. But is this the case indeed? From the theoretical point of view, an answer to this question was given by Hutchinson.[1] He studied sets closed under a union of contraction maps on the plane (specifically, similarities), showed their fractal character, proved that

0272-1716/88/1100-0026$01.00 © 1988 IEEE

they can be considered as attractors, and indicated the relationship between these sets and Koch curves. Our article applies Hutchinson's theory to computer graphics.

We present two algorithms for generating images of Koch curves. The *attracting* method is similar to a method for generating images of Julia sets termed the inverse iteration method (IIM) by Peitgen and Richter.[5] An image is obtained by plotting consecutive points attracted by the fractal. This method is relatively fast and particularly useful when studying the impact of parameter changes on the curve shapes. Numerical parameter modifications make it easy to generate new variants of known curves. Continuous parameter changes allow for animating transformations of Koch curves in a way similar to the transformations of Julia sets.[6] On the other hand, the *repelling* method makes it possible to obtain colorful images of the entire plane containing a Koch curve. This method is analogous to the method for creating colorful images of Julia sets. However, in the case of Koch curves a specific new problem occurs. There are a number of similarities involved in the iteration process, and only one should be applied at each iteration step. The problem is to select the correct transformation.

Our article extends Hutchinson's results in three directions:

- We analyze the relationship between Koch construction and iteration of similarities in a formal way, based on a definition of the Koch construction in terms of formal languages theory. Our analysis is not restricted to the limit Koch curves, but also includes their finite approximations.
- In addition to the attracting algorithm for generating images of Koch curves, which is a straightforward consequence of Hutchinson's paper, we introduce a repelling algorithm.
- We illustrate both algorithms on a number of examples using computer-generated images.

The article is organized as follows. The next section presents a formal definition of the Koch construction expressed in terms of formal languages theory. Then we show the equivalence between the Koch construction and iteration of a set of similarities on the plane. The discussion is limited here to curves that can be constructed in a finite number of steps. An extension to infinite-order curves is presented next. We recall the standard notion of the topological limit of a sequence of sets and apply this notion to define limit Koch curves and provide their algebraic characterization. The corresponding method for generating approximations of limit Koch curves is presented with examples of fractal images. Finally we introduce a dual description of the limit Koch curves, which characterizes them as repellers rather than attractors. The resulting method for generating limit Koch curves is also discussed and illustrated.

The Koch construction

To accurately state and prove theorems related to the Koch construction, we must substitute a formal definition for the intuitive descriptions usually presented in the literature. A fundamental notion is that of a vector, specified as an ordered pair (x,y) of points in the plane. (Note that throughout this article the symbols x, y, z refer to points rather than coordinates.) Unless stated otherwise, we operate on *fixed* vectors, which means that two vectors $\vec{a}=(x_1, y_1)$ and $\vec{b}=(x_2,y_2)$ are considered equal if and only if their respective endpoints coincide: $x_1=x_2$ and $y_1=y_2$. (In contrast, *free* or *abstract* vectors are considered equal if they can be made to coincide by a translation.) As usual, it is convenient to identify a vector (a pair of points) with its graphical representation (a line segment in the plane). Consequently, we write that a point x belongs to a vector $\vec{a}$ if x belongs to the line segment representing $\vec{a}$. This convention extends to sets of vectors. Thus, we assume that point x belongs to a set $\{\vec{p}_1,\ldots,\vec{p}_n\}$ when x belongs to the figure formed as the union of the line segments of the component vectors.

Definition 1. A *polyvector* is an ordered set of vectors in a plane. We write $A=\{\vec{a}_1,\ldots,\vec{a}_n\}$, or $A=\vec{a}_1\cdots\vec{a}_n$ in short. Given the plane, we denote by W and W^* the set of all vectors and the class of all polyvectors, respectively.

Definition 2. A *Koch system* is a pair $K=\langle I,P\rangle$ where $I=\vec{o}_1\cdots\vec{o}_l\in W^*$ is called the *axiom* or *initiator*, and $P=(\vec{p}, \vec{q}_1\cdots\vec{q}_m)\in W\times W^*$ is called the *production*. To specify a production, we use the notation $\vec{p}\rightarrow\vec{q}_1\cdots\vec{q}_m$. The vector $\vec{p}$ is called the *predecessor* of production P, and the polyvector $\vec{q}_1\cdots\vec{q}_m$ is called the *successor* or *generator*.

Remark 1. Definition 2 extends Mandelbrot's description of the Koch construction[2] in three directions:

- The basic elements of the construction are vectors, not line segments.
- The initiator and the production successor are arbitrary sets of vectors. They need not be of equal length, form a polygon, or even be connected.
- The predecessor of the production is an arbitrary vector. It need not be connected to the successor.

The above extensions have the following justification:

- Vector orientation plays an essential role in the Koch construction. Two Koch systems that differ only by the orientation of vectors in the initiator and/or production may generate totally different fractals. Thus, a definition of a Koch system that makes no reference to line orientation is incomplete.

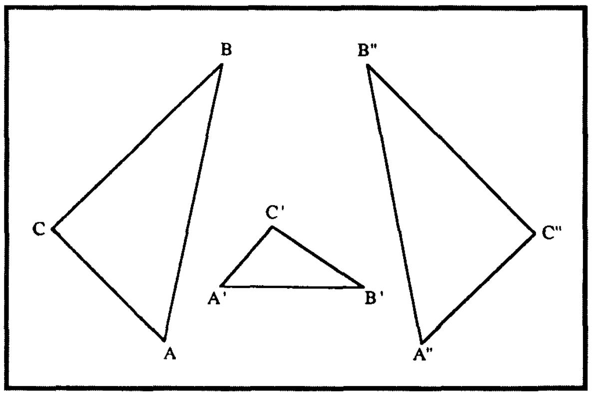

Figure 1. Illustration of the notion of direct similarity. Triangle A′B′C′ is related to ABC by a direct similarity, while triangle A″B″C″ is not directly similar to ABC because the mapping of ABC to A″B″C″ involves a reflection.

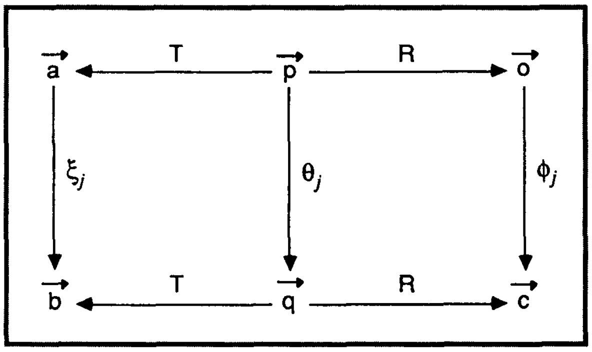

Figure 2. The relationship between mappings θ_j, ϕ_j, ξ_j, R, and T.

- When describing the construction of some fractals—for example, the dragon curve and the Gosper curve—Mandelbrot complements the specification of the initiator and the generator with additional rules of application. These rules require the starting point and the endpoint of the generator to exchange their roles in some derivation steps. By expressing productions in terms of vectors instead of line segments it is possible to incorporate the rules of application into the formal definition of the Koch system.
- Interesting modifications of fractal shapes can be obtained by allowing the vectors in the generator to be of different lengths.

In the following definitions we will refer to the notion of *direct similarity*. A direct similarity is a transformation on the plane that may change the position and size of geometric figures, but preserves their shape and orientation (which can be either clockwise or counterclockwise), as shown in Figure 1. Such similarity can be expressed as a composition of scaling, rotation, and translation; no reflections are allowed.

If a transformation T takes a figure A to the figure B, we will write $AT = B$.

Definition 3. Let $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ be the production of a Koch system K. Consider an arbitrary vector $\vec{a}$ and denote by T the direct similarity that takes vector $\vec{p}$ to the vector $\vec{a}$: $\vec{p}T = \vec{a}$. (Obviously, T is unique.) We will say that polyvector $\vec{b}_1 \cdots \vec{b}_m$ is *directly derived* from the vector $\vec{a}$ and write $\vec{a} \geqslant \vec{b}_1 \cdots \vec{b}_m$ if and only if $\vec{b}_1 \cdots \vec{b}_m = (\vec{q}_1 \cdots \vec{q}_m)T$.

Remark 2. In the case of rewriting systems that operate on strings (for example, context-free grammars), the result of applying a production $p \to q_1 \cdots q_m$ to the letter p is identical with the production's successor: $q_1 \cdots q_m$. Consequently, there is no need for distinguishing between production $p \to q_1 \cdots q_m$ and the derivation $p \geqslant q_1 \cdots q_m$. In contrast, in a Koch system the result $\vec{b}_1 \cdots \vec{b}_m$ of applying production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ to a vector $\vec{a}$ is, in general, different from the successor $\vec{q}_1 \cdots \vec{q}_m$ (since, in general, $\vec{a} \neq \vec{p}$).

