ON THE QUADRATIC MAPPING $z \to z^2 - \mu$ FOR COMPLEX μ AND z: THE FRACTAL STRUCTURE OF ITS $\mathscr{M}$ SET, AND SCALING

Benoit B. MANDELBROT

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

For each complex μ, denote by $\mathscr{F}(\mu)$ the largest bounded set in the complex plane that is invariant under the action of the mapping $z \to z^2 - \mu$. Mandelbrot 1980, 1982 (Chap. 19) reported various remarkable properties of the $\mathscr{M}$ set (the set of those values of the complex μ for which $\mathscr{F}(\mu)$ contains domains) and of the closure $\mathscr{M}^*$ of $\mathscr{M}$. The goals of the present work are as follows. A) To restate some previously reported properties of $\mathscr{F}(\mu)$, $\mathscr{M}$ and $\mathscr{M}^*$ in new ways, and to report new observations. B) To deduce some known properties of the mapping f for real μ and z, with $\mu \epsilon\ ]-\frac{1}{4}, 2[$ and $z \epsilon\ ]-\frac{1}{2}-\frac{1}{2}\sqrt{1+4\mu}$, $\frac{1}{2}+\frac{1}{2}\sqrt{1+4\mu}[$. In many ways, the properties of the transformation f are easier to grasp in the complex plane than in an interval. (This exemplifies the saying that "when one wishes to simplify a theory, one should complexify the variables".) C) To serve as introduction to some recent pure mathematical work triggered by Mandelbrot 1980. Further pure mathematical work is strongly urged.

Introduction. The illustrations are the focus of this paper, and the text is organized around the illustrations, in the form of extended comments. Additional illustrations are found in Mandelbrot 1980, 1982, 1983a, b.

1. Discussion of figs. 1a to 1e. Illustration of the action of $z \to f(z, \mu) = z^2 - \mu$ on a large complex circle. Sequences of algebraic curves approximating the repeller (Julia) sets $\mathscr{F}^*(\mu)$

A transformation becomes easier to study when one has a concrete visual feeling for its action. In the case of $z \to f(z, \mu) = z^2 - \mu$, it is known that the point at infinity is a stable fixed point of f. Hence it is an attractive point. In order for a circle of sufficiently large radius r and center O to be in the domain of attraction of ∞, a sufficient condition is $r > r_\mathrm{w} = \frac{1}{2}+\frac{1}{2}\sqrt{1+4|\mu|}$. The circle of radius r_w and center O will be denoted by W^O and called "whirlpool circle", and r_w will be called the "whirlpool radius", because the orbits of all the points outside W^O "whirl away" from O.

On the other hand, some complex z are not attracted to ∞. Examples are the bounded fixed points $z' = \frac{1}{2}+\frac{1}{2}\sqrt{1+4\mu}$ and $z'' =$

$\frac{1}{2}-\frac{1}{2}\sqrt{1+4\mu}$, and their successive pre-images. Let the maximal bounded set invariant under $f(z, \mu)$ be denoted by $\mathscr{F}(\mu)$. By definition of the whirlpool circles, $\mathscr{F}(\mu)$ is contained within W^O. Also, for every value of k, $\mathscr{F}(\mu)$ is contained in the kth pre-image of W^O under $f(z)$, i.e., the pre-image of W^O under $f_k(z)$. This last set is defined by $|f_k(z, \mu)| = r_\mathrm{w}$ and will be denoted by W^{-k}. It is an algebraic curve called "lemniscate" (Walsh 1956). The lemniscates corresponding to increasing values of k are non-overlapping and are monotonically imbedded in sequence. They can be called "parallel under $f(z, \mu)$". The set $\mathscr{F}(\mu)$ is the limit of these curves plus their interiors.

Denote by $\mathscr{F}^*$ the boundary of $\mathscr{F}$. The set $\mathscr{F}^*$ is the limit of W^{-k} for $k \to \infty$. It is the repeller set of $f(z, \mu)$ and is also called Julia set. (History: the earliest basic facts about iteration were described in Fatou 1906, and the bulk of the original theory was described near simultaneously in Julia 1918 and Fatou 1919. Since the term "Julia set" has become entrenched, I chose to honor Fatou by denoting this set by $\mathscr{F}^*$.)

Now for some illuminating illustrations. Figs. 1a to 1d represent the interiors of W^O and of several curves W^{-k} in superposition, for four selected values of μ. The goal is to demonstrate intuitively that the topology of $\mathscr{F}$ and of the Julia set $\mathscr{F}^*$
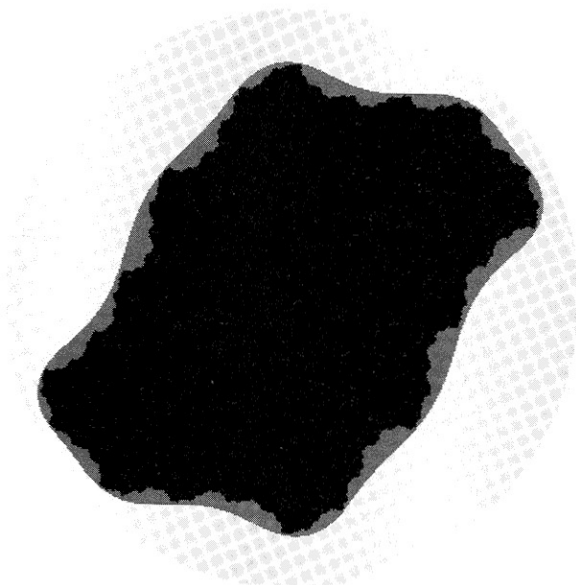
Fig. 1a. Action of $z \to z^2 - \mu$ when the Julia set is a loop-free closed fractal curve.
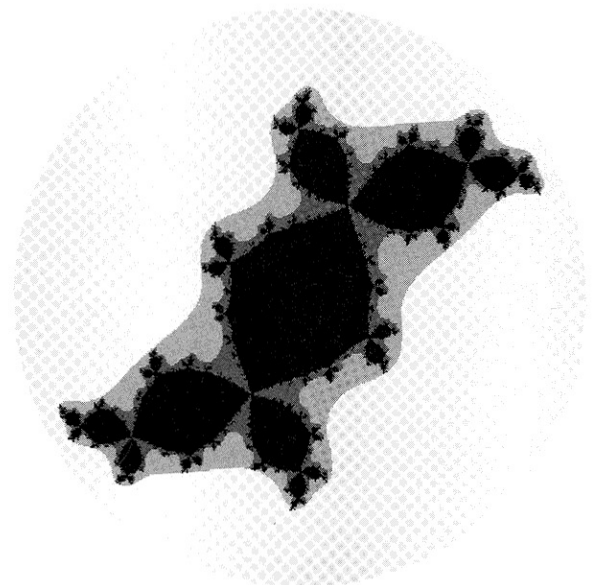

Fig. 1b. Action of $z \to z^2 - \mu$ when the Julia set is a closed fractal curve with loops.

greatly depends on the value of μ: in particular, $\mathscr{F}^*$ can be (a) a loop-free ("simple") curve that bounds a domain, (b) a curve with multiple points that bounds an infinite number of domains, (c) a tree ("branching curve without loop", "dendrite") that does not surround a domain, or (d) a totally disconnected dust. Fig. 1e represents an attractive example of $\mathscr{F}(\mu)$.

2. Discussions of figs. 2a to 2f. Classification of the values of μ by the topology of $\mathscr{F}(\mu)$. The sets $\mathscr{M}$ and $\mathscr{M}^*$. Sequences of algebraic curves approximating $\mathscr{M}^*$. The continent, islands, stellate structures, devil's causeways

This series of figures investigates in detail the set of those values of μ for which $\mathscr{F}(\mu)$ is connected.
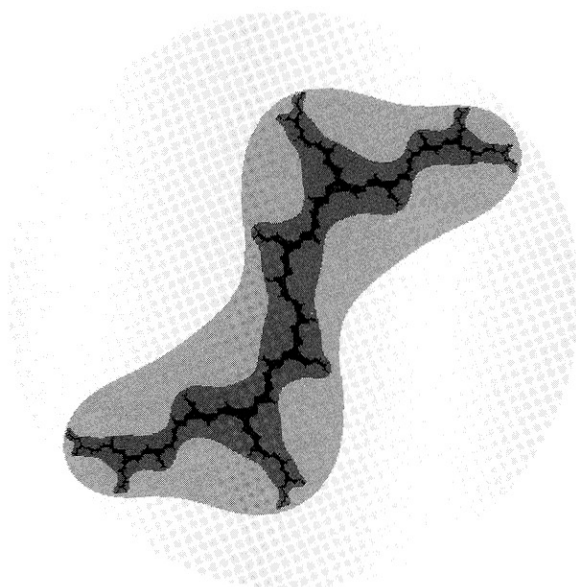

