6 Fundamental domains and examples of Kleinian groups

6.1 Fundamental domains

Let G be a Kleinian group, acting on $\mathcal{H}^3_+$, on $\hat{\mathbb{C}}$, or on $\mathcal{H}^3_+ \cup \hat{\mathbb{C}}$, and let $\Omega(G)$ be the ordinary set for the action.

Definition A *fundamental domain* for the action of G on $\Omega(G)$ is a subset F of $\Omega(G)$ such that

(i)
$$\bigcup_{g \in G} g(\bar{F}) = \Omega(G) \quad \text{and}$$

(ii)
$$g(F) \cap h(F) = \emptyset \quad \text{when} \quad g \neq h \quad (g, h \in G)$$

(where in (i), $\bar{F}$ denotes the closure of F).

Thus the images of F *tesselate* $\Omega(G)$ (they cover it without overlapping).

Example The set $\{x + iy : 0 < x < 1\}$ is a fundamental domain for the action of $z \to z + 1$ on the complex plane $\mathbb{C}$ (as indeed is the set $\{x + iy : 0 \leq x < 1\}$).

Note The precise definition of the term 'fundamental domain' varies from author to author: some require F to be closed - in which case of course one must modify condition (ii) above to require only that $g(F) \cap h(F)$ be contained in the *boundary* of both $g(F)$ and $h(F)$, rather than it be empty.

6.2 Dirichlet domains

The simplest construction of fundamental domains makes use of a metric. So for the time being we consider an action of G on $\mathcal{H}^3_+$ (or, if G is Fuchsian, on $\mathcal{H}^2_+$).

Choose $x \in \mathcal{H}^3_+$ such that for all $g \in G$ except the identity, $gx \neq x$. (Exercise: show that there are at most a discrete set of points $x \in \mathcal{H}^3_+$ which do not have this property.) Now for each $g \in G$ define the *half-space*

$$H_g = \{y \in \mathcal{H}^3_+ : d(y, x) < d(y, gx)\}$$

where $d(y, x)$ denotes the hyperbolic distance from y to x.

Definition The *Dirichlet domain centred at* x is the set

$$D_x = \bigcap_{g \in G - \{I\}} H_g$$

Thus D_x consists of those points of $\mathcal{H}^3$ which are nearer to x than they are to any gx $(g \in G - \{I\})$.

This construction was introduced by Dirichlet in the 1850's for the study of Euclidean groups, and later adapted by Poincaré for the hyperbolic case.

Proposition 6.1 *For any Kleinian group G, a Dirichlet domain D_x is a fundamental domain for the action of G on $\mathcal{H}^3_+$.*

Proof. We must prove that D_x satisfies conditions (i) and (ii) of the definition of a fundamental domain. We first observe that

$$g(D_x) = \{y : d(y, gx) < d(y, hx) \quad \forall h \in G - \{g\}\}$$

since

$$y \in g(D_x) \Leftrightarrow g^{-1}y \in D_x \Leftrightarrow d(g^{-1}y, x) < d(g^{-1}y, kx) \Leftrightarrow d(y, gx) < d(y, gkx) \; \forall k \in G - \{I\}$$

Now take any $y \in \mathcal{H}^3_+$. Take $g \in G$ (not necessarily unique) such that $d(y, gx)$ is minimal. Then $y \in g(\bar{D}_x)$ so property (i) holds. Moreover it is clear that $g(D_x) \cap h(D_x) = \emptyset$ if $g \neq h$ so property (ii) holds too. QED

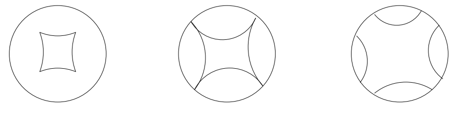

Figure 10: Polygons (in the Poincaré disc model)

Recall that a subset $X \subset \mathcal{H}^3_+$ is said to be *convex* if given any $x, y \in X$ the segment of geodesic joining x to y is entirely contained in X.