Corollary 1. Consider a Koch system K with production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$, and let $\vec{a} \to \vec{b}_1 \cdots \vec{b}_m$ be a derivation in K. Denote by θ_j the direct similarity that takes vector $\vec{p}$ to the vector $\vec{q}_j$: $\vec{p}\theta_j = \vec{q}_j$ $(j = 1, \ldots, m)$. In an analogous way, denote by ξ_j the similarity that takes vector $\vec{a}$ to the vector $\vec{b}_j$: $\vec{a}\xi_j = \vec{b}_j$. If T is the similarity that takes vector $\vec{p}$ to $\vec{a}$, then $\xi_j = T^{-1}\theta_j T$.

Proof. According to Definition 3, if $\vec{a} = \vec{p}T$, then $\vec{b}_j = \vec{q}_j T$. Thus,

$$\vec{p}(T\xi_j) = (\vec{p}T)\xi_j = \vec{a}\xi_j = $$
$$\vec{b} = \vec{q}_j T = (\vec{p}\theta_j)T = \vec{p}(\theta_j T)$$

or $\xi_j = T^{-1}\theta_j T$. □

Remark 3. In the following sections we will focus on Koch systems with the axiom limited to a single vector $\vec{o}$. In this case the derivation $\vec{o} \geqslant \vec{c}_1 \cdots \vec{c}_m$ starting from the axiom $\vec{o}$ plays a particular role which justifies the use of special symbols R and ϕ in place of T and ξ. Thus, by definition, $\vec{p}R = \vec{o}$ and $\vec{o}\phi_j = \vec{c}_j$. The relationship between different vectors and transformations discussed above is represented diagrammatically in Figure 2. Note that the mappings θ_j, ϕ_j, and R are completely defined by the Koch system K, while the mappings ξ_j and T vary from one argument vector $\vec{a}$ to another.

The next definition extends the notion of direct derivation to the predecessors that are not single vectors.

Definition 4. Let $\vec{a}_1 \cdots \vec{a}_l$ be a polyvector and $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ the production of a Koch system K. The polyvector $\vec{b}_{11} \cdots \vec{b}_{1m} \cdots \vec{b}_{l1} \cdots \vec{b}_{lm}$ is *directly derived* from the polyvector $\vec{a}_1 \cdots \vec{a}_l$ in the system K if and only if $\vec{a}_i \geqslant \vec{b}_{i1} \cdots \vec{b}_{im}$ for all $i = 1, \ldots, l$. We write

$$\vec{a}_1 \cdots \vec{a}_l \Rightarrow \vec{b}_{11} \cdots \vec{b}_{1m} \cdots \vec{b}_{l1} \cdots \vec{b}_{lm}.$$

Remark 4. Note that in the derivation

$$\vec{a}_1 \cdots \vec{a}_l \Rightarrow \vec{b}_{11} \cdots \vec{b}_{1m} \cdots \vec{b}_{l1} \cdots \vec{b}_{lm}$$

all vectors $\vec{a}_i$ ($i = 1, \ldots, l$) are substituted by their successors in a single derivation step. Consequently, Koch systems belong to the class of parallel rewriting systems. In the domain of strings, the analogous derivation type characterizes L-systems.[7,8] The relationship between Koch systems and L-systems is quite close; in fact, many Koch curves can be generated using L-systems with a geometric interpretation of string symbols.[9-12] However, a discussion of the formal aspects of this relationship is beyond the scope of this article.

Definition 5. The notion of the direct derivation is extended to the *derivation of length* $n \geq 0$ in the usual recursive way:

- For any polyvector C, $C \geqslant^0 C$.
- If $C_0 \geqslant^n C_n$ and $C_n \geqslant C_{n+1}$, then $C_0 \geqslant^{n+1} C_{n+1}$.

Definition 6. A polyvector C_n is the *Koch curve of order* n generated by a Koch system $K = \langle I, P \rangle$ if C_n is derived in K from the axiom I in a derivation of length n: $I \geqslant^n C_n$.

Finite-order Koch curves

This section presents a characterization of Koch curves in terms of algebra of relations. We show that any Koch system K corresponds to a geometric relation Φ in such a way that the Koch curve of order n generated by K can be represented as $I\Phi^n$. The formal discussion is limited to the Koch systems in which the initiator I is a single vector. A method for removing this limitation is outlined in the final section of the article.

Theorem 1. Consider a Koch system $K = \langle \vec{o}, \vec{p} \to \vec{q}_1 \cdots \vec{q}_m \rangle$. For any sequence of indices $j_1, \ldots, j_n : j_i \in \{1, \ldots, m\}$, the following equality holds:

$$\vec{o}\,\xi_{j_1} \cdots \xi_{j_n} = \vec{o}\,\phi_{j_n} \cdots \phi_{j_1}$$

where mappings ξ_j and ϕ_j are defined as in Corollary 1 and Remark 3. The operation ξ_j is assumed to be left-associative: $\vec{o}\,\xi_{j_1} \cdots \xi_{j_n} = (\ldots(\vec{o}\,\xi_{j_1}) \cdots \xi_{j_n})$.

Proof—by induction on n.

- Assuming that the sequence of zero transformations is equal to the identity mapping, for $n = 0$ the thesis is obvious.

- Assume the thesis true for an $n \geq 1$ and consider a vector $\vec{c} = \vec{o}\,\xi_{j_1} \cdots \xi_{j_n} \xi_{j_{n+1}}$. According to the inductive assumption, the vector $\vec{a} = \vec{o}\,\xi_{j_1} \cdots \xi_{j_n}$ can be expressed as $\vec{a} = \vec{o}\,\eta$, where $\eta = \phi_{j_n} \cdots \phi_{j_1}$. Furthermore, from Corollary 1 it follows that the vector $\vec{b} = \vec{a}\,\xi_{j_{n+1}}$ can be expressed as $\vec{a}\,T^{-1}\theta_{j_{n+1}}T$, where T is the direct similarity that takes the production predecessor $\vec{p}$ to the vector $\vec{a}$. The transformation T is in turn equal to the composition of the direct similarity R, which takes the predecessor $\vec{p}$ to the axiom $\vec{o}$, and the transformation η, which takes axiom $\vec{o}$ to the vector $\vec{a}$. Consequently, we obtain

$$\vec{b} = \vec{a}\,\xi_{j_{n+1}} = \vec{a}\,T^{-1}\theta_{j_{n+1}}T = \vec{a}\,(\eta^{-1}R^{-1})\theta_{j_{n+1}}(R\eta) =$$

$$(\vec{a}\,\eta^{-1})(R^{-1}\theta_{j_{n+1}}R)\eta = \vec{o}\,\phi_{j_{n+1}}\eta = \vec{o}\,\phi_{j_{n+1}}\phi_{j_n} \cdots \phi_{j_1} \quad \square$$

Interpretation. According to the above theorem, associated with a Koch system K is a set of direct similarities ϕ_j. A vector $\vec{b}$ can be derived from the axiom in a sequence of production applications if and only if it can also be obtained by transforming the axiom vector using a sequence of direct similarities ϕ_j. The similarities ϕ_j must be applied in the reversed order compared with the corresponding ξ_j mappings.

Example 1. To illustrate Theorem 1, let us introduce the following notation:

- $S(a)$ is scaling with respect to the origin of the coordinate system where parameter $a > 0$ is the scaling ratio.
- $R(\alpha)$ is rotation by angle α with respect to the origin of the coordinate system.
- $M(u,v)$ is translation by vector (u,v). (Note that here u and v are coordinates of a free vector, not endpoints of a fixed vector.)

The similarities corresponding to the Koch system presented in Figure 3a can be expressed as follows:

$$\phi_1 = S(\tfrac{1}{3})$$

$$\phi_2 = S(\tfrac{1}{3})\, R(\tfrac{\pi}{3})\, M(\tfrac{1}{3}, 0)$$

$$\phi_3 = S(\tfrac{1}{3})\, R(-\tfrac{\pi}{3})\, M(\tfrac{1}{2}, \tfrac{\sqrt{3}}{6})$$

