Self-Inverse Fractals Osculated by Sigma-Discs and the Limit Sets of Inversion Groups

Benoit B. Mandelbrot

Behind the dry facade of the hard mathematical analysis in the style of 1900 lurks a geometry of extraordinary plastic beauty and suggestive power. There was a hint of it in the great old treatise by Fricke and Klein, "Lectures on Automorphic Functions" [1], and it was fully revealed in [2]. The figures in [1] must rank among the most widely known of all mathematical illustrations, since they include the tessellations of the hyperbolic plane that the non-mathematical millions now credit to Maurits C. Escher. The present paper explores further a small corner of this universe, the geometry of the limit set of a special group based upon inversions. This limit set, a fractal curve in my terminology [2], was among those illustrated in [1]; it has been reproduced on faith by several famed books, thus helping to form the intuition of many generations of mathematicians.

Unfortunately, careful graphics performed on a computer revealed a rich structure that proves the Fricke and Klein illustration to be inaccurate and misleading—the conventional algorithm used to draw the limit set (going back to Henri Poincaré [3]) is indirect and very inefficient.

Fortunately, the careful graphics also helped me to identify a new algorithm that generates the limit set $\mathscr{L}$ for many groups $\mathscr{G}$ based upon inversions. This algorithm was first sketched in [4] and is described in [2, Chapter 18]. Since it is *not* of universal validity, it separates the groups based upon inversions into "directly osculable" and "not directly osculable" groups. (The limit set may become osculable after a change of basis.) The details of these new distinctions are, however, beyond the scope of this paper. Its sole purpose is to demonstrate this algorithm and its efficiency by describing in detail its application to the most striking of the Fricke and Klein examples.

Given a group of geometric transformations, it is interesting to identify the sets that are invariant under the action of the group, and especially the smallest among these sets. Such a question was first raised by Leibniz, who suspected that the only shapes invariant under all similarity transformations of the plane are straight lines and the whole plane (see [2], p. 419). Under the assumptions of smoothness, he was correct: these are the only connected and smooth invariants, called *self-similar sets*. Waiving the standard conditions of smoothness and restricting the similarities, however, one finds many other self-similar sets, all of which are not shapes from standard geometry, but fractal sets; for example, Cantor sets (totally disconnected sets; dusts in my terminology [2]), the boundaries of Koch snowflakes (nonsmooth curves), and Brownian motion (the best known self-similar random curve).

The present paper is devoted, in the same spirit, to a nonlinear group of transformations. Given $M \geq 3$ circles $C_m (1 \leq m \leq M)$ in the plane, to be called generat-

Benoit B. Mandelbrot

ing circles, we consider the group $\mathcal{G}$ generated by inversions with respect to these circles. Choosing the center of a circle C of radius R as origin, the inversion with respect to C is described in polar coordinates as the map $(r, \theta) \rightarrow (R^2/r, \theta)$. In what follows, the key fact about inversions is that they map circles to circles—provided we call a straight line a circle through the point at infinity.

The sets invariant under the action of this group can be called *self-inverse sets* under $\mathcal{G}$. The closed plane, including its point at infinity, is self-inverse under every group. And for many groups it is the only standard solution. One notable exception—called Fuchsian—is when all the C_m are orthogonal to a common circle Γ. In this case (which includes most cases with $M = 3$), the circle Γ is self-inverse under $\mathcal{G}$, since any circle orthogonal to C is invariant under inversion in C. It is also instructive to side track to a case excluded by the above definition. When $M = 2$ and the finite discs bounded by C_1 and C_2 do not intersect, there exist two points, γ_{12} in the finite disc bounded by C_1 and γ_{21} in the finite disc bounded by C_2, that are mutually inverse with respect to both C_1 and C_2. (If Γ is a circle orthogonal to C_1 and C_2 then the points γ_{12} and γ_{21} can be chosen on Γ.) Therefore, the set $\{\gamma_{12}, \gamma_{21}\}$, made of two points, is self-inverse. Some very special Fuchsian groups also have a self-inverse set reduced to 2 points.

On the other hand, the first student of this topic, Poincaré, observed a hundred years ago that waiving the standard assumptions of smoothness can yield strange subsets of the plane as closed self-inverse sets. Shortly before (!) Cantor introduced his set in 1884, and well before Koch introduced his nondifferentiable snowflake curve in 1904, Poincaré noted that in typical configurations of the generators C_m, the self-inverse set $\mathcal{L}$ can be either a totally disconnected set ("dust"), or a curve that is nondifferentiable (either has no tangent, or has a tangent but no curvature). (He commented that one must assume that "one can call *that* a curve". . . . Later studies of the concept of curve have confirmed that one can and that one should.)

The Limit Set

How do we find a self-inverse set? The answer is that one can start with any set S, and enlarge it just enough to *make* it self-inverse. For any set S the *clan* of $\mathcal{G} S$ is defined to be the union of all transforms $g(S)$ for $g \epsilon \mathcal{G}$—this is usually called the orbit of S, but I prefer the term "clan". Of course the closure of $\mathcal{G} S$ is a closed, self-inverse set. So is the subset $(\mathcal{G} S)'$ consisting of all limit points of $\mathcal{G} S$. Recall that a point P is a limit point of a set if every deleted neighborhood of P—i.e., the neighborhood minus P itself—intersects the set; the set of limit points is called the derived set.

In particular, one can start with any point P_0, then form the clan $\mathcal{G} P_0$ and take the derived set $(\mathcal{G} P_0)'$. The amazing fact is that, under wide conditions, the derived set is a) independent of the point P_0, so that it can be denoted by $\mathcal{L}$; and b) of zero area (= planar Lebesgue measure). The derived set $\mathcal{L}$ has several very important characteristic properties. Not only is it self-inverse, but it is the *minimal* self-inverse set. It is (by construction) the *limit set* of $\mathcal{G}$. Furthermore, it is the set on which the group $\mathcal{G}$ is *continuous*, and outside of which $\mathcal{G}$ is discontinuous. We give one final useful characterization: it is clear that $\mathcal{G}$ includes an infinity of inversions in addition to its base, and that $\mathcal{L}$ is the derived set of the centers of these inversions. To my knowledge, this last statement is not put in this form in the literature, but it can be seen to be equivalent to Poincaré's construction of $\mathcal{L}$.

What is the shape of $\mathcal{L}$? Our knowledge of the Cantor and Koch sets benefits from the availability of a multiplicity of direct and transparent constructions. To the contrary, the shape of the minimal self-inverse $\mathcal{L}$ has remained elusive. To close this gap, I have devised a new construction which involves the self-inverse open sets obtained as clans of open discs; I call them "sigma-discs", or "σ-discs", where the letter σ is self-explanatory; it indicates that a σ-disc is the denumerable union of discs. The complement of a self-inverse σ-disc is also a self-inverse set, and the minimal self-inverse $\mathcal{L}$ can ordinarily be represented as the complement of a finite union of σ-discs. Each of these σ-discs can be said to *osculate* $\mathcal{L}$. This "osculating σ-disc" construction can be made recursive: as the recursion advances, it tends to outline $\mathcal{L}$ very rapidly—much more rapidly than the classical construction of Poincaré.