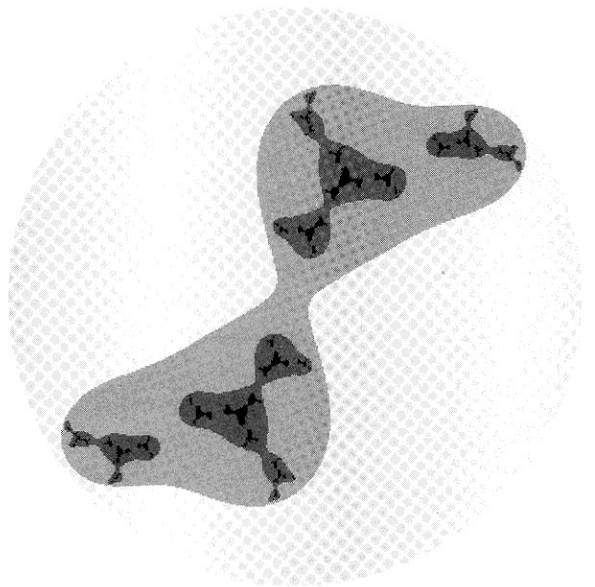Fig. 1c. Action of $z \to z^2 - \mu$ when the Julia set is a fractal tree.

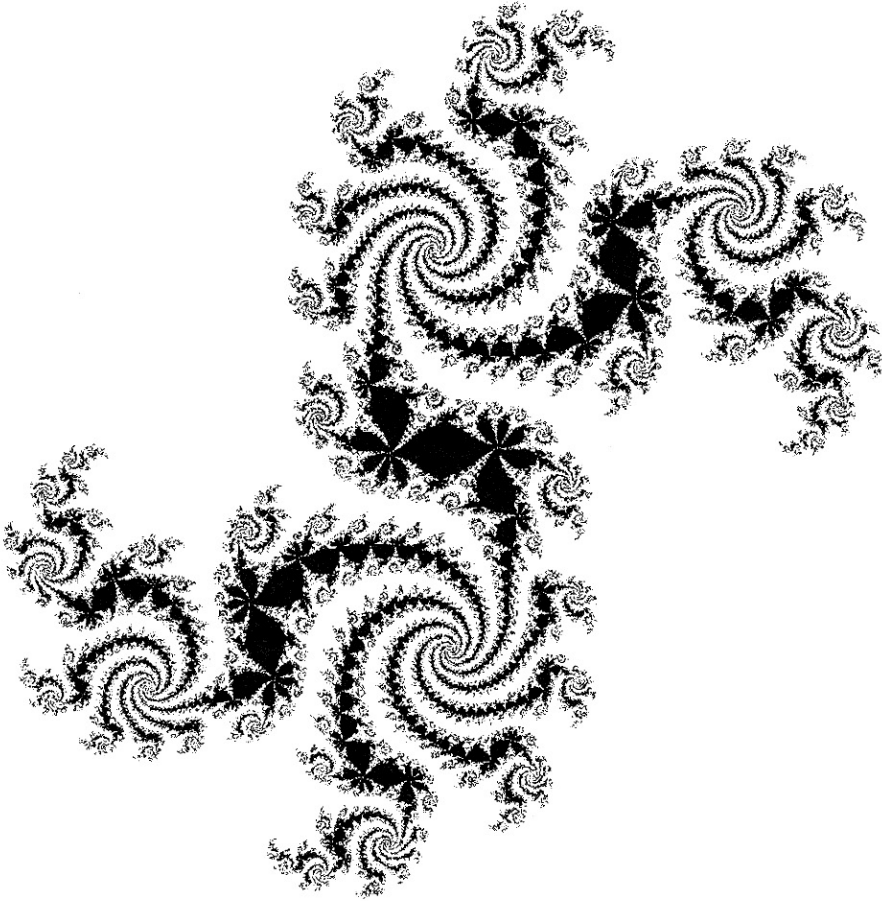Fig. 1d. Action of $z \to z^2 - \mu$ when the Julia set is a fractal dust ("Cantor set").

Fig. 1e. Interior of a Julia set (repeller set) of $z \to z^2 - \mu$ after two successive sixfold bifurcations.

This set is to be denoted by $\mathcal{M}^*$, and $\mathcal{M}$ will denote the set of values of μ for which $\mathcal{F}(\mu)$ has interior points, that is, includes domains. On a graph, e.g., on fig. 2a, the $\mathcal{M}$ set and the $\mathcal{M}^*$ set cannot be distinguished, but they turn out to be significantly different in structure.

Construction of the $\mathcal{M}$ set and of the $\mathcal{M}^*$ set. To follow the method used in figs. 1a to 1d would be cumbersome and unreliable, but is not necessary because Gaston Julia gave the following direct criterion. The set $\mathcal{F}(\mu)$ is disconnected if and only if the sequence of iterates of $z = 0$, beginning with $-\mu, \mu^2 - \mu$ and $((\mu^2 - \mu)^2 - \mu$, converges to infinity. For this to be the case, a necessary and

sufficient condition is that $|f_k(0, \mu)|$ must exceed for some value of k the whirlpool radius $r_W = \frac{1}{2} + \frac{1}{2}\sqrt{1 + 4|\mu|}$ derived in section 1. If $|\mu| > 2$, this condition is satisfied for $k = 1$. Hence the $\mathcal{M}^*$ set is entirely contained within the closed disc $|\mu| \leqslant 2$. Furthermore, the program is simplified (though the runs become a bit longer) if r_W is replaced by a uniform threshold equal to 2. For each k, one can draw in the μ-plane the set defined by $|f_k(0, \mu)| = 2$. As we know from section 1, such a set is an algebraic curve called "lemniscate". Here, all the lemniscates include the point $\mu = 2$, but otherwise they are non-overlapping and monotonically imbedded in sequence. The $\mathcal{M}^*$ set is the limit of these curves plus their interiors.
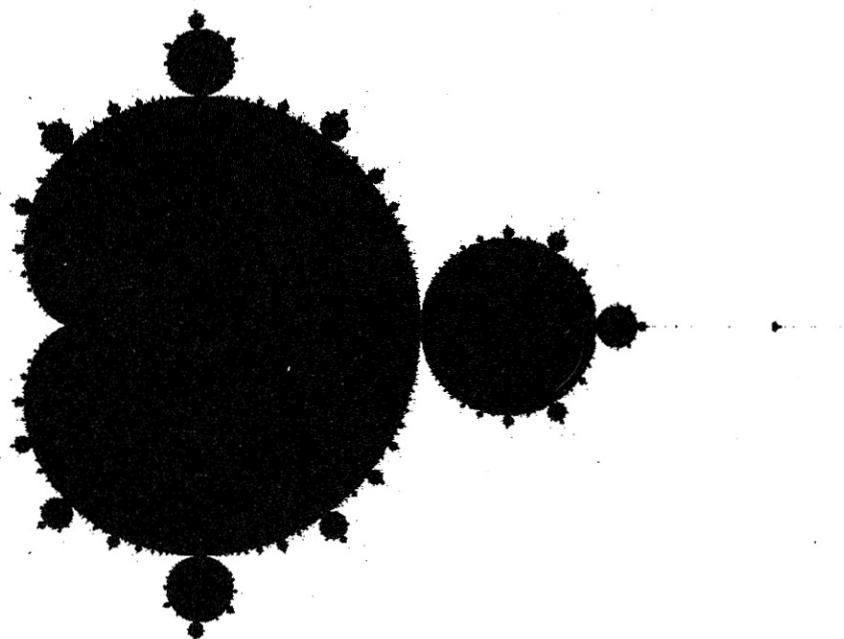
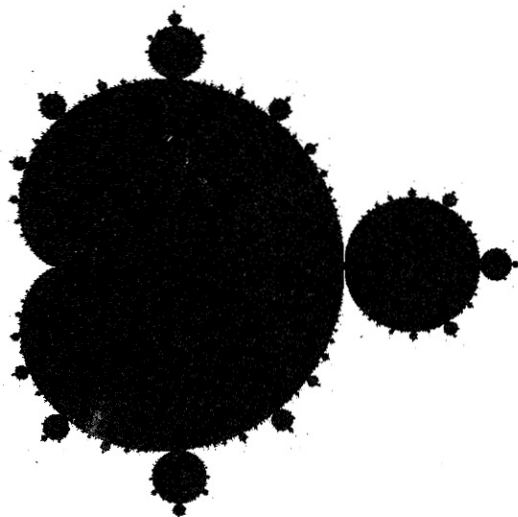Fig. 2a. Overall view of the $\mathscr{M}$-set of $z \rightarrow z^2 - \mu$.