$$\phi_4 = S(\tfrac{1}{3})\, M(\tfrac{2}{3}, 0)\,.$$

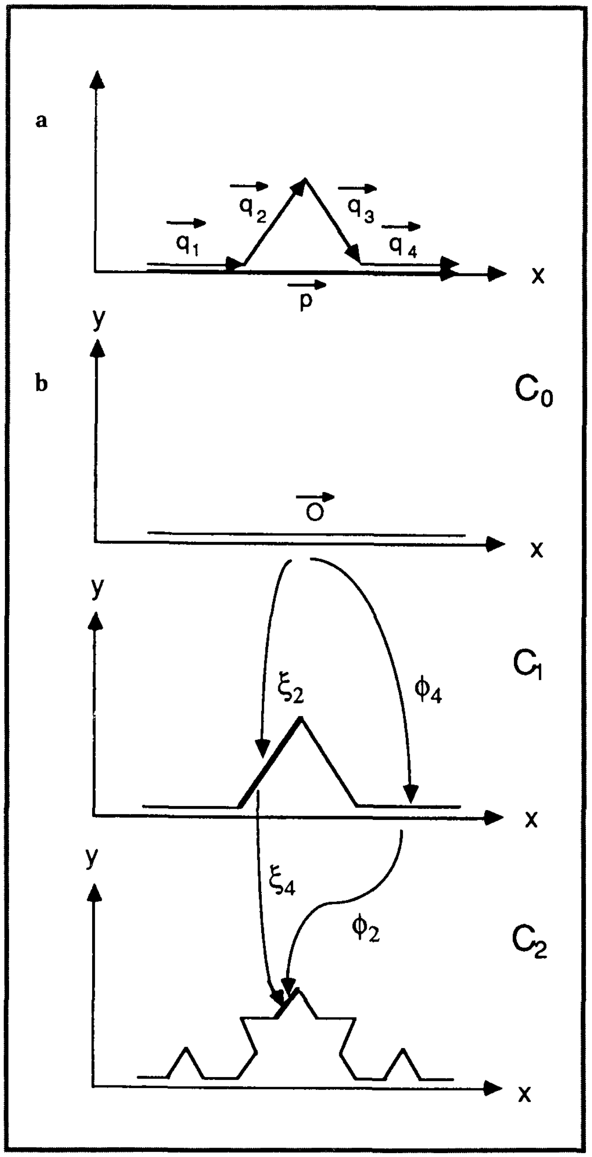

Figure 3. (a) The production of a Koch system. (b) Two methods for obtaining a vector $\vec{b} \in C_2$: a sequence of productions and a sequence of similarities.

Figure 3b shows that a vector $\vec{b} \in C_2$ can be derived from the axiom using mappings $\xi_2\xi_4$, or obtained as the image of the axiom using similarities $\phi_4\phi_2$. Operations ϕ_j are applied in the reversed order compared with the corresponding operations ξ_j.

Remark 5. The specification of similarities ϕ_j by a composition of more primitive operations has an intuitive geometric appeal—it is conceptually close to the specification of symmetries in terms of rotations, translations, reflections, and glide reflections. This emphasizes the relationship between fractal and "classic" geometry: Koch curves can be perceived as symmetric patterns that admit similarities as symmetries. The concept of considering similarities as symmetries is certainly not new. The extensive study of "patterns and tilings" by Grunbaum and Shephard provides several examples of so-called "similarity patterns" obtained by overlaying smaller and smaller copies of a given motif.[13] However, all these patterns use exactly one similarity. The possibility of generating many interesting patterns using two or more similarities went unnoticed there.

Since each sequence of n similarities ϕ_i takes the axiom $\vec{a}$ to a vector $\vec{b}$ that belongs to the Koch curve C_n, *and each vector of C_n corresponds to some sequence* of such transformations, the following corollary holds.

Corollary 2. Consider a Koch system $K = \langle \vec{a}, \vec{p} \to \vec{q_1} \cdots \vec{q_m} \rangle$, and let Φ denote the union of the similarities ϕ_j associated with K:

$$\Phi = \bigcup_{j=1}^{m} \phi_j \ .$$

For any $n = 0,1,2,\ldots$ the Koch curve of order n generated by K can be expressed as $C_n = \vec{a}\Phi^n$.

Interpretation. According to the above corollary, a Koch curve of order n can be obtained recursively, starting from $C_0 = \vec{a}$ and using the formula $C_{i+1} = C_i\Phi$ to progress through the sequence of Koch curves of consecutive orders. Note that in general the relation Φ is not monotonic; i.e., $C_i\Phi$ is *not* a superset of C_i. Consequently, the curve C_{i+1} cannot be obtained simply by adding new vectors to C_i. Some if not all vectors of C_i must also be erased.

In the following sections we will introduce the notion of a limit Koch curve and we will show that, by operating on points instead of vectors, it is possible to generate the limit Koch curves in a monotonic process with no erasing.

The topological limit of a sequence of sets

In this section we recall some basic topological notions.[1,14] To start with, let us assume that all sets considered are closed sets on the plane **P**.

Definition 7. Let $\varrho(x,y)$ denote the Euclidean distance between points x,y. The distance between point x and set Y is defined as

$$\rho(x,Y) = \inf_{y \in Y} \rho(x,y)$$

The *half-distance* between set X and set Y is equal to

$$\rho'(X,Y) = \sup_{x\in X} \rho(x,Y) \quad .$$

Note that, in general, $\varrho'(X,Y) \neq \varrho'(Y,X)$. The *distance* between sets X and Y is the greater of the two half-distances:

$$\rho(X,Y) = \max \{\rho'(X,Y), \rho'(Y,X)\} \quad .$$

The function $\varrho(X,Y)$ satisfies the distance axioms in the space of all closed nonempty subsets of the plane **P** and is called the *Hausdorff metric* on this space. Note that for any set families $X_1,\ldots,X_m$ and $Y_1,\ldots,Y_n$ the following inequality holds:

$$\rho \left[\bigcup_{i=1}^{m} X_i \, , \bigcup_{j=1}^{n} Y_j \right] \leq \qquad (1)$$

$$\max \{\rho(X_i, Y_j) : 1 \leq i \leq m, \; 1 \leq j \leq n\}$$

Definition 8. A set A such that

$$\lim_{n\to\infty} \rho(A_n, A) = 0$$

is called the *topological limit* of the sequence of sets $A_0, A_1, A_2, \cdots$. It is known that if a topological limit exists, it is unique. Consequently, we can use notation $A = \mathrm{Lt} A_n$.

Definition 9. Consider a function f: $\mathbf{P} \to \mathbf{P}$. The *Lipschitz constant* of f is defined as

$$\mathrm{Lip}\ (f) = \sup_{x\neq y} \frac{\rho(f(x), f(y))}{\rho(x,y)} \quad .$$

We will use the following properties of Lip (f):

- For any points $x,y \in \mathbf{P}$

$$\rho(f(x), f(y)) \leq \mathrm{Lip}\ (f)\ \rho(x,y)$$

- If f: $\mathbf{P} \to \mathbf{P}$ and g: $\mathbf{P} \to \mathbf{P}$, then

$$\mathrm{Lip}\ (fg) \ = \mathrm{Lip}\ (f)\ \mathrm{Lip}\ (g)$$

- If f is a similarity, then

$$\mathrm{Lip}\ (f^{-1}) = \frac{1}{\mathrm{Lip}\ (f)} \qquad (2)$$

A function f is called a *contraction* if $\mathrm{Lip}(f) < 1$.

The limit Koch curves

This section characterizes limit Koch curves as sets invariant with respect to unions of similarities and is based on the work of Hutchinson.[1]

Definition 10. The *spread* σ of a Koch system K is the distance between the axiom $\vec{\sigma}$ and its direct successor $\vec{c}_1 \cdots \vec{c}_m$:

$$\sigma = \rho(\vec{\sigma}, \vec{c}_1 \cdots \vec{c}_m)$$

where

$$\vec{\sigma} \Rightarrow \vec{c}_1 \cdots \vec{c}_m \quad .$$

Lemma 1. The distance between two consecutive curves C_n and C_{n+1} generated by a Koch system K with production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ satisfies the inequality

$$d_n \leq \sigma \frac{\mathrm{length}(\vec{a}_{\max})}{\mathrm{length}(\vec{\sigma})}$$

where $\vec{a}_{\max}$ is the longest vector in the polyvector C_n.

Proof. Consider derivation $\vec{a} \Rightarrow \vec{b}_1 \cdots \vec{b}_m$, which results from the application of production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ to a vector $\vec{a}$. According to Figure 2, the vectors $\vec{a}, \vec{b}_1, \ldots, \vec{b}_m$ are related to the vectors $\vec{\sigma}, \vec{c}_1, \ldots, \vec{c}_m$ by a similarity $R^{-1}T$, hence

$$\frac{\rho(\vec{a}, \vec{b}_1 \cdots \vec{b}_m)}{\rho(\vec{\sigma}, \vec{c}_1 \cdots \vec{c}_m)} = \frac{\mathrm{length}(\vec{a})}{\mathrm{length}(\vec{\sigma})} \quad .$$

The longer the vector $\vec{a}$ is, the bigger the value of both ratios. Taking into account Inequality 1 from the previous section, we obtain

$$\rho(C_n, C_{n+1}) \leq \max \{\rho(\vec{a}_j, \vec{b}_{j_1} \cdots \vec{b}_{j_m}) : \vec{a}_j \in C_n \;\&\; \vec{a}_j \Rightarrow \vec{b}_{j_1} \cdots \vec{b}_{j_m}\} = \sigma \frac{\mathrm{length}(\vec{a}_{\max})}{\mathrm{length}(\vec{\sigma})} \quad \square$$

Definition 11. The *contraction ratio* of a production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ is defined as

$$\gamma = \frac{\mathrm{length}(\vec{q}_{\max})}{\mathrm{length}(\vec{p})}$$

where $\vec{q}_{\max}$ is the longest vector of the generator $\vec{q}_1 \cdots \vec{q}_m$.