Proposition 6.2 *A Dirichlet domain D_x for a Kleinian group G is convex and locally finite (i.e. each compact subset K of $\mathcal{H}^3_+$ meets only finitely many $g(D_x)$).*

Proof. Convexity is obvious since D_x is defined to be an intersection of half-spaces, each of which is convex. For local finiteness, take the Poincaré disc model of $\mathcal{H}^3_+$ and without loss of generality take x to be the origin and K to be the closed ball with centre the origin and (hyperbolic) radius ρ. We claim that if g is any element of G such that $gD_0 \cap K$ is non-empty then $d(0, g0) \leq 2\rho$, which will prove local finiteness since G, being discrete, contains only finitely many elements with $d(0, g0) \leq 2\rho$ (else the orbit of 0 would have an accumulation point in $\mathcal{H}^3_+$, contradicting discontinuity of the action of G there). To prove the claim, take any $y \in gD_0 \cap K$; then $d(0, y) \leq \rho$ (since $y \in K$) and $d(g0, y) \leq d(0, y)$ (since $y \in gD_0$) so $d(0, g0) \leq \rho + \rho = 2\rho$. QED

Definition A convex region P obtained as the intersection of countably many half spaces H_j in $\mathcal{H}^3_+$, with the property that any compact subset of P meets only finitely many of the hyperplanes ∂H_j is called a *polyhedron* (and a subset of $\mathcal{H}^2_+$ with the analogous property is called a *polygon*).

Thus Proposition 6.2 says that a Dirichlet domain is a polyhedron. Note that the proposition does not say that D_x has only *finitely many faces*, at least it only says this when D_x is *compact*. When D_x has finitely many faces (for some x) we say that G is *geometrically finite*.

Now consider any point y on the boundary of D_x, so y is on the boundary of H_g for one of the half-spaces defining D_x, in other words $d(y, x) = d(y, gx)$ for some $g \in G$. Then

$$d(g^{-1}y, g^{-1}x) = d(y, x) = d(y, gx) = d(g^{-1}y, x)$$

so $g^{-1}y$ also lies in the boundary of D_x. Thus each face of D_x is carried to another face of D_x by an appropriate element of G. We call these elements *side-pairing transformations*.

Example Consider the standard action of $PSL(2, \mathbb{Z})$ on the complex upper half-plane. Then for any point iv on the imaginary axis, with $v > 1$, the Dirichlet domain is the region $\{z \in \mathcal{H}^2_+ : |z| > 1, \ |Re(z)| < 1/2\}$ illustrated in Figure **??**, and the side-pairing transformations on this domain are $T : z \to z + 1$, $S : z \to -1/z$. (Proof: exercise.)

6.3 Poincaré's Polyhedron Theorem

We have seen that given a Kleinian group G, Dirichlet's construction allows us to find a fundamental domain on which G acts by side-pairing transformations. Poincaré's Polyhedron Theorem takes us in the opposite direction: given a convex polyhedron in $\mathcal{H}^3$ (or polygon in $\mathcal{H}^2$) and a set of side-pairing transformations for that polyhedron it gives us necessary and sufficient conditions for the group generated by those transformations to be discrete (i.e. Kleinian) and for the given polyhedron to be a fundamental domain for the group action. The precise conditions, though conceptually straightforward, are a little cumbersome to state, so we shall restrict ourselves to the two-dimensional case for most of the time. Our main concern (in the next subsection) will be to understand examples.

Let P be a polygon in $\mathcal{H}^2_+$. Note that the definition allows various possibilities. P may be compact (as on the left in Figure 10), it may have *ideal vertices* (vertices on the boundary of $\mathcal{H}^2_+$, as in the middle in Figure 10), or may have infinite area (as on the right in Figure 10).

Definition A *side-pairing* transformation of P is an isometry g_s of $\mathcal{H}^2_+$, sending one side s of P bijectively to another, s', and such that $g_s(P) \cap P = s'$.

Notation

For x_j a vertex of P which lies inside $\mathcal{H}^2_+$ (and so the two edges of P meeting at x_j meet at a non-zero angle), we let N_j denote an ϵ-neighbourhood (in the hyperbolic metric) of x_j intersected with P.

For y_j be an ideal vertex of P (so the two edges of P 'meeting' at y_j have angle zero between them), we let N'_j denote an ϵ-neighbourhood (in the Euclidean metric) of y_j intersected with P.