Fuchsian groups with $M = 3$ for which $\mathcal{L}$ is the circle Γ

It was noted that when $\mathcal{G}$ is Fuchsian, the circle Γ is self-inverse. A case when Γ is the minimal self-inverse set occurs when $M = 3$ and each of three circles C_1, C_2 and C_3 is tangent to the other two (Fig. 1). By inverting

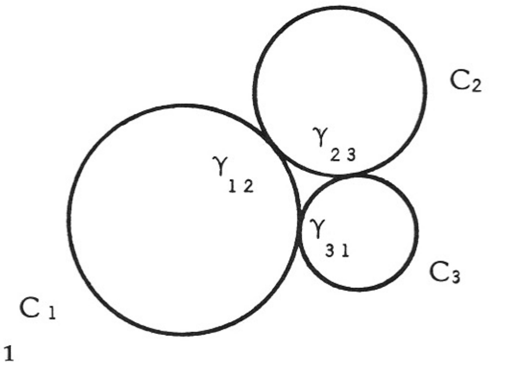

Figure 1

the plane about the point of tangency γ_{12} of C_1 and C_2, one achieves the situation of Figure 2, where C_1 and C_2 are parallel straight lines and C_3 is tangent to both. The

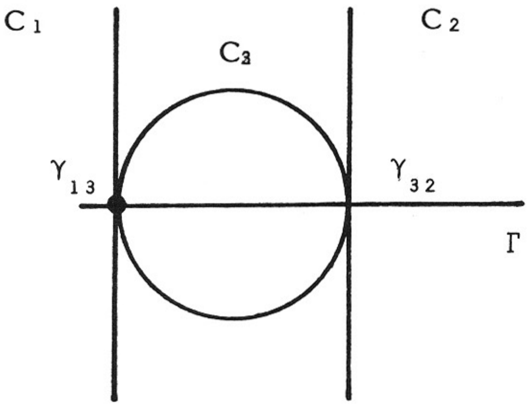

Figure 2

group is also transformed: it is now generated by reflections in C_1 and C_2 and inversion in C_3. It is easily seen that $\mathcal{L}$ is Γ.[1]

Some Fuchsian groups with $M = 3$ for which $\mathcal{L}$ is a fractal dust. Osculating σ-intervals

Figure 3 gives an illustrative example where $M = 3$ and the group is Fuchsian but $\mathcal{L}$ is a very small subset

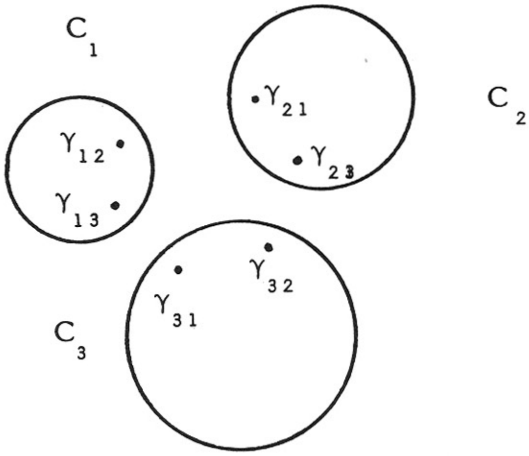

Figure 3

of Γ: we let Γ be the orthogonal circle and assume that the finite discs bounded by C_1, C_2 and C_3 are nonintersecting. It was noted that there are points γ_{12} and γ_{21} on Γ that are mutually inverse with respect to both C_1 and C_2. Now invert the plane about γ_{12}, obtaining Figure 4. Since γ_{12} goes to infinity, the circle Γ becomes a straight line; C_1 and C_2 become two circles with the common center γ_{21}. It is easy to see that the half of Γ that does not include the points γ_{13} and γ_{31} and the points γ_{23} and γ_{32} fails to belong to $\mathcal{L}$. We use this half line to define the open interval $]\gamma_{12},\gamma_{21}[$.

Thus, the self-inverse set $\mathcal{G}$ $]\gamma_{12},\gamma_{21}[$, which is a σ-interval (a union of nonoverlapping open intervals) lies entirely in the complement of $\mathcal{L}$. The same is true of the σ-intervals $\mathcal{G}$ $]\gamma_{13},\gamma_{31}[$ and $\mathcal{G}$ $]\gamma_{23},\gamma_{32}[$. Also, but not quite so obviously, the complement of these three σ-intervals is $\mathcal{L}$. It is a fractal dust of zero length (linear Lebesgue measure), and provides a self-inverse version of the usual self-similar Cantor set.

One can delineate the shape of $\mathcal{L}$ with rapidly increasing accuracy by injecting the intervals of the complement of $\mathcal{L}$ in order of decreasing length, or—more simply—in order of increasing length of the shortest words in $\mathcal{G}$ that generate these intervals from one of the $]\gamma_{ij},\gamma_{ji}[$. Although I find it hard to believe that the preceding algorithm is new, I do not recall having seen it described anywhere.

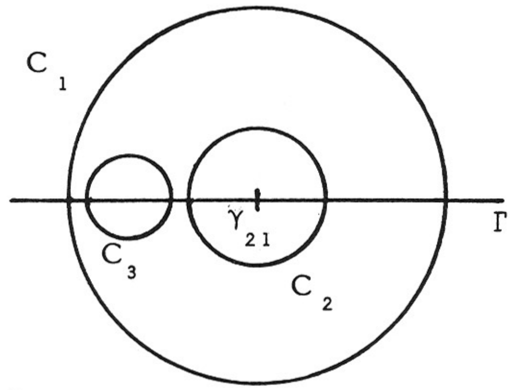

Figure 4

Fricke and Klein: Misled and Misleading

The limit set is relatively easy to find for a group generated by inversions in three circles, but what about larger configurations? That is the point of the new algorithm for $\mathcal{L}$. It too is so completely elementary that

[1] The group of Figure 2 is a special "modular" subgroup. Take the center of C_3 as origin and its radius as 1. Then each real x_1 can be written uniquely as $i_1 + \rho_1$, where i_1 is a signed even integer and ρ_1E lies in the open interval $]-1,1[$. Similarly, $1/\rho_1 = i_2 + \rho_2$. When x_1 is rational, every ρ_k is also rational; as k increases, the denominators of ρ_k decrease and eventually, for $k = K$, reach $\rho_K = 0$ or 1. Call K the depth of x_1. For irrational numbers, $K = \infty$; for integers $K = 0$. It is easy to prove that every rational x_1 with $\rho_K = 0$ is the transform of $x_K = 0$ by a word of the group $\mathcal{G}$, while every rational x_1 with $\rho_K = 1$ is the transform of $x_K = 1$ when $i_K/4$ is an integer, and of $x_K = -1$ otherwise. Furthermore, every x_l with $\rho_K = 0$ is obviously the limit point of x_1's with $\rho_K + 1 = 1$. Conclusion: whenever the abscissa of P_0 is rational, the set $\mathcal{G}P_0$ is dense in Γ.