Lemma 2. Assuming the notation of Corollary 1 and Remark 3, the following equality holds:

$$\gamma = \max\{\mathrm{Lip}(\theta_j): 1 \le j \le m\} = \max\{\mathrm{Lip}(\phi_j): 1 \le j \le m\} = \max\{\mathrm{Lip}(\xi_j): 1 \le j \le m\} \ .$$

Proof. The equality $\gamma = \max\{\mathrm{Lip}(\theta_j): 1 \le j \le m\}$ results directly from Definition 11. Furthermore, taking into account Equation 2 in the previous section, we obtain

$$\mathrm{Lip}(\phi_j) = \mathrm{Lip}(R^{-1}\theta_j R) = \mathrm{Lip}(\theta_j) \ .$$

Using the same argument for ξ_j, we conclude that

$$\mathrm{Lip}(\theta_j) = \mathrm{Lip}(\phi_j) = \mathrm{Lip}(\xi_j)$$

for any $j \in \{1,\ldots,m\}$, so the thesis holds.□

Lemma 3. The length of any vector $\vec{b}$ in the polyvector C_n satisfies the inequality

$$\mathrm{length}(\vec{b}) \le \mathrm{length}(\vec{\sigma})\gamma^n$$

Proof. According to Theorem 1, if $\vec{b} \in C_n$, then there exists a sequence of n transformations $\phi_{i_n} \cdots \phi_{i_1}$ such that $\vec{b} = \vec{\sigma}\phi_{i_n} \cdots \phi_{i_1}$. From Lemma 2 it follows that $\mathrm{Lip}(\phi_i) \le \gamma$ for all functions ϕ_i under consideration. Consequently, $\mathrm{length}(\vec{b}) \le \mathrm{length}(\vec{\sigma})\gamma^n$.□

Definition 12. Consider a sequence of polyvectors C_n generated by a Koch system K using derivations of length $0,1,2,\ldots$. A set $C_\infty = \mathrm{Lt}C_n$ is called the *limit curve* generated by the system K.

Theorem 2. Consider a Koch system $K = \langle \vec{a}, \vec{p} \to \vec{q}_1 \cdots \vec{q}_m \rangle$. If the contraction ratio γ of the production $\vec{p} \to \vec{q}_1 \cdots \vec{q}_m$ is less than 1, then the limit curve C_∞ exists and is bounded.

Proof. Consider a sequence $C_n, C_{n+1}, \ldots, C_p$ of polyvectors generated by the Koch system K. According to Lemmas 1 and 3, the distance d_i between consecutive polyvectors C_i and C_{i+1} satisfies the inequality $d_i \le \sigma\gamma^i$. The distance between polyvectors C_n and C_p does not exceed the sum of distances $d_n + d_{n+1} + \cdots + d_{p-1}$:

$$\rho(C_n, C_p) \le \sum_{i=n}^{p-1} \sigma\gamma^i = \sigma\gamma^n \frac{1-\gamma^{p-n}}{1-\gamma} \ .$$

Since $\gamma < 1$ and $p > n$, we obtain

$$\rho(C_n, C_p) \le \sigma\gamma^n \frac{1}{1-\gamma} \ .$$

The above formula shows that the distance $\varrho(C_n, C_p)$ tends to zero with $n \to \infty$; hence according to the Cauchy criterion there exists the limit set C_∞ such that

$$\lim_{n\to\infty} \rho(C_n, C_\infty) = 0 \ .$$

Or $C_\infty = \mathrm{Lt}C_n$. Furthermore,

$$\rho(C_0, C_\infty) \le \sigma \frac{1}{1-\gamma} \ ,$$

so C_∞ is bounded.□

Theorem 3. Consider a Koch system $K = \langle \vec{a}, \vec{p} \to \vec{q}_1 \cdots \vec{q}_m \rangle$. The contraction ratio γ is assumed to be less than 1. Let Φ denote the union of the similarities ϕ_i associated with K:

$$\Phi = \bigcup_{j=1}^{m} \phi_j \ .$$

The limit curve C_∞ generated by K has the following properties:

1. $C_\infty \Phi = C_\infty$.
2. For any nonempty set X on the plane, if $X\Phi \subset X$, then $C_\infty \subset X$.
3. For any point x in the plane

$$\lim_{n\to\infty} \rho'(x\Phi^n, C_\infty) = 0 \ .$$

Proof.

1. $$C_\infty = \mathop{\mathrm{Lt}}_{n\to\infty} \vec{a}\Phi^n = \mathop{\mathrm{Lt}}_{n\to\infty} \vec{a}\Phi^{n+1} = (\mathop{\mathrm{Lt}}_{n\to\infty} \vec{a}\Phi^n)\,\Phi = C_\infty\Phi \ .$$

2. Let X be an arbitrary nonempty set closed with respect to Φ. To show that $C_\infty \subset X$ we will consider a point $x \in X$ and a vector $\vec{b} \subset C_n$ $(n \ge 0)$. According to Theorem 1, $\vec{b}$ is the image of the axiom $\vec{\sigma}$ with respect to some sequence of transformations included in Φ^n: $\vec{b} = \vec{\sigma}\phi_{i_n}\phi_{i_{n-1}} \cdots \phi_{i_1}$. Let y denote the image of x with respect to the same sequence of transformations, $y = x\phi_{i_n}\phi_{i_{n-1}} \cdots \phi_{i_1}$, and assume that the distance between the axiom $\vec{\sigma}$ and the set $\{x\}$ is equal to d_0. According to Lemma 2, the distance between the vector $\vec{b}$ and the set $\{y\}$ satisfies the inequality $\varrho(\vec{b}, \{y\}) \le d_0\gamma^n$. Since the set X is assumed to be closed with respect to all transformations included in Φ, y belongs to X. Thus, the half-distance $d_n = \varrho'(\vec{b}, X)$ is less than or equal to $d_0\gamma^n$. Considering that the contraction ratio γ is less than one, the half-distance d_n between an arbitrary vector $\vec{b} \subset C_n$ and the set X tends to zero when n tends to infinity. Thus, the limit set C_∞ is a subset of X.

3. Consider an arbitrary point x in the plane, and a point z that belongs to C_∞. Denote by d_0 the distance $\varrho(x,z)$. Following the same arguments as in a previous part of this proof, we obtain

$$\rho(x\phi_{j_n} \cdots \phi_{j_1}, z\phi_{j_n} \cdots \phi_{j_1}) \leq d_0\gamma^n ,$$

where $\phi_{j_n} \cdots \phi_{j_1}$ is an arbitrary sequence of transformations included in Φ^n. According to part 1 of this proof, $z\phi_{j_n} \cdots \phi_{j_1} \in C_\infty$, thus $\varrho(x\phi_{j_n} \cdots \phi_{j_1}, C_\infty) \leq d_0\gamma^n$, or $\varrho'(x\Phi^n, C_\infty) \leq d_0\gamma^n$. Considering that the contraction ratio γ is less than one, the thesis is obtained.□

Corollary 3. For any Koch system K with a contraction ratio $\gamma<1$ and any point x in the plane,

$$C_\infty = \mathrm{Lt}\, x\Phi^n .$$

Proof. Following the same argument as in the proof of part 1 in Theorem 3, we find that the set $X = \mathrm{Lt}x\Phi^n$ has the property $X\Phi = X$. Thus, according to part 2, $C_\infty \subset X$. On the other hand, from part 3 it follows that $X \subset C_\infty$. Thus, $X = C_\infty$.□

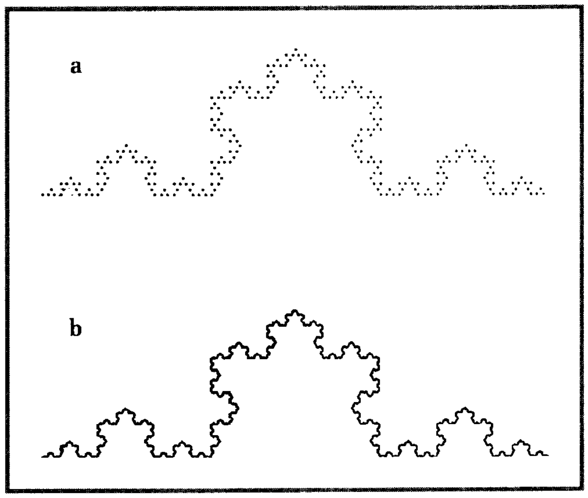

Figure 4. Two approximations of the limit snowflake curve obtained by the attracting method.