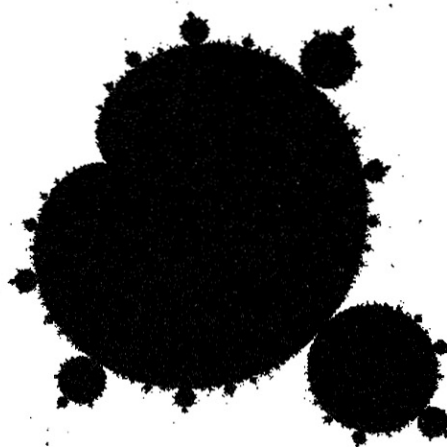Fig. 2b. Detail of an $\mathscr{M}$-island: the "speck" to the right of fig. 2a.

Fig. 2c. Detail of an $\mathscr{M}$-island: the "speck" at the bottom of fig. 2a.

As is known, μ is called superstable of minimal period k if $f_k(0, \mu) = 0$ but $f_h(0, \mu) \neq 0$ for every $h < k$. It follows that $f_{nk}(0, \mu) = 0$ for every integer n, hence all superstable μ's fail to iterate to infinity, meaning that they belong to the $\mathcal{M}^*$ set. Walsh 1956 reports that a lemniscate cannot contain a loop within a loop: it is necessarily either a single loop, or a finite union of loops with non-overlapping interiors. It turns out that in the present case, all the lemniscates are single loops for all values of k, hence $\mathcal{M}^*$ is a connected set. But before we tackle this point, other features of the $\mathcal{M}^*$ set must be considered.

The continental subset of $\mathcal{M}$. For reasons that will transpire momentarily, the structure of fig. 2a is clarified by positioning the grid of μ's so that real valued μ's are not tested. A first glance reveals that the great bulk of the black points lie in a large and very highly structured "continent". It has a striking "cactus tree" structure, which I propose to describe as a "molecule" made of an infinity of "atoms". At the center is a "seed atom", which has the shape of a cardioid, and contains all the μ's for which $f(z, \mu)$ has a single stable limit point besides ∞. The exactly circular atom straight to the right from the cardioid contains all the μ's for which $f(z, \mu)$ has a stable cycle of period 2, and the near circular atoms that follow to the right correspond to stable cycles of periods 4, 8, etc. The points where these atoms join are the μ's corresponding to the basic real μ bifurcations. Other near circular atoms that touch the cardioid correspond to cycles of order $k > 2$.

The shape of the $\mathcal{M}^*$ set near the value μ_∞ at the rightmost tip of the continental subset. Scaling property in the plane, and its use to rederive (as corollary) the known scaling property of bifurcations on the real line. Consider the sequence of atoms that converge to the tip of the continent. They seem essentially alike, and seem tangent to two straight half-lines that are symmetric with respect to the real axis and converge to the value μ_∞ defined as the accumulation point of bifurcations.

This is a geometric property of scaling, more precisely, of asymptotic geometric scaling.

An inferred consequence is that these atoms' horizontal intercepts decrease geometrically at each bifurcation. This inference is of course well-known to be true, having been discovered by Grossman and Thomae and by Feigenbaum.

To verify the identity of the atom shapes by a more exacting test, $\mathcal{M}^*$ was redrawn by replacing the parameter μ by $v = \log(\mu - \mu_\infty)$. The cardioid shaped seed atom is thereby made much smaller, and the other atoms indeed become near identical.

The big island to the right of the continent. Other islands. In addition to the continent, the $\mathcal{M}^*$ set contains a number of scattered specks. It is hoped that these specks escaped the watchful eye of the editors and the printers of the present Proceedings. The reason why I am concerned is that their counterparts on page 250 of Mandelbrot 1980 came to be erased, on the firm assumption that they could only be dirt!

In fact they are very real, and it may be useful to devote a few lines to telling how I discovered them. Examining my first rough graph of $\mathcal{M}^*$, I too took most of them to be dirt. But the biggest one, positioned to the right of the continent, looked too big to be spurious, and it was easy to verify that it intersects the axis of real μ's along the interval, discovered by Myrberg and Metropolis, Stein and Stein (see Collet and Eckmann 1980), for which $f(z, \mu)$ has a stable cycle of period 3. I had this speck examined in closeup, fig. 2b and it was revealed to be essentially a downsized version of the continent. Other Myrberg intervals that I examined in closeup were also revealed to intersect very small downsized versions of the continent.

Thus, the rightmost tip of the continent continues along the real axis by a peculiar causeway. Because of analogy with the Devil's Staircase (Mandelbrot 1982, page 83), I propose to call it the "Devil's Stepstones". The metaphor starts with large stones set in a stream to accommodate ordinary super giants, then smaller stones are set to accommodate ordinary giants, small giants, super-
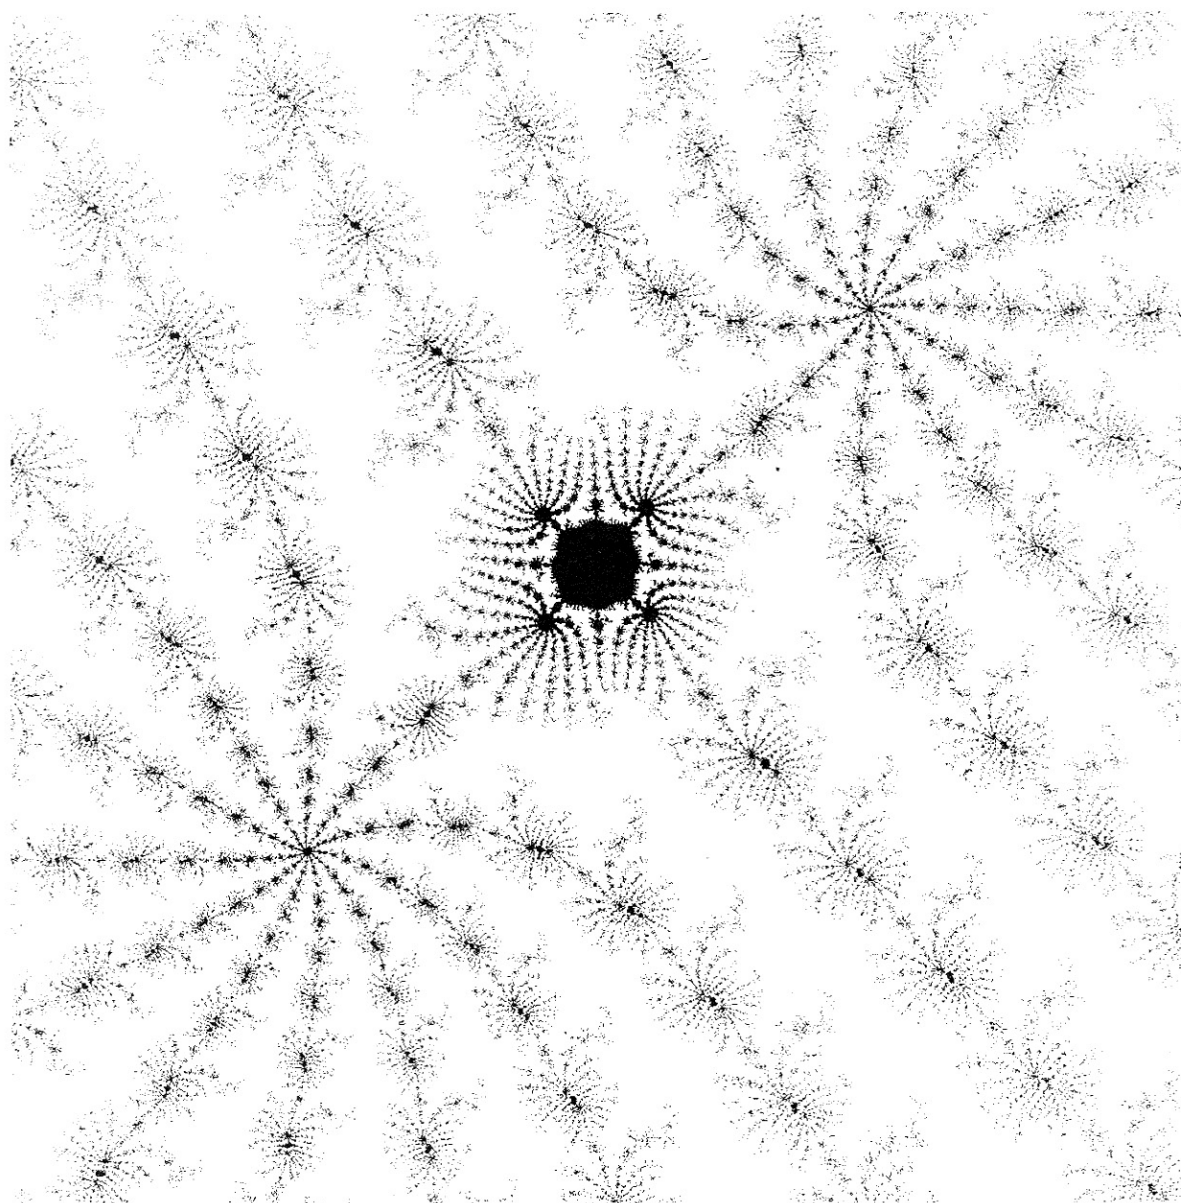
Fig. 2d. Detail of the $\mathcal{M}$-web offshore the continent on fig. 2a.

men, and so on down to devilishly tiny beasts. Ultimately the stones leave no gap of positive width, however small. The real axis runs along the center of this causeway, in a way that is familiar to students of the real transform $f(z, \mu)$.