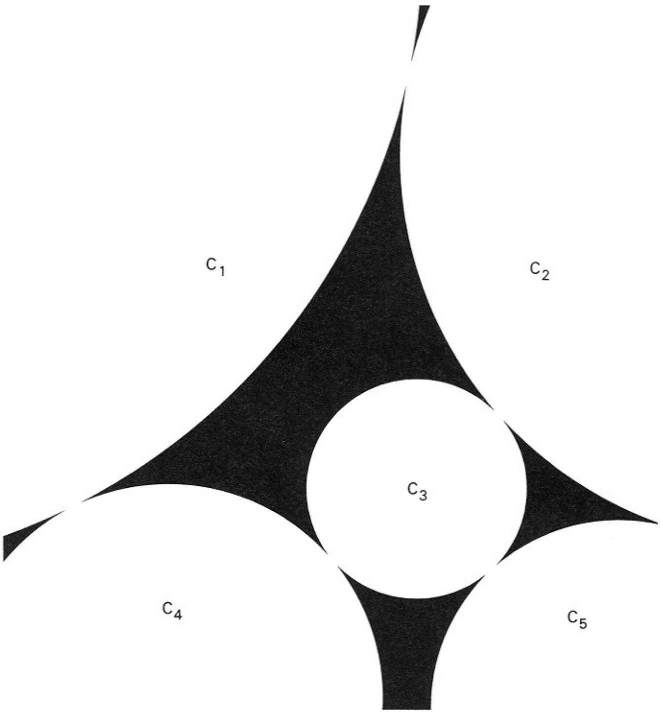

Figure 5

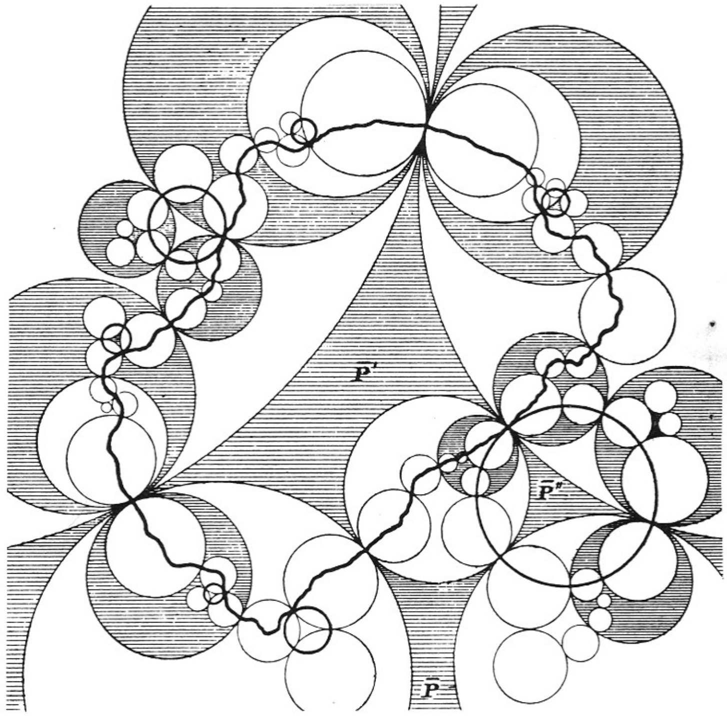

Figure 6

it might have been (but was not!) recognized in the 1880's, when Henri Poincaré and Felix Klein first tackled this topic. Hence it is best to present it against the classic background of those illustrations in Fricke and Klein [1] that purport to represent the limit sets $\mathcal{L}$ of several special inversion groups.

The main discrepancy does *not* lie in the fact that the true $\mathcal{L}$'s involve detail that no one would attempt to draw by hand. Even if detail is erased (which is best done by stopping the new algorithm after a small number of stages) the "old $\mathcal{L}$" seems cruder: structureless and lawless.

Perhaps it is best to tackle a single example in detail in order to illustrate my construction. Our point of departure is Figure 5, which reproduces the 5 circles used in Figure 156 of [1]. Next, Figure 6 reproduces the corresponding "old $\mathcal{L}$" as claimed by Fricke and Klein: it consists of a wiggly curve together with a collection of circles, shown in heavy outline. The "true $\mathcal{L}$" is shown in diverse guises in many of the later figures in this article. On Figure 7, the true $\mathcal{L}$ is shown as the boundary of a black background domain.

It now seems obvious that Figure 6 was drawn by a hapless draftsman (legend has it that he was an engineering student in Fricke's class), who had been instructed how to determine a few points of $\mathcal{L}$ exactly, and was then left to draw "some very wiggly and complicated curve" passing through these points. As Fricke did not know what to expect, the draftsman received no explicit directions.

Observe that the black domain on Figure 5, viewed as open (not including the generating circles), splits into three disjoint maximal open domains, to be denoted as $\mathcal{D}_3$ (a triangle), $\mathcal{D}_4$ (a quadrilateral), and $\mathcal{D}_5$ (a pentagon containing the point at infinity). This suggests that we might begin to investigate $\mathcal{L}$ by investigating the limit sets of two subgroups of $\mathcal{G}$: first, the Fuchsian subgroup $\mathcal{G}_3$, generated by the inversions in the 3 circles bounding $\mathcal{D}_3$, and second, the subgroup $\mathcal{G}_4$, generated by the inversions in the 4 circles bounding $\mathcal{D}_4$. The limit set of a subgroup is, after all, contained in the limit set of the group.

The Fuchsian Subgroup $\mathcal{G}_3$

One point at which the old $\mathcal{L}$ and the true $\mathcal{L}$ agree is that both include a large circle, to be denoted as $\mathcal{L}_3$, and smaller circles. The circle $\mathcal{L}_3$ is orthogonal to the 3 generating circles C_2, C_3, and C_5 that bound the domain $\mathcal{D}_3$, and from our previous discussion it is the limit set of the subgroup $\mathcal{G}_3$.

Now suppose we take $\mathcal{L} = (\mathcal{G}P_0)'$, where P_0 is an arbitrary point of $\mathcal{L}_3$. A fortiori, we have the representation $\mathcal{L} = (\mathcal{G}\mathcal{L}_3)'$, which is extravagantly redundant, yet very useful in describing the limit set. An approximate idea of $\mathcal{L}$ based upon this representation is given by Figure 8. It limits itself to the elements of $\mathcal{G}$ that are obtained as products of fewer than a certain large number K of inversions.