Interpretation. Parts 1 and 2 of Theorem 3 characterize the limit Koch curve C_∞ as the smallest nonempty set invariant with respect to a union of similarities. Part 3 characterizes the set C_∞ as an *attractor*. The iterative application of the transformations included in Φ can be considered as a *dynamic process*[2,5] that describes evolution of the set of points S_n in time. The process starts with a one-element set $S_0 = \{x\}$. The subsequent sets S_n get closer and closer to the limit set C_∞ regardless of the selection of the initial point x. Thus, C_∞ attracts points from the entire plane. Corollary 3 further specifies that all points of C_∞ will be reached by applying some (possibly infinite) sequences of transformations from Φ to an arbitrary starting point x.

The attracting method for Koch curve generation

Theorem 3 and Corollary 3 suggest a simple method for generating finite approximations of the limit Koch curves.

- Start from a set $S_0 = \{z\}$ where point z is known to belong to C_∞.
- Given set S_n, construct set S_{n+1} by applying all transformations $\phi_i \subset \Phi$ to all points $z_k \in S_n$. Repeat this step for consecutive values of n until the desired number of points approximating C_∞ is reached.

Note that all generated points z_k belong to the limit curve C_∞, so all calculated points contribute to the approximation of C_∞ and no erasing occurs.

The above method assumes that an initial point $z \in C_\infty$ is known. Two approaches can be used to find such a point.

- Solve for z any of the equations $z\phi_i = z$, where $\phi_i \subset \Phi$. Since each mapping ϕ_i is a contraction, it has a unique fixed point z. Furthermore, $z = z\phi_i^n$ for any $n \geq 0$; thus according to Corollary 3, $z \in C_\infty$.
- Choose an arbitrary point x in the plane, and apply to it a sequence of transformations $\phi_{i_1} \cdots \phi_{i_n} \subset \Phi^n$. According to part 3 of Theorem 3, if n is sufficiently large, the resulting point z will be arbitrarily close to the curve C_∞. Consequently, z can be used as the starting point for curve generation. Because of the attracting nature of the process Φ, the impact of the error in choosing the initial point will further decrease as the iteration continues.

Example 2. Figure 4 shows two approximations of one branch of the snowflake curve. The relation Φ used for iteration is the union of similarities ϕ_1 through ϕ_4 from Example 1.

Example 3. Figure 5 shows four curves generated by a pair of mappings

$$\phi_1(z) = z\gamma e^{i\frac{\pi}{4}} \qquad \phi_2(z) = z\gamma e^{i\frac{3\pi}{4}} + i$$

for different values of parameter γ. A point is colored red if ϕ_1 is the last transformation used; otherwise it is colored green. Modification of the numerical parameters reveals interesting variations of the basic dragon curve shape.

Remark 6. The use of complex variables emphasizes an analogy between the Koch curves and the Julia sets. For

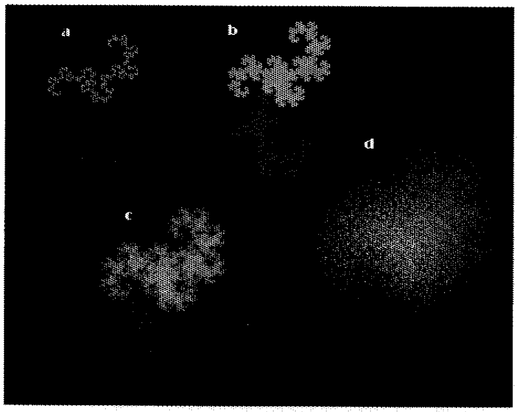

Figure 5. Variations of the dragon curve: (a) $\gamma=0.65$, (b) $\gamma=0.7071=\sqrt{2}/2$, (c) $\gamma=0.75$, (d) $\gamma=0.85$.

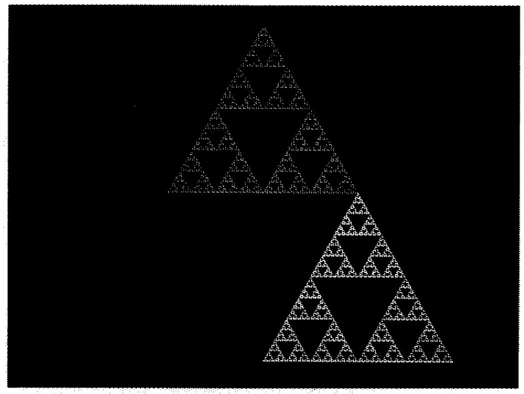

Figure 6. The Sierpinski gasket approximated using the attracting method.

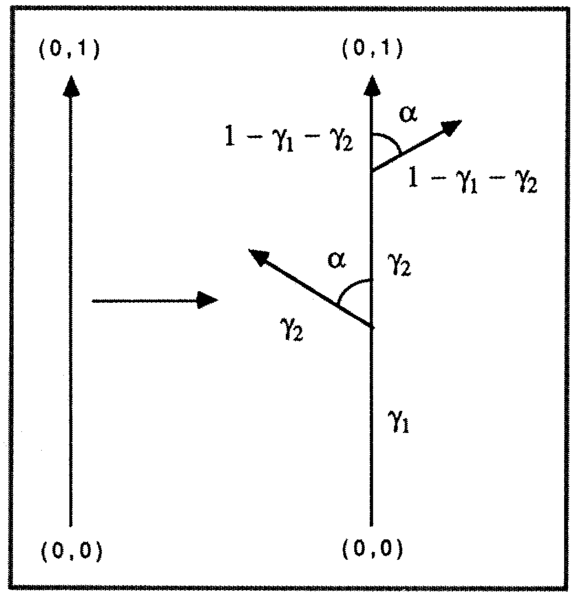

Figure 7. Production of the Koch system generating twiglike shapes.

example, a Julia set can be generated using the inverse iteration method, by iterating two mappings:

$$f_1(z) = +\sqrt{z+1} \qquad f_2(z) = -\sqrt{z+1} \quad .$$

The particular mappings generating a Koch curve and a Julia set are different, but the underlying iterative algorithm is the same.

Example 4 (based on work by Demko et al.[15]). Figure 6 shows the curve generated by the union of three transformations:

$$\phi_1(z) = \frac{1}{2}z+\frac{1}{2}i$$

$$\phi_2(z) = \frac{1}{2}z+\frac{1}{2}e^{-i\frac{5}{6}\pi}$$

$$\phi_3(z) = \frac{1}{2}z+\frac{1}{2}e^{-i\frac{1}{6}\pi} \quad .$$

As previously, the point colors indicate the last transformation used. Note that the figure obtained is the Sierpinski gasket.[2] Transformations ϕ_1 through ϕ_3 provide an interesting characterization of this well-known curve: The Sierpinski gasket is the smallest nonempty set closed with respect to three scaling transformations. Their centers (fixed points) lie at the vertices of an equilateral triangle and the scaling ratios are equal to 1/2.

Example 5. Figure 7 shows the production of a Koch system generating twiglike shapes.[12] The corresponding similarities are given below:

$$\phi_1(z) = \gamma_1 z$$

$$\phi_2(z) = \gamma_2 z + \gamma_1 i$$

$$\phi_3(z) = (1-\gamma_1-\gamma_2)z + (\gamma_1+\gamma_2)i$$

$$\phi_4(z) = \gamma_2 z e^{i\alpha} + \gamma_1 i$$

$$\phi_5(z) = (1-\gamma_1-\gamma_2)z e^{-i\alpha} + (\gamma_1+\gamma_2)i$$

Figure 8 presents the images resulting from iterating the similarities ϕ_1 through ϕ_5 for three different values of parameters γ_1 and γ_2. In all cases, $\alpha = \pi/6$. As in Example 3, modification of the numerical parameters reveals interesting shape variations.

The repelling method for Koch curve generation

The sets of equations considered in the previous section defined Koch curves as attractors of dynamic processes. In this section we address the problem of describing Koch curves as repellers. The basic concept is to use reciprocal mappings ϕ_j^{-1} instead of the functions ϕ_j. However, the repelling algorithm for Koch curve generation is more complicated than its attracting counterpart: For all points and at all iteration steps it requires a careful selection of the applicable mapping.

Theorem 4. Consider a Koch system K with the contraction ratio $\gamma < 1$ and let Φ denote, as previously, the union of similarities $\phi_1, \ldots, \phi_m$ associated with K. A point x belongs to the limit curve C_∞ if and only if there exists a function $\phi_j \subset \Phi$ such that $x\phi_j^{-1} \in C_\infty$.