At this point, I traced several puzzling observations to the same source. The first observation was that for periods 1, 2 and 3, each superstable μ is the "nucleus" of an atom known to belong to the continent or an island off the real axis. However, two superstable μ's of period 4 remained "unattached", and for higher periods the number of "unattached" superstable μ's kept increasing rapidly. One may have argued that some atoms con-
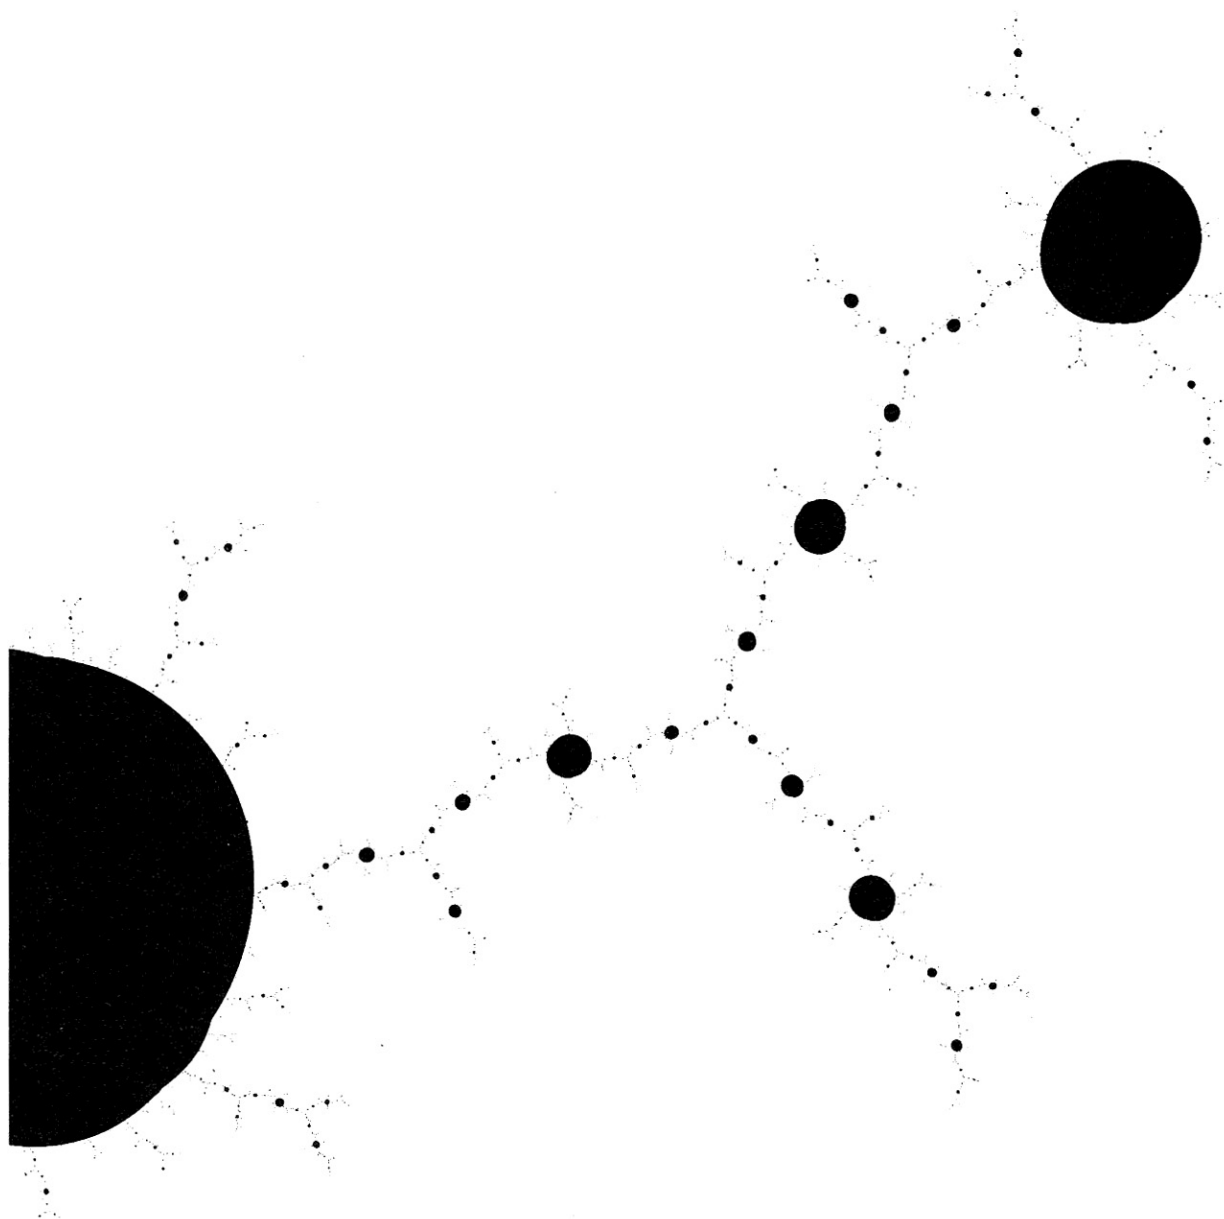
Fig. 2e. The interior of the Julia set when μ is near the nucleus of an offshore "island" of $\mathcal{M}$.

tain multiple-root nuclei, or several distinct nuclei, but these atoms should have looked different from atoms of smallest period 1, 2 or 3, while in fact all atom shapes fell into either of the two patterns exemplified by the seed cardioid and the circle to the right of it.

The second puzzling observation was that except for the point μ_∞, the tips of the continent gave no evidence of being followed by Devil's Stepstones.

The third puzzling observation was already mentioned; when the $\mathcal{M}$-set was traced with low precision on a medium-tight lattice, it seemed surrounded by unattached specks of dirt.

A close-up view of the big speck to the right of $\mathcal{M}^*$ (fig. 2c) settled the three puzzles together; most specks did not vanish but turned out to be islands identical to the continent in their topology and overall form. It soon became clear that these
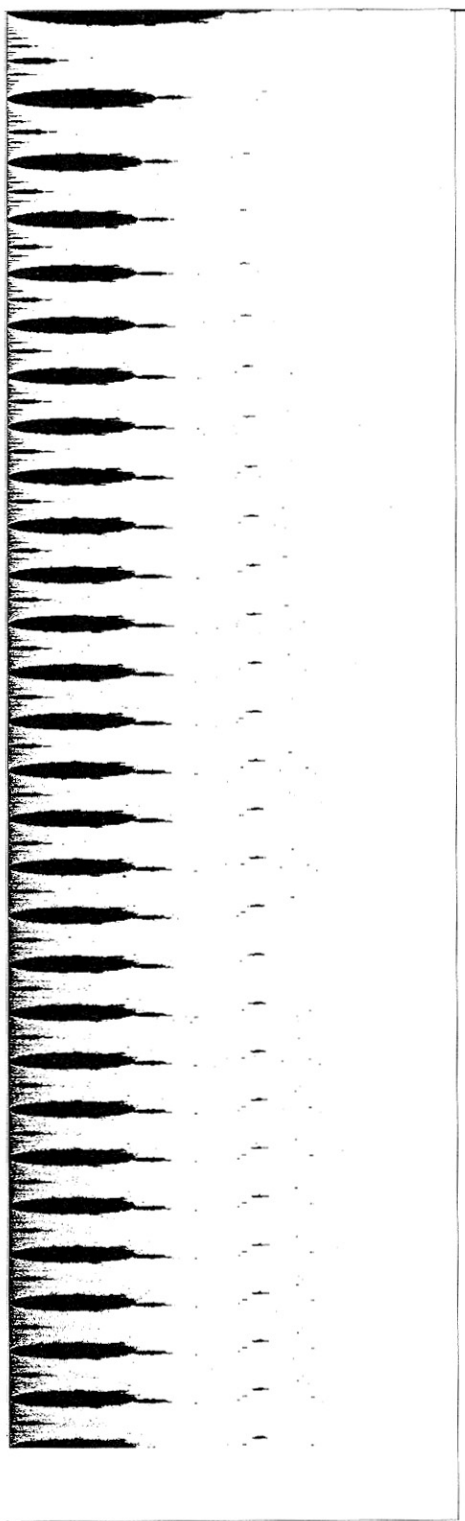
Fig. 2f. Detail of the $\mathcal{M}$ set of $z \to \lambda z (1 - z)$ after inversion and compression.