Large circles are represented by thin curves, and circles whose radius is smaller than the large circles'

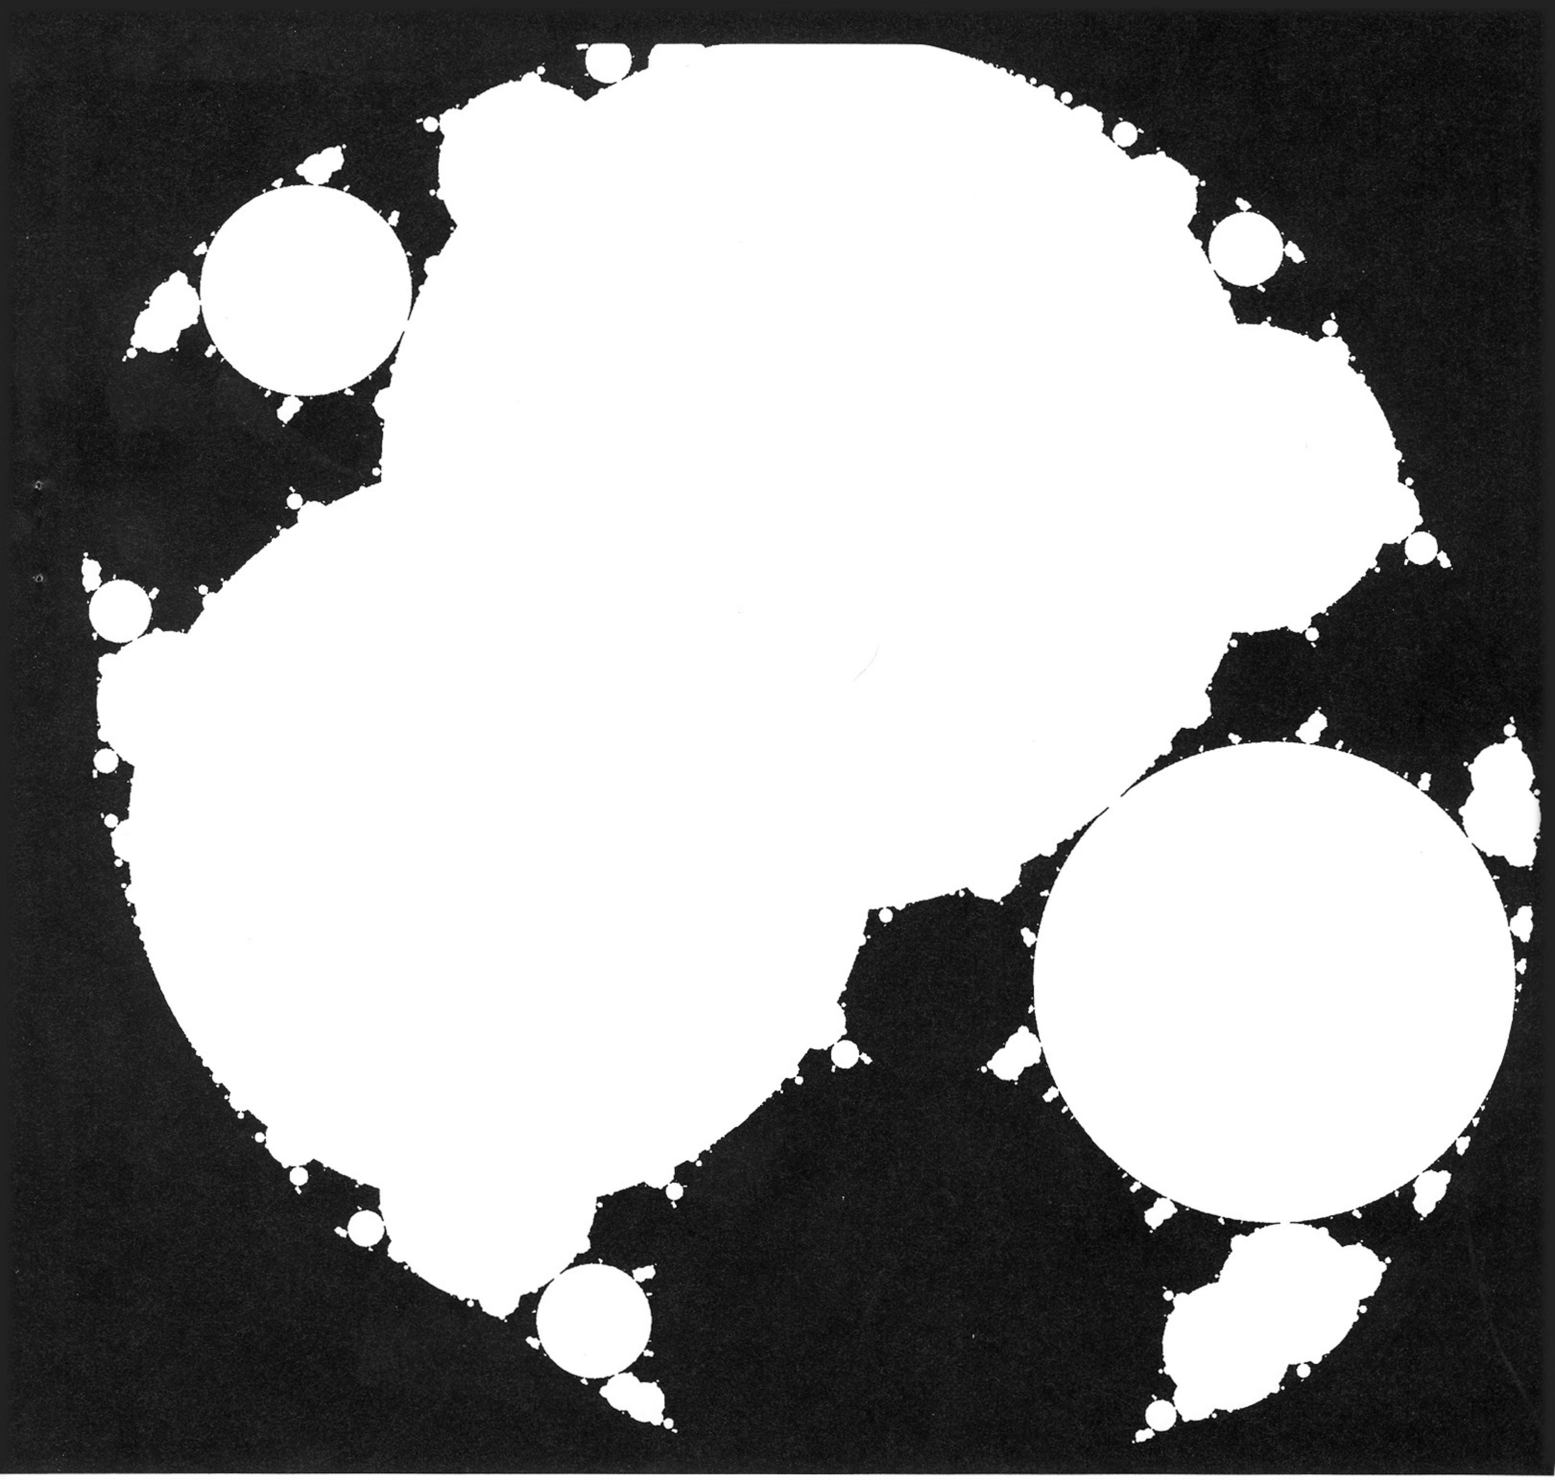

Figure 7

thickness, are represented by points. Were the algorithm pushed further, all these points would bleed together to form an infinitely ramified curve. Had we abstained altogether from plotting the large circles, the result would not change perceptibly, because every point on every circle can be obtained in alternative fashion as the limit of points belonging to smaller circles, that is, of points which stand for small circles.