Proof. According to part 1 of Theorem 3, $C_\infty = C_\infty \Phi$. Thus, for any point $x \in C_\infty$ there exists a point $y \in C_\infty$ and a transformation $\phi_j \subset \Phi$ such that $y\phi_j = x$, or $x\phi_j^{-1} = y \in C_\infty$. On the other hand, if $x\phi_j^{-1} \in C_\infty$, then $x\phi_j^{-1}\phi_j = x \in C_\infty$.□

Theorem 5. Consider a Koch system K with the contraction ratio $\gamma < 1$ and let Φ denote the union of similarities $\phi_1, \ldots, \phi_m$ associated with K. If a point x does not belong to the limit curve C_∞, then for any infinite sequence of transformations $\phi_{j_1}^{-1}\phi_{j_2}^{-1}\cdots$

$$\lim_{n \to \infty} \rho(x\phi_{j_1}^{-1} \cdots \phi_{j_n}^{-1}, C_\infty) = \infty \ .$$

Proof. Consider a sequence of similarities $\phi_{i_n} \cdots \phi_{i_1}$ and a point $v \in C_\infty$. Since the set C_∞ is closed with respect to all similarities $\phi_j \subset \Phi$ (Theorem 3, part 1), the point $u = v\phi_{i_n} \cdots \phi_{i_1}$ also belongs to C_∞. Now let us consider point $y = x\phi_{i_1}^{-1} \cdots \phi_{i_n}^{-1}$. According to Definition 9 and Lemma 2, $\varrho(x,u) \le \gamma^n \varrho(y,v)$, or

$$\rho(y, v) \ge \gamma^{-n} \rho(x, u) \ .$$

The distance $\varrho(x,u)$ is greater than zero, because $x \notin C_\infty$ and the set C_∞ is closed. Thus, $\gamma^{-n}\varrho(x,u) \to \infty$ with $n \to \infty$, and consequently $\varrho(y,v) \to \infty$. Since $v \in C_\infty$ and C_∞ is bounded (Theorem 2), the distance between $y = x\phi_{i_1}^{-1} \cdots \phi_{i_n}^{-1}$ and C_∞ tends to infinity with $n \to \infty$.□

From Theorems 4 and 5 it follows that a point x belongs to the curve C_∞ if and only if there exists an infinite sequence of transformations $\phi_{i_1}^{-1}\phi_{i_2}^{-1}\cdots$ that *does not* take x to infinity. This observation can be used as a basis for generating images of Koch curves, although only finite transformation sequences and finite distances on the plane can be considered in practice. The algorithm proceeds as follows:

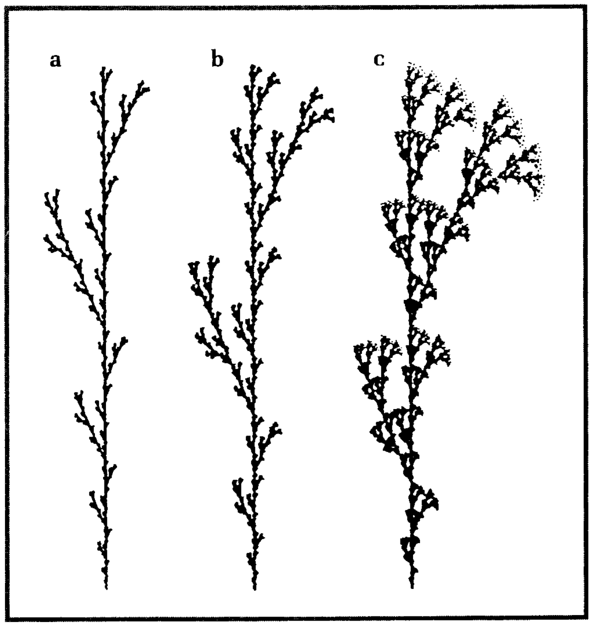

Figure 8. "Twigs" generated using the attracting method: (a) $\gamma_1 = 0.5$, $\gamma_2 = 0.3$, (b) $\gamma_1 = \gamma_2 = 1/3$, (c) $\gamma_1 = 0.2$, $\gamma_2 = 0.3$. In all three cases, $\alpha = \pi/6$.

- Define a window on the image plane to establish the area of interest within which the curve will be traced. Subdivide this window into an array of sample points that will correspond to the pixels on the screen (for example, each sample will represent one pixel if no oversampling is used.) Assume the maximum length N of the transformation sequence considered. Define a "large" circle Ω (including the curve C_∞), which will be used to test whether points tend to infinity.
- Partition the plane into regions D_j such that for any $x \in C_\infty \cap D_j$ the function $\phi_j^{-1} \subset \Phi^{-1}$ takes point x to some point of C_∞: $x\phi_j^{-1} \in C_\infty$. According to Theorem 4, at least one such function ϕ_j^{-1} exists; hence this partition is feasible.
- For each sampling point x_0 calculate a sequence of points $x_0, x_1, x_2, \ldots$ according to the rule

$$\text{if} \quad x_n \in D_j \quad \text{then} \quad x_{n+1} = x_n \phi_j^{-1} \ .$$

Stop this iteration if the index n reaches limit N or x_n falls outside of the circle Ω. Assign a color to the point x_0 according to the final value of n.

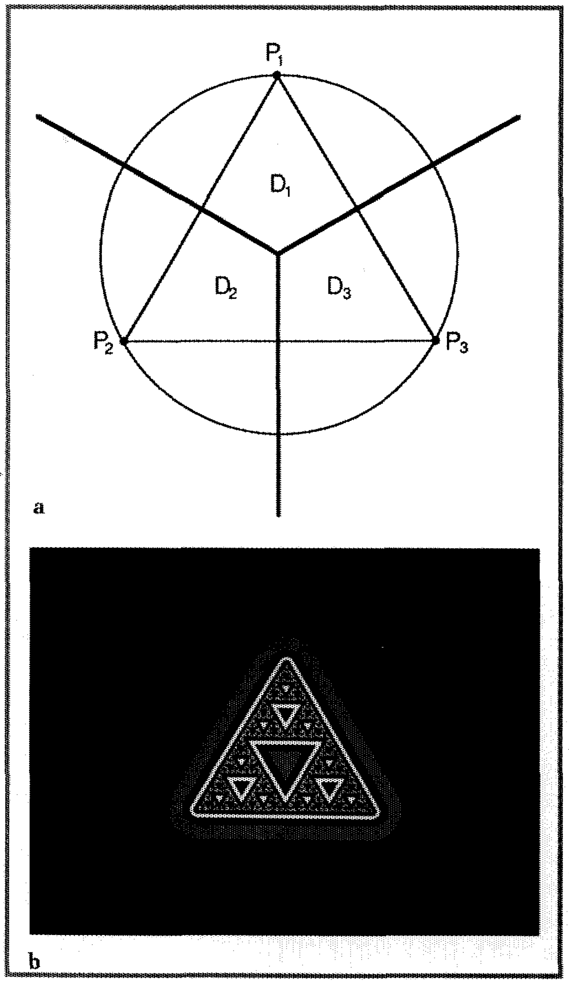

Figure 9. (a) Relative position of the scaling centers P_1 through P_3 and domains D_1 through D_3 for generating the Sierpinski gasket (b) using the repelling method. All scaling ratios are equal to 2.

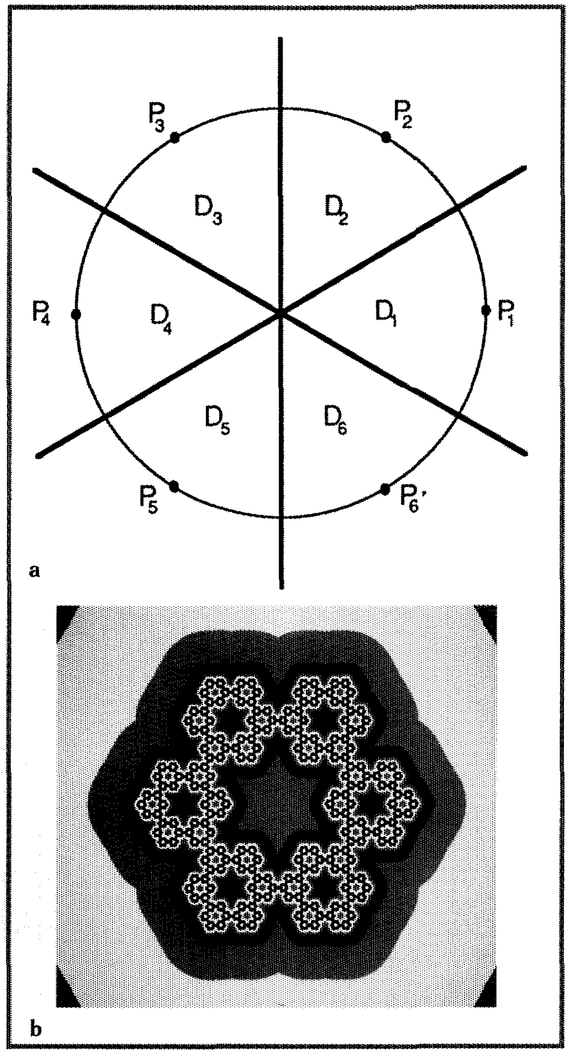

Figure 10. (a) Relative position of the scaling centers P_1 through P_6 and domains D_1 through D_6 for generating the "multisnowflake" curve (b) using the repelling method. All scaling ratios are equal to 3.