islands do not scatter around haphazardly, but form "stellate" arrays (see Mandelbrot 1980, 1982 for details). An array close to a point of bifurcation of order 11 is seen on fig. 2d. Further closeups revealed increasing numbers of increasingly small islands between larger islands along each "ray". This was reminiscent of the above-mentioned fact that increasingly small islands are "pierced through" by the real axis. At this point, the growing analogy with the real axis suggested that the islands in every ray are linked together by curves that are counterparts of the real axis, but could not be seen because curves other than the axis nearly always fall between the lattice points used in computation. This implies that the stellate structure reflects an underlying tree structure, and that the $\mathcal{M}^*$ set is connected. It may be recalled that $\mathcal{M}^*$ is approximated by a sequence of lemniscates, which (according to general results) might have split into separate loops. The fact that $\mathcal{M}^*$ remains connected implies that the approximating lemniscates are single loops. This was verified to be the case until high order pre-images of the circle $|\mu| = 2$, parts of which started falling between the lattice points.

Needless to say, computer based observations do not provide a substitute for actual proofs. In some cases, full proofs bring rigor without additional insight, but mathematical study of my observations on the $\mathcal{M}^*$ set turned out to be fruitful and useful, witness Douady and Hubbard 1982 and forthcoming works. Needless to say (again) this mathematical study could not have been undertaken without my computer based observations.

The inverse of bifurcation: the notion of confluence. The literature of bifurcation never seems to refer to the opposite effect that is observed, say, when μ starts with a value in an atom other than a seed atom, and changes continuously, without leaving the island, until it reaches a seed atom. Mandelbrot 1980 gave to this inverse operation the name *confluence*. The point is that the continent is the domain of confluence to a stable limit point, and each island is a domain of confluence to a periodic cycle, but not of confluence to a limit point.

The atoms' and the islands' intrinsic coordinates. Homologous points. Given an atom of minimal period k, denote by z_μ any point in the stable cycle corresponding to the parameter value μ. We know that the complex number $f'_k(z_\mu, \mu)$ is less than 1 in modulus. Its real and imaginary parts form intrinsic coordinates for the point μ within the atom to which it belongs. Two points having identical intrinsic coordinates can be called homologous within their atoms. The set of μ's for which $f'_n(z_\mu, \mu)$ is real will be called the atom's "spine". It runs from a point where $f'_k(z_\mu, \mu) = 1$ (which is a cusp in the case of seed atoms), to a point of bifurcation of order 2, where $f'_k(z_\mu, \mu) = -1$.

Furthermore, each atom's position in its island can be identified by an "address", namely the sequence of integers that identify the sequence of bifurcations that lead to this atom starting from the seed cardioid. Each bifurcation is indeed marked by a rational number n_i/m_i, with $m_i \geqslant 2$ and $0 < n_i < m_i$. Thus, it suffices to write these n_i and m_i in sequence separated by commas. One can agree that the seed cardioid's address is 0 (and other addresses may, but need not, start by 0). The combination of the address of the atom and of the value of $f'_k(z_\mu, \mu)$ forms an intrinsic coordinate for a point μ within the island to which μ belongs. Two points having identical intrinsic coordinates can be called "homologous" within their islands.

An island's spine combines its seed cardioid's spine with the spines of atoms corresponding to bifurcation into $m_i = 2$. Every island spine's endpoint is homologous to the tip μ_∞ of the continental subset of the $\mathcal{M}^*$ set.

"Universality class" argument to explain why the islands are alike. Assume that μ is near a superstable value μ^* of minimum period k. We wish to determine the shape of the atom nucleated by μ^*.

The lowest order terms in the expansion of $f_k(z, \mu)$ near $z = 0$ and $\mu = \mu^*$ can be written as $\beta_k z^2 + \gamma_k(\mu - \mu^*)$. Now let us state and test out a brutal "universality class" argument, then a milder version of it.

The brutal argument claims that the shape of the atom nucleated by μ^* depends only on the lowest

terms in the expansion of $f_k(z, \mu)$ near $z = 0$ and $\mu = \mu^*$. If this were the case, μ^* would nucleate a cardioid-shaped seed atom. This atom and the molecule grown upon it would be identical to the continent, except for its size being reduced in the ratio $1/\beta_k \gamma_k$. The milder argument agrees to take account of a few higher order terms near $z = 0$ and $\mu = \mu^*$, while continuing to disregard the behavior of $f_k(z, \mu)$ far from $z = 0$ and $\mu = \mu^*$. This milder argument suggest the following properties. A) The atom nucleated by μ^* is the seed atom of a molecule, and its shape resembles the continental cardioid except for some mild non-linear deformation. B) Other atoms obtained by bifurcation are arrayed around this seed as on the continent, except again for a mild deformation.

Inspection of the actual $\mathcal{M}$ set indicates that when the prediction A) is correct, B) is also correct. Moreover, A) can only fail by A) and B) being replaced by the following properties. A') The atom is not a seed and its shape is near circular. B') Other atoms obtained by bifurcation are arrayed around the atom nucleated by μ^*, in the same way as their counterparts are arrayed around the continent, except for a transformation that straightens out the cusp.

Example: the superstable values for $n = 2$ are the roots of $\mu^2 - \mu = 0$, that is $\mu^* = 0$ and $\mu^* = 1$. Near $\mu^* = 1$, $f_z(z, \mu) = (z^2 - \mu)^2 - \mu = z^4 - 2\mu z^2 + \mu^2 - \mu \sim -2z^2 + (\mu - 1)$. This suggests an atom equal to the basic cardioid downsized in the ratio of $\frac{1}{2}$ and translated to the right by 1. But the actual atom is bigger in every direction and happens to be precisely a disc.

One may expect to find that the condition of validity of the milder universality class predictions A) and B) is that $\beta_k \gamma_k$ be large.

A further universality class argument (not well developed as yet) suggests that atoms increasingly removed from the seed of their islands tend to the universal shape.

"Universality class" argument to explain the shape of $\mathscr{F}^*(\mu)$ when μ^* lies in an island. The brutal universality class argument also makes a prediction concerning $\mathscr{F}^*(\mu)$: C) The Julia set $\mathscr{F}^*(\mu^*)$, call it

a "little dragon", obtains by reducing in the ratio β_k the Julia set which the full $\mathscr{F}(z, \mu)$ predicts for the point that lies in the continent and is homologous to μ, namely $\mu' = \beta_k \gamma_k (\mu - \mu^*)$. The milder universality class argument makes the prediction C′) The portion of the Julia set $\mathscr{F}^*(\mu)$ near $z > 0$ obtains by reducing in the ratio β_k the Julia set which the full $\mathscr{F}(z, \mu)$ predicts for the point that lies in the continent and is homologous to μ.

Inspection of actual sets $\mathscr{F}^*(\mu)$ indicates that the portion of $\mathscr{F}^*(\mu)$ near $z = 0$ is indeed the little dragon predicted by the milder C′). But the brutal prediction C) gives a quite incorrectly pallid idea of the structure of the whole of $\mathscr{F}^*(\mu)$. This set does *not* reduce to the little dragon near $z = 0$, but is made up of an infinity of mildly deformed replicas of this little dragon.

These replicas from Devil's Stepstones with the same structure we already encountered in the shape of $\mathscr{M}^*$. (Fig. 2e is relative to a case where μ is very close to the nucleus of the cardioid on fig. 2c.) That is, these replicas are strung along a tree. As to this tree's shape, it brings in something entirely foreign to the universality argument. Indeed, this shape is determined by μ, and not merely by the point in the continent that is homologous to μ. This shape varies fairly slowly with μ, and is approximately determined by μ^*.