The formula $\mathscr{L} = (\mathscr{G}\mathscr{L}_3)'$ suffices to represent $\mathscr{L}$, but it is cumbersome. The sequence $\mathscr{G}_K\mathscr{L}_3$, where $\mathscr{G}_K$ is the collection of products of at most K inversions, converges to $\mathscr{L}$ more rapidly than the sequence $\mathscr{G}_K P_0$ relative to any single point P_0. Nevertheless, it converges very slowly: without the exceptional computer facilities marshalled for Figure 7, this algorithm would only yield a very loose idea of $\mathscr{L}$.

The Chain-based Subgroup $\mathscr{G}_4$

The most striking discrepancy between the old $\mathscr{L}$ and the true $\mathscr{L}$ concerns the limit set $\mathscr{L}_4$ corresponding to the subgroup $\mathscr{G}_4$ based on the 4 circles, C_1, C_2, C_3, and C_4, that touch the domain $\mathscr{D}_4$. In the terminology of

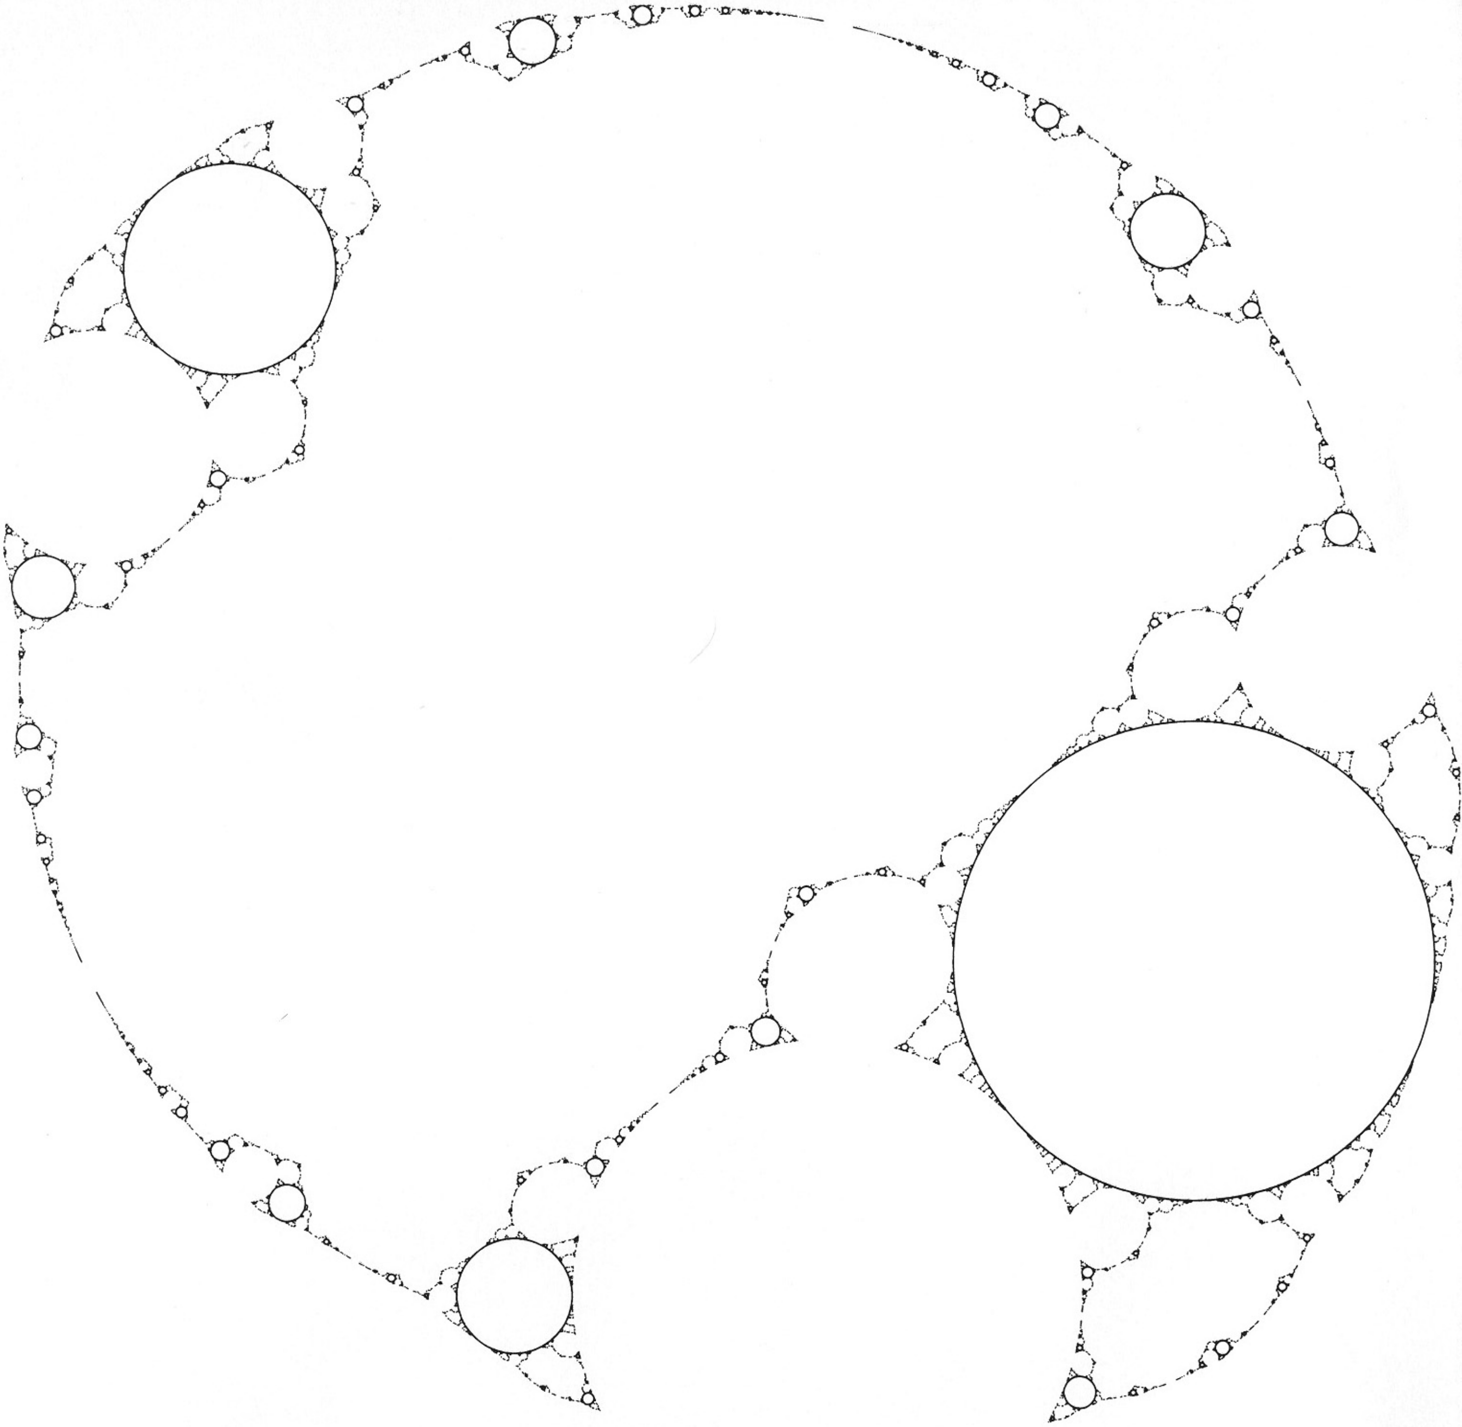

Figure 8

[2], these circles form a connected "Poincaré chain", in which each link is tangent to exactly two neighbors. (As a result, the limit set $\mathscr{L}_4$ is a Jordan curve.) The old $\mathscr{L}_4$, separated from the rest of Figure 6, is shown on Figure 9, and the true $\mathscr{L}_4$, as constructed by my new algorithm, is shown on the back cover.

How do we determine $\mathscr{L}_4$? The idea is (once again) to determine what is *not* in $\mathscr{L}_4$.

We begin with the observation that for any three circles C_i, C_j, C_k there is a common orthogonal circle Γ_{ijk} that passes through the points of tangency of C_i, C_j, *and* C_k. In this case the Γ_{ijk} are distinct; that is, $\mathscr{G}_4$ is not Fuchsian.

Each circle Γ_{ijk} divides the plane into two discs—one bounded and one unbounded. One of these two open discs contains no points of tangency of the four circles; we denote it by Δ_{ijk}. To see this we can invert the plane about a circle centered on the point of tangency γ_{ij}. The configuration of circles become that in Figure 10, and Γ_{ijk} becomes a horizontal line: the transform of Δ_{ijk} is the half-plane above or below the transform of Γ_{ijk}.