The justification of the above method is straightforward. If, after N iterations, a point x is taken out of the circle Ω that contains the curve C_∞, x does not belong to C_∞. On the other hand, if after N iterations x stays within Ω, it is assumed that $x \in C_\infty$. In fact, in this latter case x can be at some small distance from C_∞, but if the parameters are properly chosen, the error will be negligible compared with the screen resolution.

Example 6. Let us apply the repelling method to generate an image of the Sierpinski gasket. According to Example 4, the gasket is invariant with respect to three scalings ϕ_1 through ϕ_3, with centers at the vertices of an equilateral triangle and the scaling ratios equal to 1/2. Obviously, the reciprocal transformations ϕ^{-1}_1 through ϕ^{-3}_1 are also scalings, with the same centers and the scaling ratio equal to 2. To determine their domains D_i, let us refer to Figure 6. It shows that the gasket can be divided into three subgaskets (red, green, and yellow), which belong to different domains. Thus the domain boundaries can be defined as the bisectors of the line segments connecting pairs of the scaling centers (Figure 9a).

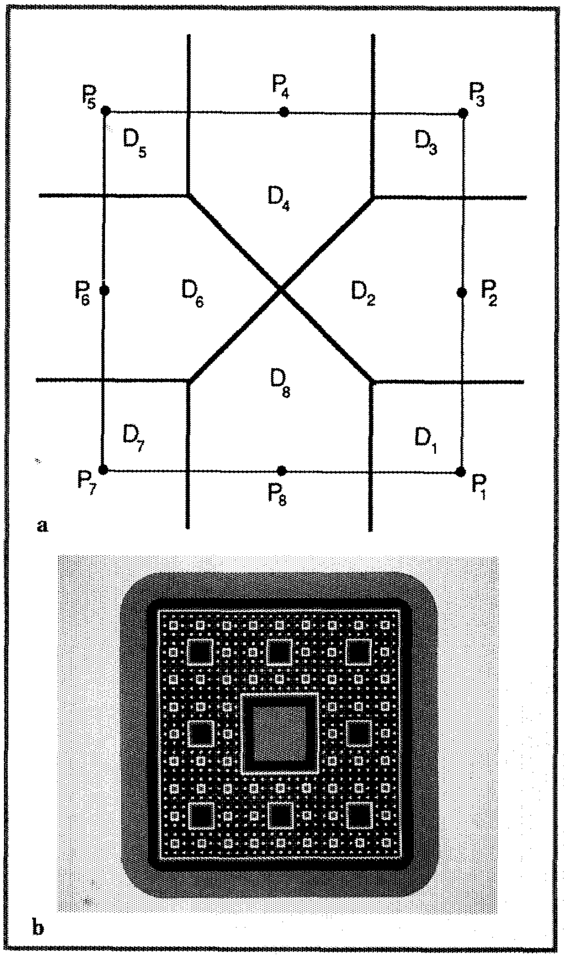

Figure 11. (a) Relative position of the scaling centers P_1 through P_8 and domains D_1 through D_8 for generating the Sierpinski carpet (b) using the repelling method. All scaling ratios are equal to 3.

The resulting image of the Sierpinski gasket obtained using the repelling method is shown in Figure 9b.

Example 7. The concept of sets of scalings can also be used to generate other fractals. Figures 10a and 11a define the scalings and their respective domains, which were used to generate images shown in Figures 10b and 11b. The fractal in Figure 10b consists of an infinite number of snowflake curves. The fractal in Figure 11b is the Sierpinski carpet.

Example 8. Partition of the plane into domains D_1 through D_5, corresponding to the "twig" from Example 5, is given in Figure 12a. The parameter values are $\gamma_1 = \gamma_2 = 1/3$ and $\alpha = \pi/3$. The exact positions of domain boundaries were arbitrarily chosen from the range of possibilities that satisfy the condition

$$\text{if} \quad z \in D_j \cap C_\infty \quad \text{then} \quad z\,\phi_j^{-1} \in C_\infty \ .$$

Figure 12b shows the repelling twig.

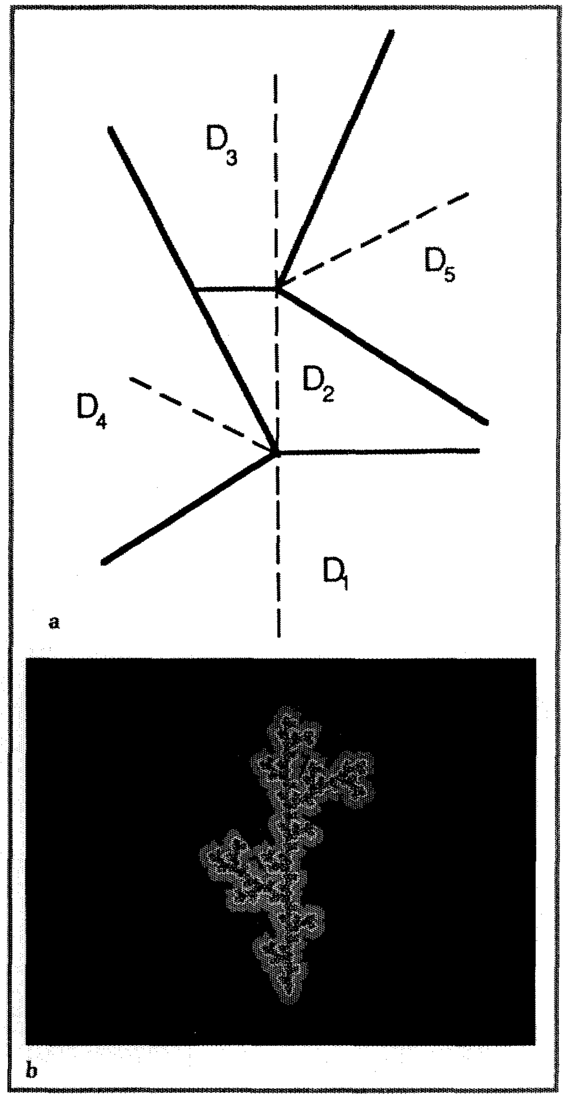

Figure 12. (a) Partition of the plane into domains D_1 through D_5 for generating a twig (b) using the repelling method. Transformations ϕ_1^{-1} through ϕ_5^{-1} are the reciprocals of the transformations from Example 5. For orientation, the dashed lines show the corresponding first-order curve C_1.

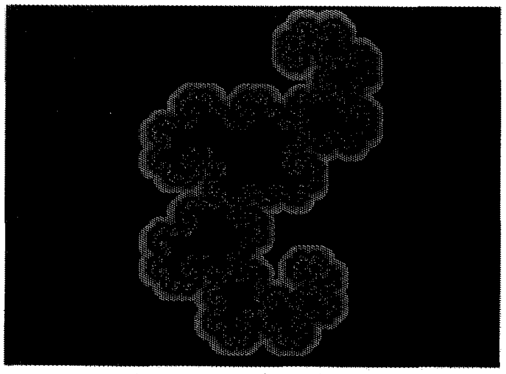

Figure 13. The dragon curve generated using the repelling method.

In principle, the repelling method for generating Koch curves is analogous to the widely used method for generating colorful images of Julia sets. However, the necessity of partitioning the plane into domains D_j can make it difficult to apply to some Koch curves. For example, refer to Example 3 and Figure 5. According to the coloring rules assumed, red points belong to the domain D_1 and green points belong to the domain D_2. Figure 5b indicates that the boundary between these two domains can itself be a fractal line. Domain definition in the cases illustrated in Figures 5c and 5d appears to be even more enigmatic.

If the plane cannot be easily subdivided into domains, the fractal can still be generated using a "brute force" variant of the repelling method. The idea is to keep track of *all* points resulting from the repetitive application of transformations $\phi_j^{-1} \in \Phi^{-1}$ to the sampling point x_0, as long as they stay within the circle Ω. Formally, if x_0 is the initial sampling point, the set X_n of points considered after the nth iteration step is given by the recursive formula

$$X_0 = \{x_0\}$$
$$X_{n+1} = X_i \, \Phi^{-1} \cap \Omega$$

The iteration stops if the set X_n becomes empty or the index n reaches a limit N. As previously, the final value of n determines the color of the point x_0.

Example 9. Figure 13 shows the repelling version of the dragon curve from Figure 5b.

Conclusions

This article presents two methods for generating Koch curves. They are analogous to the commonly used iterative methods for producing images of Julia sets. The attracting method is based on a characterization of Koch curves as the smallest nonempty sets closed with respect to a union of similarities on the plane. This characterization was first studied by Hutchinson. The repelling method is in principle dual to the attracting one, but involves a nontrivial problem of selecting the appropriate transformation to be applied at each step. Both methods are illustrated with a number of computer-generated images.