To introduce an even rougher but useful approximation, let us begin by bringing in the parameter value μ'' corresponding to the tip of the island containing μ^*. This point is homologous to the classical real point μ_∞ at the tip of the continent. In the next section's discussion of the Julia sets $\mathscr{F}^*(\mu)$, we shall see that μ_∞ is among the values of μ for which $\mathscr{F}^*(\mu)$ is a tree having a real interval as spine, and other ribs but no flesh. When μ'' is the tip of an island other than the continent, $\mathscr{F}^*(\mu'')$ is also a tree (though it contains no straight interval). Now we come back to $\mathscr{F}^*(\mu)$: it is found that the replica dragons belonging to this set string along a tree approximated by $\mathscr{F}^*(\mu'')$.

Rough estimates of the counterparts of the ratio δ for bifurcations of order > 2. For the purpose of

this subsection, it is best to change the coordinates by replacing the parameter μ by the parameter $\lambda = 1 \pm \sqrt{1 + 4\mu}$. This corresponds to the mapping $z \to f^*(z, \lambda) = \lambda z(1 - z)$. The corresponding transform of the $\mathscr{M}$ set is shown on page 250 of Mandelbrot 1982. The continent is no longer of the same shape as the islands, since, instead of being seeded by a cardioid, it is now seeded by two discs. But this transformed shape has its own assets. A first advantage, which had been ascertained with pen and paper, is that the bifurcation from a stable fixed point to a cycle of period m recurs at the points where either λ or $2 - \lambda$ is of the form $\exp(2\pi \, in/m)$, with n an integer less than m.

A second advantage only transpired after the $\mathscr{M}$ set had been computed and could be examined. It was observed that the tips of the "major" sprouts around the circle $|\lambda - 2| = 1$, defined as the sprouts rooted at the points $\exp(2\pi \, i/m)$, appear to be placed along a larger circle whose diameter begins at $\lambda = 1$ and ends somewhere beyond $\lambda = 3$. This suggests that one perform an inversion of the $\mathscr{M}$ set with respect to $\lambda = 1$, with $\lambda = 3$ remaining fixed. This inversion should yield sprouts placed between parallel lines. Furthermore, the transform of the root of the mth major sprout should lie at a vertical distance from $\lambda = 3$ equal to $2 \tan(\frac{1}{2}\pi - \frac{1}{2}2\pi/m) = 2 \cotan(\pi/m)$. For large m, this yields $2m/\pi$, i.e., a series of equally spaced points.

The inverted $\mathscr{M}$ set shown in fig. 2f is plotted using very different units along the two axes, so that the graph remains legible yet covers many values of m. The above hunch is confirmed, except for $m = 2$ and 3. That is, an extrapolation from sprouts with a higher value of m would yield a smaller sprout for $m = 2$. Denote by A the height of the inverted sprouts for larger m. Assuming circular atoms, these properties of inversion yield the result that the relative linear size of the sprout of order m is $A \sin^2(\pi/m)[2 - A \sin^2(\pi/m)]^{-1}$. This is roughly the ratio of successive absolute changes in μ between bifurcations of order m, that is, the mth counterpart of the $1/\delta$ ratio of Grossman and Thomae and of Feigenbaum. In fact, this ratio is for all m close to m^{-2}.

3. Discussion of figs. 3a to 3f. The repeller stack for real μ and complex z. Illustrations of the influence of the value of μ on the shape and the topology of the repeller (Julia) set $\mathscr{F}^*(\mu)$

In figs. 3a to 3d, the horizontal coordinates x and y are the real and imaginary parts of z, and the vertical coordinate is μ. The figures represent a stack of Julia sets $\mathscr{F}^*(\mu)$ for μ ranging from $-\frac{1}{4}$ to 2. The goal is to show that the shape of $\mathscr{F}^*(\mu)$ varies continuously, while the topology of $\mathscr{F}(\mu)$ moves around discontinuously. The stack was sliced along the plane $xO\mu$, and the two halves have been separated.
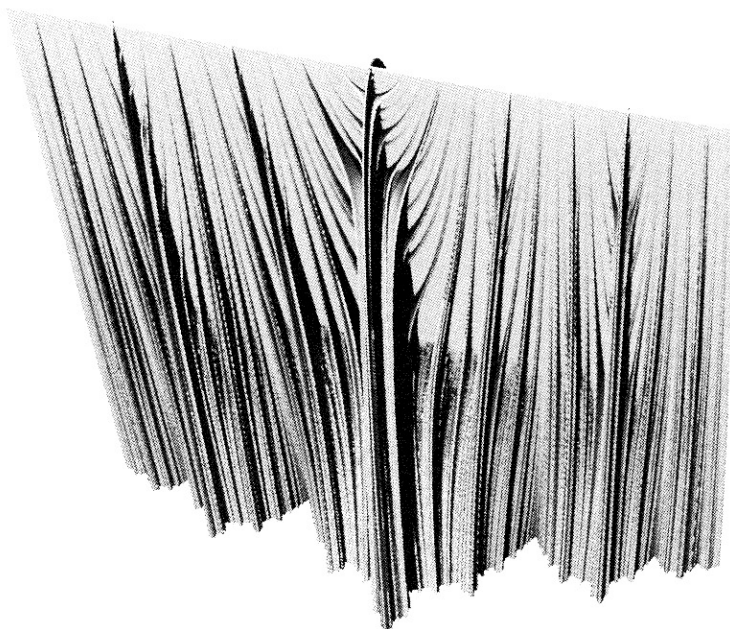

Fig. 3a. Perspective view of the top portion of the repeller stack of Julia sets for real μ (vertical) and complex z (horizontal).

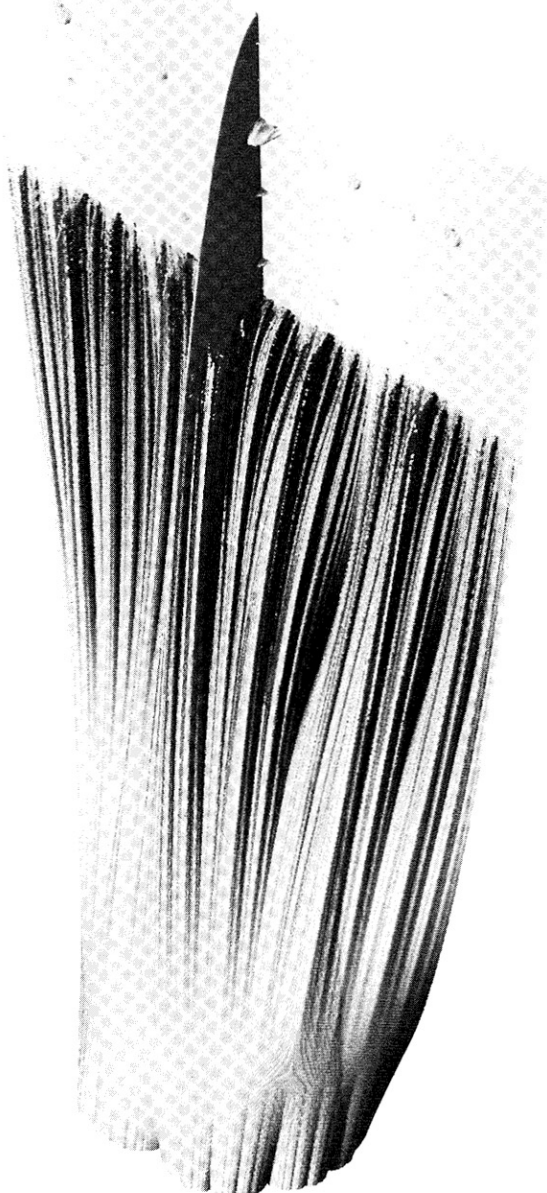Fig. 3b. Top portion of the "veils" within fig. 3a.