Theorem 6.4 (Poincaré's Polygon Theorem) *Let P be a polygon in $\mathcal{H}^2_+$, equipped with a set of side-pairing transformations g_s, one for each side of P and with $g_{s'} = g_s^{-1}$ if g_s pairs s with s'. If there exists a real $\epsilon > 0$ such that:*

• *for each vertex $x_0 \in \mathcal{H}^2_+$ of P there are vertices $x_1, ..., x_n$ of P (not necessarily all different) and isometries $f_0 = I, f_1, ...f_n, f_{n+1} = I$ such that*

(i) each $f_{j+1} = f_j g_s$ for some s, and

(ii) $f_j(N_j)$ are non-overlapping and have union the disc centre x_0 radius ϵ

and

• *for each ideal vertex y_0 of P there are ideal vertices $y_1, ..., y_n$ of P and isometries $f_0 = I, f_1, ...f_{n+1}$ with f_{n+1} fixing y_0 and parabolic, and such that*

(i)$'$ each $f_{j+1} = f_j g_s$ for some s, and

(ii)$'$ the $f_j(N'_j)$ are contiguous and non-overlapping

then

the group G generated by the side-pairing transformations g_s is discrete, P is a fundamental domain for the action of G on $\mathcal{H}^2_+$, and all relations in G are consequences of cycles $f_{n+1} = I$ corresponding to vertices of P in $\mathcal{H}^2_+$.

For a proof of this theorem see Beardon's book on discrete groups, or Ratcliffe or Maskit.

Comments

1. The condition on ideal vertices does not introduce any new relations, but it does ensure that P/G is *complete* (or equivalently that the translates of P cover the whole of $\mathcal{H}^2_+$).

2. The version for $\mathcal{H}^3_+$ (Poincaré's Polyhedron Theorem) is analogous. Now the 'sides' that are paired by the g_s are two-dimensional *faces* and instead of conditions (i) and (ii) we ask that neighbourhoods of *edges* of P fit together neatly (neighbourhoods of vertices then automatically fit together properly). Around edges which end at ideal vertices we ask that there be parabolic cycles (as in (i)$'$,(ii)$'$ above). Each edge which has ends inside $\mathcal{H}^3$ gives rise to a relation between the g_s and all relations in G are consequences of these.

6.4 Examples of Fuchsian and Kleinian groups

Examples in $PSL(2, \mathbb{R})$ (Fuchsian groups)

1. $PSL(2, \mathbb{Z})$ *(the modular group)*

Take our standard fundamental domain with side-pairings given by $S : z \to -1/z$ and $T : z \to z + 1$. Around $x_1 = (1 + i\sqrt{3})/2$ the picture is just that around $x_0 = (-1 + i\sqrt{3})/2$, conjugated by T. The vertex $y_0 = \infty$ is ideal, and T is parabolic ($z \to z + 1$). Poincaré's Polygon Theorem tells us that

$$PSL(2, \mathbb{Z}) = < S, T : S^2 = I, (ST)^3 = I >$$

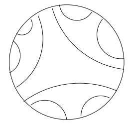

Figure 11: A truncated triangle and its first three reflections (Poincaré disc model)

2. *Surface groups*

Let P be a regular octagon with vertex angles all $\pi/4$. (To find such an octagon in the Poincaré disc model, just take a small regular octagon centred at the origin and blow it up steadily in size until the angles are $\pi/4$: this case must occur, by continuity, since in the limiting case when all vertices are ideal the angles are 0). Mark a pairing of the sides of P by labelling pairs of (oriented) sides such a way that one circuit anticlockwise around the boundary reads $ABA^{-1}B^{-1}CDC^{-1}D^{-1}$. Now think of A as an isometry carrying the first side marked A to the second side marked A etc. Then P is a fundamental domain for the group

$$G = < A, B, C, D : [A, B][C, D] = I >$$

(where $[A, B][C, D] = ABA^{-1}B^{-1}CDC^{-1}D^{-1}$). Note that $\mathcal{H}^2_+/G$ is a surface of genus two. (Higher genus surfaces may be obtained similarly.)

Comment. The octagon need not be regular: all that is really needed is that the angles add up to 2π and that the sides paired be of the same length. This is the beginning of the Teichmüller theory of hyperbolic stuctures on surfaces.