The Koch systems discussed in this article have the axiom limited to a single vector and use only one production. These restrictions can be removed by grouping all vectors into classes. The applicable production is then determined by the class a given vector belongs to. Each production also specifies the target classes for all resulting vectors. A corresponding approach can be applied to generate Koch curves by function iteration. In this case, a point in the plane is characterized by its position and an attribute or state. A typical transformation ϕ_i has a form "if point x is in state s_p, then take it to point y and make the state of the result equal to s_q." For an example of an image generated using this technique, see Figure 14. This branching shape belongs to a class termed "nonuniform fractals" by Mandelbrot[2] and cannot be generated by a Koch system with a single production. A formal characterization of Koch systems with multiple productions is left for further research.

There are also many other problems open for further research. Some of them follow:

- The repelling method for generating Koch curves presented in the previous section relies on a partition of the plane into domains D_j. However, domains D_j are defined only for the points that belong to the curve C_∞, and an arbitrary partition can be assumed outside of it. Are some of these partitions more "natural" than others? What is the impact of the partitions used on the images generated by the repelling method?

- The correspondence between Koch curves and Julia sets would be even more convincing if it could be illustrated by a *continuous* transformation of a Koch curve (such as the dragon curve) into a Julia set (such as the self-squared dragon).

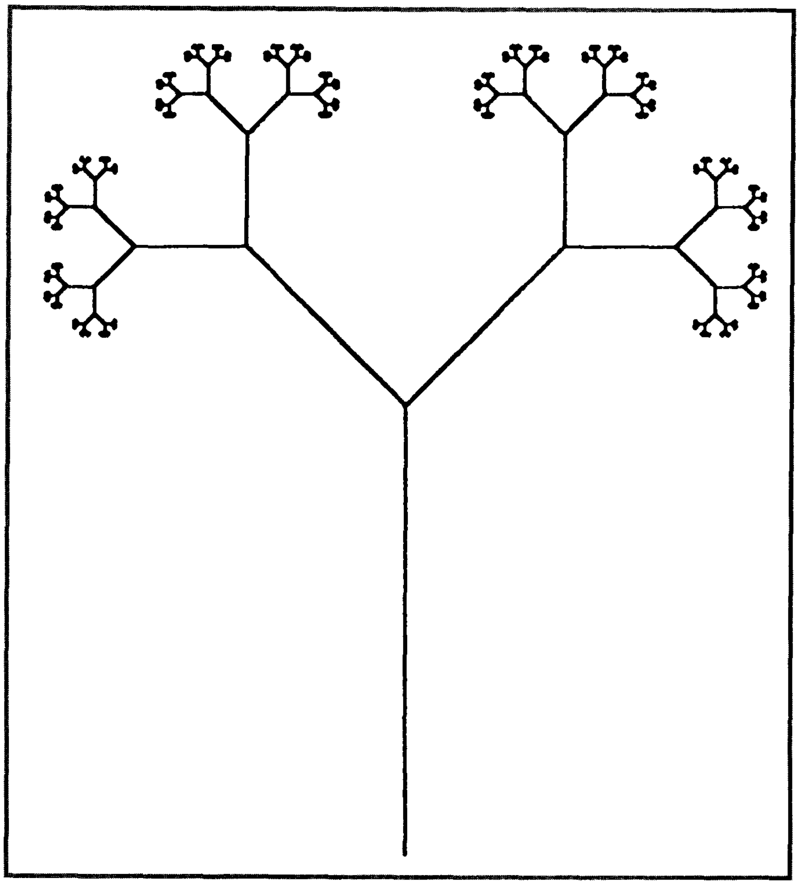

Figure 14. An example nonuniform fractal generated using the attracting method.

- This article shows that the usual description of the Koch curves in terms of an iterative geometric construction can be replaced by an algebraic characterization. A "dual" question applies to the Julia sets. Their known descriptions refer to the function iteration. Is it possible to define Julia sets by geometric constructions?

- Our results are related to the theory of iterated function systems originated by Barnsley and Demko[16] (further results also appear in the literature[15,17]). However, these systems operate in a probabilistic manner, while our approach is purely deterministic. Recently, a deterministic variant of iterated function systems has also been investigated by M. Barnsley, as he noted to us in a personal communication (August 1988). It would be interesting to study the relationship between his approach and our approach.

Finally, we would like to convey our impression of the general character of the reported research. We find it remarkable that it combines notions from areas of mathematics and computer science that traditionally have been perceived as quite unrelated. To name a few, we draw on results of the theory of formal languages, geometry, topology, and complex analysis, and we illustrate them using computer-generated images of fractals. Extrapolating from this experience, we believe that fractals may have great and still largely unexploited educational potential as a visually appealing method for illustrating various concepts of mathematics and computer science. Interestingly, the educational applications were also presented as the original motivation of Koch's work. ■

Acknowledgment

Benoit Mandelbrot kindly brought to our attention the paper by Hutchinson. Lynn Mercer and Jim Hanan volunteered valuable assistance in preparing illustrations included in this article. This research could not have been done without the support and facilities provided by the Department of Computer Science at the University of Regina. Support from the Natural Sciences and Engineering Research Council of Canada (grant number A0324) is also gratefully acknowledged.

References

1. J.E. Hutchinson, "Fractals and Self-Similarity," *Indiana Univ. Mathematics J.*, Vol. 30, No. 5, Sept./Oct. 1981, pp. 713-747.
2. B.B. Mandelbrot, *The Fractal Geometry of Nature*, W.H. Freeman, San Francisco, 1982.
3. H. von Koch, "An Elementary Geometric Method for Studying Some Questions in the Theory of Planar Curves" (in French), *Acta Mathematica* Vol. 30, 1906, pp. 145-174.
4. A.R. Smith, "Plants, Fractals and Formal Languages," *Computer Graphics* (Proc. SIGGRAPH), Vol. 18, No. 3, July 1984, pp. 1-10.
5. H. O. Peitgen and P.H. Richter, *The Beauty of Fractals*, Springer-Verlag, Berlin, 1986.
6. A. Norton et al., "Dynamics of $e^{i\theta} x(1-x)$" (film), *SIGGRAPH Video Review*, Issue 25, ACM, New York, 1986.
7. A. Lindenmayer, "Mathematical Models for Cellular Interaction in Development," Parts I and II, *J. Theoretical Biology*, Vol. 18, No. 3, Mar. 1968, pp. 280-315.
8. G. Rozenberg and A. Salomaa, *The Mathematical Theory of L Systems*, Academic Press, New York, 1980.
9. A.L. Szilard and R.E. Quinton, "An Interpretation for DOL Systems by Computer Graphics," *The Science Terrapin*, Vol. 4, 1979, pp. 8-13.
10. F.M. Dekking, "Recurrent Sets," *Advances in Mathematics*, Vol. 44, No. 1, Apr. 1982, pp. 78-104.
11. F.M. Dekking, "Recurrent Sets: A Fractal Formalism," Report 82-32, Dept. of Mathematics and Informatics, Delft Univ. of Technology, The Netherlands, 1982.
12. P. Prusinkiewicz, "Graphical Applications of L-Systems," *Proc. Graphics Interface, Vision Interface 86*, Canadian Information Processing Soc., Toronto, 1986, pp. 247-253.
13. B. Grunbaum and G. Shephard, *Tilings and Patterns*, W.H. Freeman, New York, 1987.
14. K. Kuratowski, *Introduction to Set Theory and Topology*, Pergamon Press, Oxford, UK, and Polish Scientific Publishers, Warsaw, 1972.
15. S. Demko, L. Hodges, and B. Naylor, "Construction of Fractal Objects with Iterated Function Systems," *Computer Graphics* (Proc. SIGGRAPH), Vol. 19, No. 3, July 1985, pp. 271-278.
16. M.F. Barnsley and S. Demko, "Iterative Function Systems and the Global Construction of Fractals," *Proc. Royal Soc. London A*, Vol. 399, No. 1817, June 1985, pp. 243-275.
17. J. Levy-Vehel and A. Gagalowicz, "Shape Approximation by a Fractal Model," *Eurographics 87*, North-Holland, Amsterdam, 1987, pp. 159-180, 572.

Przemyslaw Prusinkiewicz is a visiting associate professor of computer science and mathematics at Yale University, on leave from the University of Regina, Canada. His research interests are computer graphics, interactive techniques, and computer music. Prusinkiewicz received his MS and PhD in computer science from the Technical University of Warsaw, Poland. He is a member of ACM and IEEE Computer Society.

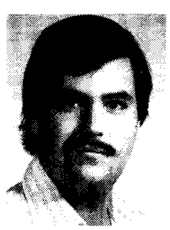

Glen Sandness is completing a master's program in computer science at the University of Regina. His research interests include computer graphics and interface design for handicapped persons. Sandness received his BSc from the University of Regina in 1985. He is a member of the ACM student chapter at Regina.

Prusinkiewicz can be reached at the Department of Mathematics, Yale University, Box 2155, Yale Station, New Haven, CT 06520.

Sandness can be contacted at the Department of Computer Science, University of Regina, Regina, Saskatchewan, Canada S4S 0A2.