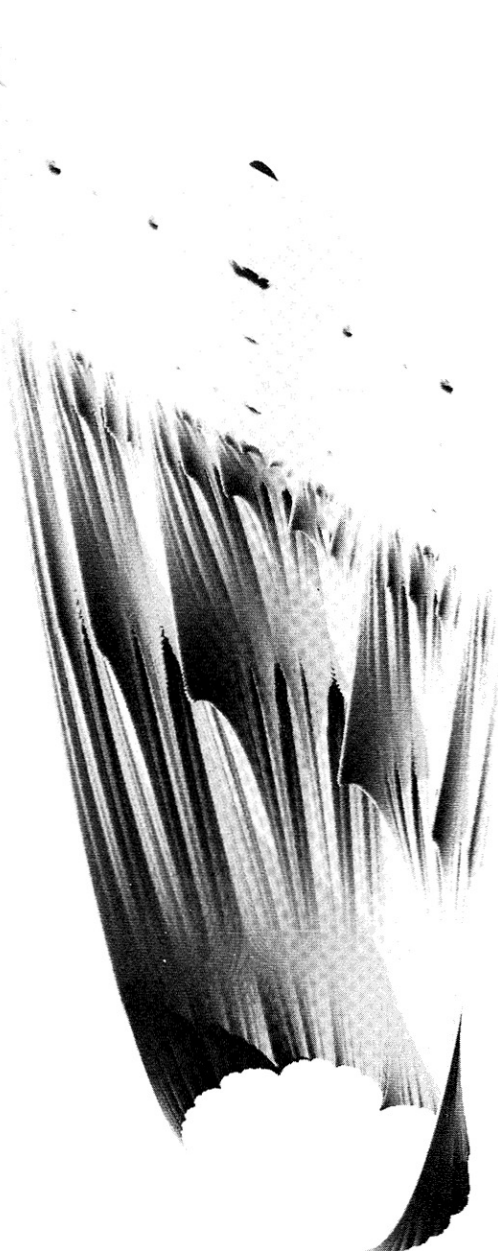Fig. 3c. Perspective view of the repeller stack minus the web: outside view.

Fig. 3d. Perspective view of the repeller stack: inside cut showing bifurcations.
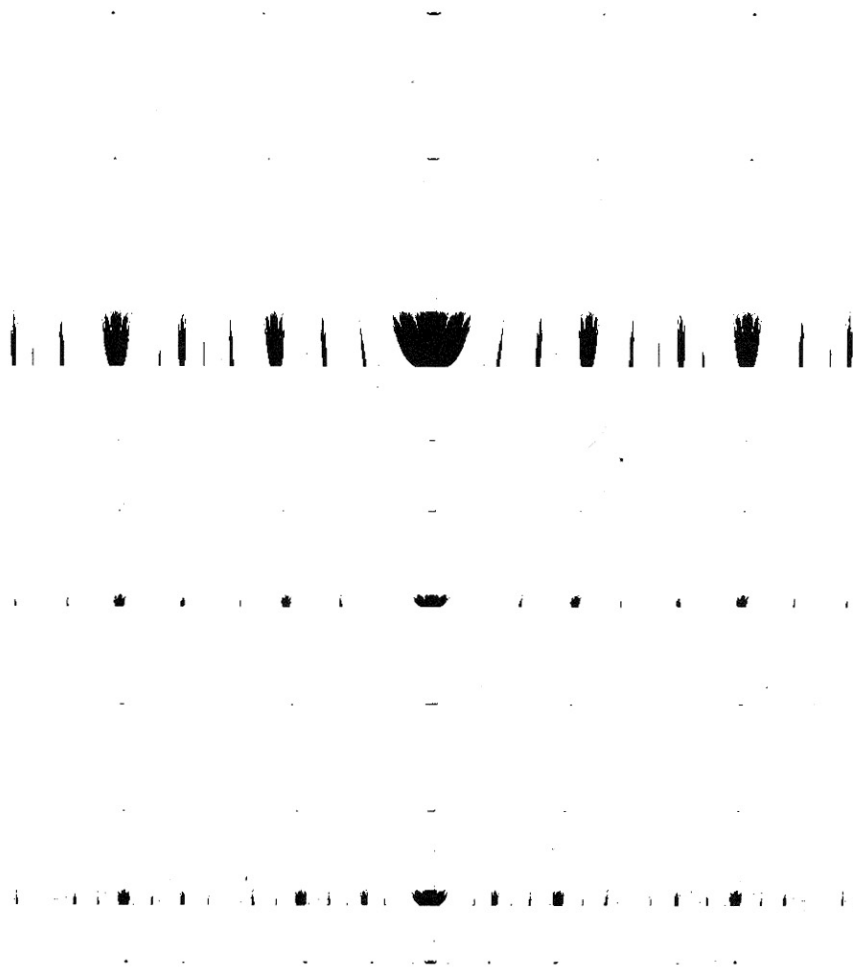
Fig. 3e. Detail of the top portion of fig. 3c: cut along the plane $y = 10^{-3}$ with x and μ as coordinates.

Computation of the stack. The theory of $\mathscr{F}^*(\mu)$ involves two "proof-of-existence" constructions. The first is used in figs. 1a to 1d. The second consists in tracing the pre-images of the unstable fixed point $z' = \frac{1}{2} + \frac{1}{2}\sqrt{1 + 4\mu}$. This second construction is efficient only if μ is near 0, i.e., when $\mathscr{F}^*(\mu)$ is an uncomplicated loop. In general, either construction requires prohibitively long computer runs to yield an acceptable approximation. For the sake of efficient computation, it was found best to devise several alternative constructions and to use them in combination. After the fact, these programs turned out to help in understanding the facts. Fig. 3a combines some "veils" and a "shell", while fig. 3b represents the veils alone (with fewer stages for the sake of clarity), and figs. 3c and 3d represent the shell alone.

The ribs and veils. For each μ, the backbone of the horizontal section of the stack is the real interval from $]-z', z'[$. The other ribs are the pre-images of $]-z', z'[$ under $f(z)$; fig. 3b, shows them up to order 8. Since $]-z', z'[$ is well known to fail to converge to ∞ under $f(z)$, the backbone and the ribs belong to the set $\mathscr{F}(\mu)$ and can be said to form its "skeleton". The ribs corresponding to different
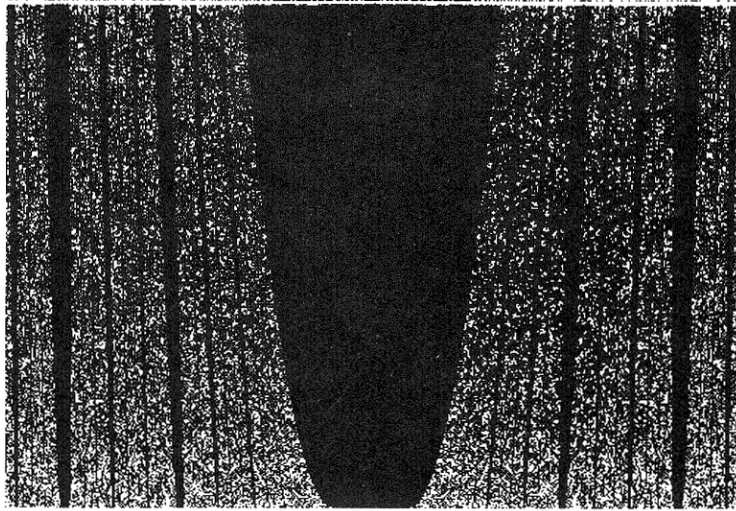
Fig. 3f. Even finer detail of the top portion of fig. 3c: cut along the plane $y = 10^{-10}$ with x and μ as coordinates.

μ's merge together to form a series of "veils". They include a square wall in the plane $y = 0$, and a rounded wall in the plane $x = 0$. Moving through a superstable μ, the veils change from hanging on the rounded wall to hanging on the square wall, or from hanging on a high-order veil to hanging on one of lower order. The pre-images of the unstable fixed point z' are the rib tips. The precise relationship between the ribs' closure and the set $\mathscr{F}(\mu)$ depends on the value of μ.

Superstable μ's. For superstable values of μ, the ribs' closure is a domain, and is identical to $\mathscr{F}(\mu)$. Hence, (by the same anatomical analogy) $\mathscr{F}(\mu)$ can be said to include no proper flesh. The obvious example is $\mu = 0$, when $\mathscr{F}(\mu)$ is the disc of unit radius, and the kth order ribs are segments joining 0 to the points of the form $\exp(2^{-k}\pi \text{ i}/n)$, with n an integer and $0 < n < 2^{k+1}$.

Chaotic μ's. For the chaotic values of μ, the closure of the ribs is a tree, and is again identical to $\mathscr{F}(\mu)$. The obvious example (though a degenerate one) is $\mu = 2$. Indeed, the set $\mathscr{F}(2)$ and its ribs both reduce to the backbone $[-2, 2]$. To obtain $\mathscr{F}(2)$, it is

obviously faster to draw the backbone than to use the proof of existence construction that dots $\mathscr{F}(2)$ with the dense pre-images of $z = 2$.