Now it is not hard to see that an inversion in one of the four circles cannot carry a point of tangency inside Δ_{ijk}. It follows that no transform of Δ_{ijk} contains a transform of a tangent point. If we use the tangent

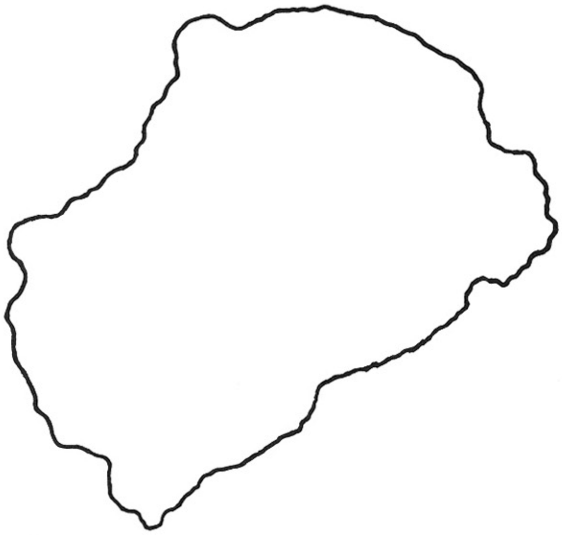

Figure 9

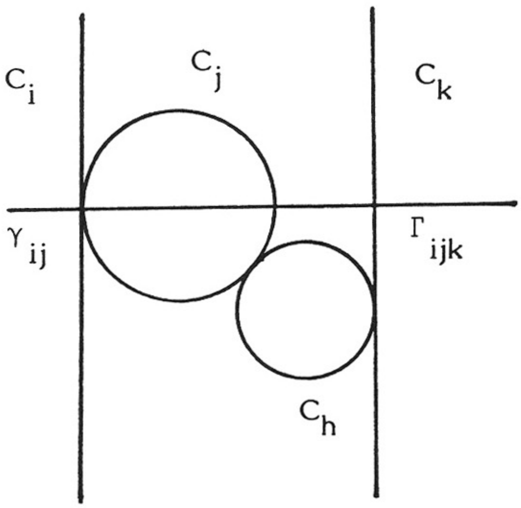

Figure 10

point γ_{ij} to compute $\mathscr{L}$, then the union of all transforms of the four discs Δ_{ijk} must be in the complement of $\mathscr{L}$.

On the other hand, the boundary of Δ_{ijk} (the circle Γ_{ijk}) contains the limit set of the Fuchsian sub-group generated by inversions in C_i, C_j, and C_k. Therefore, any open disc containing Δ_{ijk} must intersect $\mathscr{L}$.

Now consider the 4 open discs Δ_{ijk} and their clans separately. One of the initiators, namely the disc Δ_{124}, is unbounded, and it intersects the disc Δ_{234}. (In the present configuration, Δ_{234} is also unbounded—in fact, Γ_{234} is nearly a straight line—but in other configurations of 4-link Poincaré chains Δ_{234} may be bounded.) The discs Δ_{124} and Δ_{234} are easily identified on Figure 11. Together, they provide a first approximation of the outside of $\mathscr{L}_4$. It is analogous to the approximation of the Koch snowflake curve $\mathscr{K}$ in Figure 43 of [2], using the regular convex hexagon that circumscribes $\mathscr{K}$.

The other initiator discs Δ_{ijk}, namely, the discs Δ_{123} and Δ_{143}, are bounded and intersect each other. They are easily identified on Figure 12. Together, they provide a first approximation of the inside of $\mathscr{L}_4$. It is analogous to the approximation of $\mathscr{K}$ by two triangles forming the inscribed regular star hexagon (Plate 43 [2]).

A second approximation of the outside of $\mathscr{L}_4$, also clearly seen in Figure 11, is achieved by adding to Δ_{124} and Δ_{234} their inverses in C_3 and C_1, respectively. The result is analogous to the second circumscribed approximation of $\mathscr{K}$ on Plate 43 of [2]. The corresponding second approximation of the inside of the $\mathscr{L}_4$ is achieved by adding to Δ_{123} and Δ_{143} their inverses in C_4 and C_2, respectively. The result is analogous to the second inscribed approximation of $\mathscr{K}$ (Plate 43[2]).

The complement of the σ-disc (denumerable union of discs) made up of the four clans $\mathscr{G}\ \Delta_{ijk}$ squeezes down to the curve $\mathscr{L}_4$. The union of the four "initiator discs" alone provides a useful approximation of the complement of $\mathscr{L}_4$. The approximations using the product of K or fewer inversions converge rapidly to $\mathscr{L}_4$.

Fractal Osculation. Osculating Discs

The fact that $\mathscr{L}_4$ is not intersected by any of the four open discs Δ_{ijk}, with indices associated with $\mathscr{L}_4$, but is intersected in more than one point by the circle bounding every Δ_{ijk}, suggests that $\mathscr{L}_4$ and Δ_{ijk} be called *osculating*.

In its standard context in differential geometry, the notion of osculation is linked to the concept of curvature. To the first order, a standard curve near a regular point P is approximated by the tangent line. To the second order, it is approximated by the circle, called "osculating", which has the same tangent and the same curvature.

The circles tangent to the curve at a point P can be indexed by u, the inverse of the distance from P to the circle's center. The index of the osculating circle will be written as u_0. If $u < u_0$, a small portion of curve centered at P lies entirely on one side of the tangent circle, except for P itself, while if $u > u_0$ it lies entirely on the other side, except for P. We say u_0 is a *critical value* or a *cut*.

For fractals, the definition of osculation by curvature is meaningless. However, there is an infinity of points where the limit set $\mathscr{L}$ of any Poincaré chain squeezes between two discs tangent to each other. For example, the point of tangency of the generating circles C_i and C_j belongs to $\mathscr{L}$, and $\mathscr{L}$ squeezes between two discs Δ_{ijk} and Δ_{ijh}. It is tempting to call both of these discs *osculating*.