Whenever μ is close to either a superstable or a chaotic value, the maximal invariant set $\mathscr{F}(\mu)$ is rapidly approximated by only a few levels of ribs. Since $\mathscr{F}(0)$ is simply a disc and all the other superstable or chaotic μ's fall in $]1, 2[$, it was found best to draw a few levels of ribs for every μ in $]1, 2[$ (anyhow the cost of "unnecessary" computation is less than the cost of determining whether or not the computation was worth performing).

Stable but not superstable μ's. The remaining real values of $\mu \epsilon] -\frac{1}{4}, 2[$ are the μ's for which a stable fixed point or a finite period exists but is not superstable. For these values of μ, the set of ribs is not dense in a domain and does not remain a tree even in the limit. To describe the resulting structure of $\mathscr{F}^*(\mu)$, let us mix the previous anatomical metaphor with a botanical one: we can say that for these μ's, the trees' branches join asymptotically to form a "canopy". Clearly, an inspection of the pre-images of the unstable fixed point z' could not distinguish between the cases when the branch tips

are disconnected and form dusts, and cases where the branch tips are connected.

The shell. As already mentioned, the proof of existence construction of $\mathscr{F}^*(\mu)$ via the pre-images of z' is efficient when $\mathscr{F}^*(\mu)$ is an uncomplicated loop, that is, for μ near $\mu = 0$. Whenever $\mathscr{F}^*(\mu)$ is even moderately kinky, the cusp shaped kinks remain unfilled even after other portions of $\mathscr{F}^*(\mu)$ have been covered many times over. (For the cognoscendi: the reason is that this method reconstitutes the invariant measure on the Julia set, and this measure can be extraordinarily uneven.)

An efficient graphic method is one that spends roughly equal times on each portion of $\mathscr{F}^*(\mu)$. The shell in fig. 1 was drawn by the following shell generator (Norton, 1982). Each horizontal plane was covered with a square lattice and the position of a lattice point was saved in computer memory whenever, a), its kth iterate falls within a circle of radius 2, and, b), the kth iterate of at least one of its neighbors falls outside of that circle. These points are identified by a search method that starts with the unstable fixed point z', and is very efficient, because the number of wasted points (tested but not saved) is only a small multiple of the number of points that are saved.

Unfortunately, whenever $\mathscr{F}^*(\mu)$ is *very* kinky, as is the case for $\mu \epsilon\,]1, 2[$, the shell generator misses many points in $\mathscr{F}^*(\mu)$. It misses "A-pieces" that (by definition) are so thin that they squeeze between the lattice points. And it misses "B-pieces" that (by definition) are large but connect to the unstable fixed point z' through A-pieces.

When the shell is examined from the inside, fig. 2d, these A- and B-pieces above do not matter, because they would be hidden anyhow. And inclusion of the ribs would hide the evidence discussed below. When, to the contrary, the shell is examined from the outside, Figs. 3a, b, c, A- and B-pieces do matter. Luckily, many of the points missed by the shell generator are picked up by the rib generator, and the combination of the two yields a sensible idea of the outside shape of the stack. (Note that I had not originally planned to split this figure

open, but an inside view was computed by mistake, a *felix culpa*.)

Basic observations. As μ increases, the repeller set varies continuously, but its topological characteristics change back and forth. The largest invariant set varies continuously within each Myrberg interval of μ. However, let μ decrease through the value $-\frac{1}{4}$, or through a value homologous to $-\frac{1}{4}$ within an island. The result is that a continuous canopy becomes punctured, leaving a dust and allowing the flesh to "evaporate". On opposite sides of a chaotic μ, the tree tips combine into canopies in different fashions.

Near $\mu = 0$, the shell is extremely smooth. More generally, as μ moves away from $\mu = 0$, the unsmoothness of $\mathscr{F}^*(\mu)$ increases very slowly and gradually. This led me to conjecture that the fractal dimension of $\mathscr{F}^*(\mu)$ is a very regular function of μ: infinitely differentiable and perhaps analytic. This hunch was proven true in Ruelle 1982.

Interpretation of the bottom portion of the inside view of the shell. Bifurcations. Aside from a clearly visible circle for $\mu = 0$, the most striking feature of this view resides in rows upon rows of protuberances. The lowest row lies at height $\mu = \frac{3}{4}$. For real μ below $\frac{3}{4}$, the mapping $f(z)$, whether in real or complex z, has a stable limit point. For $\mu = \frac{3}{4}$, this stable limit bifurcates into a stable limit cycle of period 2. At the same time, one sees that $\mathscr{F}^*(\mu)$ changes from being a simple loop to being an infinitely knotted one. (The protuberances are denumerable, and have only two limit points: the fixed point z' and $-z'$.) The next highest row of protuberances marks the second bifurcation, and so on.

The reader is surely acquainted with Robert May's tree diagram, May 1976, which maps the variation with μ of the values of x for all the points in the corresponding stable cycle. Were it superposed on fig 2d, this tree would be rooted at the sharp tip to the bottom left, and each branch point would hang on a suitable protuberance. One may extend May's diagram to map the variation with μ

of the "real preorbit of the cycle", defined as the set of real z that eventually fall exactly in the cycle. The resulting preorbit map diagram would be made of many trees, with a branch hanging on every protuberance of fig. 2d.

Interpretation of the top portions of the inside and outside views of the shell. The Myrberg intervals of μ. The top of the stack is characterized by a nearly blank wall, which we know from the veil generating construction. However, this wall is interrupted by mysterious hanging "knobs" forming horizontal strips. Each strip corresponds to a Myrberg interval of values of μ.

Interpretation of some horizontal or vertical sections of the stack. When μ lives in a Myrberg interval, a horizontal section of the stack is a tree formed of Devil's Stepstones. Since the stepstones vary continuously with μ, those which intersect the plane $x = 0$ would form a kind of Devil's Corduroy.

Now take vertical sections. The $y = 0$ section of the whole stack is bounded to the side by the half parabola $\mu = x^2 - 2x$ for $x > \frac{1}{2}$, the half parabola $\mu = x^2 + 2x$ for $x < -\frac{1}{2}$, and the segments {from $(\mu = -\frac{1}{4}, x = -\frac{1}{2})$ to $(\mu = -\frac{1}{4}, x = \frac{1}{2})$} and {from $(\mu = 2, x = -2)$ to $(\mu = 2, x = 2)$}. The bottom of this vase-shaped outline is filled solid and the top is surmounted by Myrberg strips.

Now consider analogous vertical sections of the top of the shell for $y = 10^{-3}$ (fig. 3e) and of a detail for $y = 10^{-10}$ (fig. 3f). Here, one sees a large number of black shapes, each of them a deformed version of the vase shape of the overall stack. When μ is such that the iterates of $f(z, \mu)$ are not chaotic, the intersection of $\mathcal{F}^*(\mu)$ and the wall $y = 0$ is a denumerable set. Its points of accumulation number 2 when μ lies in the continent, and

are themselves denumerable when μ lies in an island. When μ is chaotic, the intersection of $\mathcal{F}^*(\mu)$ and the wall is an interval.

Acknowledgment

References

P. Collet and J.P. Eckman, Iterated Maps on the Interval as Dynamical Systems, Birkhauser (Boston, 1980).

A. Douady and J. Hubbard, Comptes Rendus (Paris) 294-I (1982) 123–126.

P. Fatou, Sur les solutions uniformes de certaines équations fonctionnelles, Comptes Rendus (Paris) 143 (1906) 546–548.

P. Fatou, Sur les équations fonctionnelles, Bull. Société Mathématique de France 47 (1919) 161–271; 48 (1920) 33–94; 48 (1920) 208–314.

G. Julia, Mémoire sur l'itération des fonctions rationnelles, J. de Mathématiques Pures et Appliquées 4 (1918) 47–245. Reprinted (with related texts) in: Julia Oeuvres (1968) I 121–319.

B.B. Mandelbrot, Fractal aspects of the iteration of $z \to \lambda z(1 - z)$ for complex λ and z. Nonlinear Dynamics, R.H.G. Helleman, ed., Annals of the New York Academy of Sciences 357 (1980) 249–259.

B.B. Mandelbrot, The Fractal Geometry of Nature (Freeman, San Francisco, 1982).

B.B. Mandelbrot, Forthcoming article in Scientific American (1983a).

B.B. Mandelbrot, Forthcoming monograph (1983b).

R. May, Simple mathematical models with very complicated dynamics, Nature 261 (1976) 459.

V.A. Norton, Generation and display of geometric fractals in 3-D, Computer Graphics 16 (1982) 61–67.

D. Ruelle, Repellers for real analytic maps, Ergodic theory and dynamical systems 2 (1982) 99–107.

J.L. Walsh, Interpolation and approximation by rational functions in the complex domain, American Mathematical Society Colloquium Publication, number 20 (1956).