Figure 12

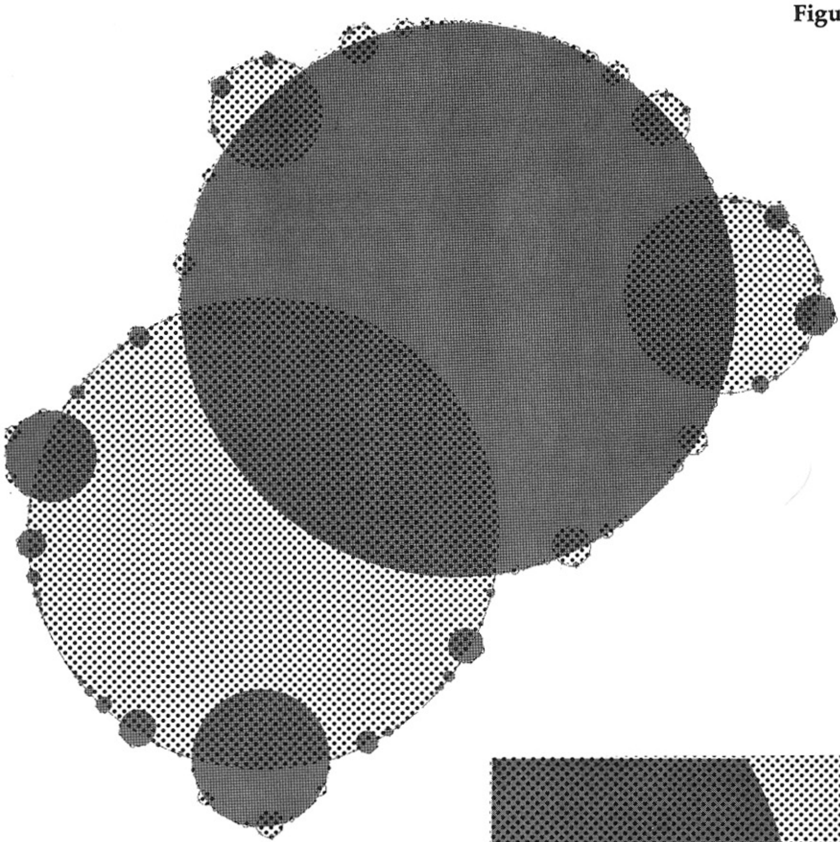

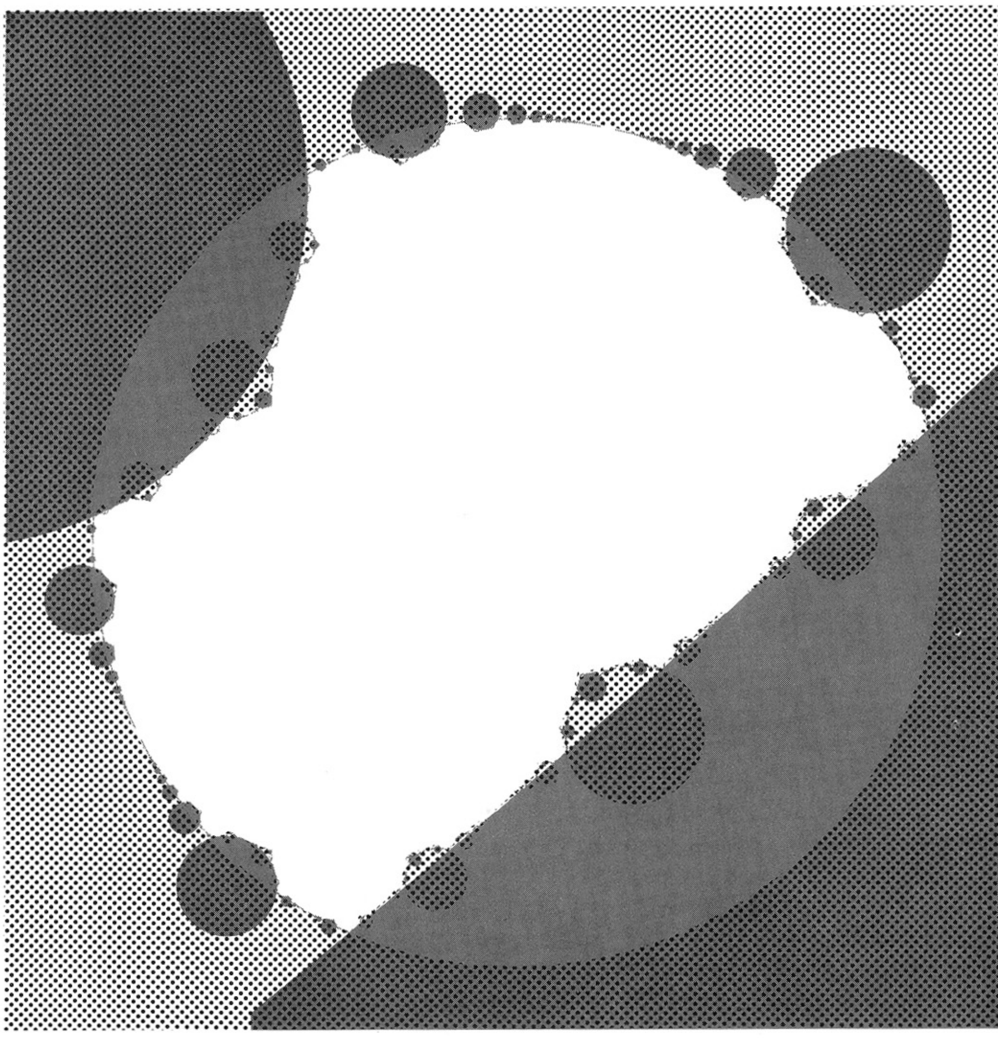

Figure 11

To pinpoint this notion, we take a point P where $\mathscr{L}$ has a tangent, and start with the definition of ordinary osculation based on critically (= cut). The novelty is that, as u varies, the single critical u_0 is replaced by two distinct values, u' and $u'' > u'$, defined as follows: For all $u < u'$, $\mathscr{L}$ lies entirely on one side of our circle except for P, while for all $u > u''$, $\mathscr{L}$ lies entirely to the other side, except for P; and for $u' < u < u''$, parts of $\mathscr{L}$ are found on both sides of the circle. I suggest that the circles of parameters u' and u'' be *both* called *fractally osculating* to $\mathscr{L}$. The open discs bounded by the osculating circles and not intersecting $\mathscr{L}$, will be called osculating discs. It may happen that one or two osculating circles degenerate to a point.

As is well known, standard osculation is a local concept, since its definition is independent of the curve's shape away from P. By contrast, I have defined fractal osculation globally, because this is all that was needed here; a local version is defined without difficulty.

The Group $\mathscr{G}$. Classification of the circles Γ_{ijk} as either osculating or intersecting the set $\mathscr{L}$

We found that for the subgroup $\mathscr{G}_4$ every one of the $4!/3!1! = 4$ circles Γ_{ijk}, defined an orthogonal to 3 of the generating circles C_m, is osculating to $\mathscr{L}_4$. For other configurations of generating circles C_m, including the configuration in Figure 5 and Poincaré chains with $M > 4$, the situation is more complex.

For example, consider the complement of the limit set $\mathscr{L}$. We know already that one needs the clan of Δ_{235} under $\mathscr{G}_3$ to generate the inside of $\mathscr{L}_3$, and that one needs the clans of Δ_{123} and Δ_{234} under $\mathscr{G}_4$ to generate the inside of $\mathscr{L}_4$, and we shall see that one needs the clans of Δ_{235}, Δ_{123} and Δ_{234} under $\mathscr{G}$ to generate the inside of $\mathscr{L}$ (Figure on this issue's back cover). Similarly we shall see that one needs the clans of Δ_{345}, Δ_{145}, and Δ_{125} under $\mathscr{G}$ in order to generate the outside of $\mathscr{L}$ (Figure on this issue's front cover).

However, the circles based on the triplets 234 and 135 behave in a *totally different* way: In either case, some of the remaining circles C_m are positioned, respectively, in the interior disc $\mathscr{J}\Gamma_{ijk}$ and the exterior disc $\mathscr{E}\Gamma_{ijk}$. Two open discs, namely $\mathscr{G}\Gamma_{ijk}$ and its inverse in the circle C_m located within $\mathscr{J}\Gamma_{ijk}$, cover the whole plane; the same is true of $\mathscr{E}\Gamma_{ijk}$. *A fortiori*, the clans $\mathscr{G}\mathscr{J}\Gamma_{ijk}$ and $\mathscr{G}\mathscr{E}\Gamma_{ijk}$ are both identical to the whole plane. Hence, when $ijk = 234$ or 135, neither the interior nor the exterior of Γ_{ijk} can serve as the initiator disc of an osculating σ-disc. It continues to be true, however, that the complement of $\mathscr{L}$ is a σ-disc, and that it is obtained as the union of clans of the form $\mathscr{G}\Delta_{ijk}$. Moreover, whenever Δ_{ijk} *can be* defined, *it is needed* for the construction, and whenever Δ_{ijk} *cannot be* defined, *one can do without it*. (Nevertheless, when C_i, C_j and Ck form a Fuchsian subgroup having a full circle Γ_{ijk} as its limit set, the construction of $\mathscr{L}$ is made much faster by including $\mathscr{G}\Gamma_{ijk}$.)

Now to the statement of a general rule. In order for the triplet i, j, k to be needed in the case of the group $\mathscr{G}$, the key facts are as follows: A) There exists a unique circle orthogonal to C_i, C_j, and C_k; B) Either the inside of Γ_{ijk} or the outside of Γ_{ijk} fails to contain any of the points of tangency between any two circles C_m and C_n.

More general non-Fuchsian groups

This last rule applies to more general groups $\mathscr{G}$ based upon inversions. The first step is to replace the points of tangency when C_m and C_n are not tangent. As noted early in this paper, when C_m and C_n fail to overlap, there are two points, γ_{mn} in C_m and γ_{nm} in C_n, which are mutually inverse with respect to both C_m and C_n. When C_m and C_n overlap, let γ^*_{mn} and γ^*_{nm} denote their intersection points. Now requirement B) at the end of the preceding section generalizes as follows: the circle Γ_{ijk} is osculating if, for all m and n, either its inside or its outside fails to include γ_{mn}, and includes either the point γ^*_{mn} or the point γ^*_{nm}, but not both.

Fuchsian groups and osculating intervals

When there is a circle Γ orthogonal to all the C_m, my algorithm always yields Γ itself, but $\mathscr{L}$ may be either Γ, or a fractal dust subset of Γ. In addition, the group may be everywhere continuous, a case which I call chaotic and which is not investigated in this paper; if so, $\mathscr{L}$ is the whole plane. In order to construct $\mathscr{L}$ when it is not the plane, we use osculating intervals.

When both intervals bounded by γ_{ij} and γ_{ji} contain no other γ_{mn}, $\mathscr{L}$ reduces to γ_{ij} and γ_{ji}. This case plays for the Fuchsian groups the same degenerate role as the Fuchsian groups play for the other groups based upon inversions.

Acknowledgment. The illustrations in this paper were drawn using computer programs written for this purpose by Dr. V. A. Norton.

References

1. Robert Fricke, Felix Klein (1897) *Vorlesungen über die Theorie der automorphen Funktionen*. Leipzig: Teubner. (1965) Reprint. Johnson.
2. Benoit Mandelbrot (1982) *The Fractal Geometry of Nature*. San Francisco: W. H. Freeman. (Supersedes *Fractals: Form, Chance and Dimension*, W. H. Freeman, 1977.)
3. Henri Poincaré (1914). *Oeuvres*. Gauthier-Villars.
4. Benoit Mandelbrot (1981) *Mathematical Calendar*. New York-Heidelberg-Berlin: Springer-Verlag.

IBM Thomas J. Watson Research Center
Yorktown Heights, New York 10598