3. *Triangle groups*

Consider a triangle in $\mathcal{H}^2_+$ with angles π/p, π/q, π/r, where p, q, r are positive integers such that $1/p+1/q+1/r < 1$. We can always draw such a triangle in $\mathcal{H}^2_+$ by taking a small Euclidean triangle at the origin in the Poincaré disc model and gradually enlarging it until the angles are those desired. The (hyperbolic) area of such a triangle is π minus the angle sum. Now let G be the group generated by reflections in the sides of the triangles, and let G_0 be its orientation-preserving subgroup (products of even numbers of reflections). G_0 has generators $g_1 = R_2R_3$ and $g_2 = R_3R_1$. By Poincaré's Theorem G_0 is discrete, a quadrilateral made up of the initial triangle and one of its reflections is a fundamental domain for G_0, and a presentation for G_0 is

$$G_0 = < g_1, g_2 : g_1^p = g_2^q = (g_1g_2)^r = I >$$

(Note that if $1/p + 1/q + 1/r > 1$ we can construct a *spherical triangle* and the group G_0 is then *finite*.)

4. *Limit sets of triangle and truncated triangle groups*

When the fundamental polygon for G is compact, the limit set of G is the entire boundary circle S^1 of the Poincaré disc (the translates of P get smaller and smaller in the Euclidean metric as we move towards the boundary circle, so the orbit of any point inside the disc accumulates everywhere on S^1).

When the fundamental domain has ideal vertices the limit set remains the entire circle, but we can go further and take for example a 'truncated triangle' for our polygon P (see Figure 11). As before let G be the group generated by reflections R_1, R_2, R_3, and G_0 be the orientation-preserving subgroup (generated by R_2R_3, R_3R_1). Now R_2R_3 is hyperbolic and the 'gap' between its fixed points is in $\Omega(G_0) \subset S^1$. hence $\Lambda(G_0) \neq S^1$, so $\Lambda(G)$ has empty interior in S^1. Hence $\Lambda(G)$ is totally disconnected, but $\Lambda(G)$ is infinite, perfect, closed and bounded, so $\Lambda(G)$ is a Cantor set. Note that G_0 is freely generated by R_2R_3 and R_3R_1: there are no vertices so no relations.

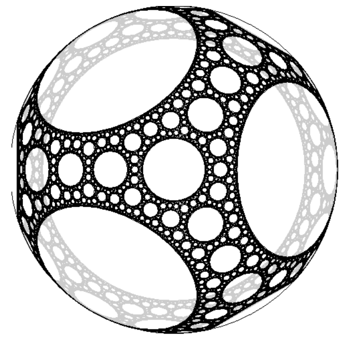

Figure 12: The limit set of a truncated tetrahedron group (picture by McMullen)

Examples in $PSL(2,\mathbb{C})$ (Kleinian groups)

1. *Tetrahedron groups*

Our 'polygon' now becomes a tetrahedron in $\mathcal{H}^3_+$ rather than a triangle in $\mathcal{H}^2_+$, and we consider the group G generated by reflections in its faces, and the orientation preserving subgroup G_0.

A tetrahedron in $\mathcal{H}^3_+$ is determined by its six *dihedral angles* (the angles between adjacent faces). To satisfy the conditions of Poincaré's Theorem we require them all to be of the form π/n with n integer.

A vertex inside $\mathcal{H}^3_+$ must have $1/p_1 + 1/p_2 + 1/p_3 > 1$, an ideal vertex must have $1/p_1 + 1/p_2 + 1/p_3 = 1$, a truncated vertex must have $1/p_1 + 1/p_2 + 1/p_3 < 1$ and where there is a truncated vertex the tetrahedron must meet the boundary of $\mathcal{H}^3_+$ in a $\pi/p_1, \pi/p_2, \pi/p_3$ triangle.

One can show that all combinations of dihedral angles are actually realised by tetrahedra or truncated tetrahedra. If all the vertices are internal or ideal then $\Lambda(G) = \hat{\mathbb{C}}$. If one or more vertices is truncated then $\Lambda(G)$ is a circle-packing (we get a circle as limit set for the triangle group around the truncated vertex, and then other elements of G move this circle around). See Figure 12 for a picture on the Riemann sphere and see Bullett and Mantica (Nonlinearity 1992) for more pictures and explanations.

2. *'Strings of beads'*

Here $C_1, ..., C_n$ are circles in $\hat{\mathbb{C}}$, each of the same size, touching the circle on each side and orthogonal to the unit circle S^1. Let R_m denote inversion in C_m, and extend R_m to a reflection in the hemisphere H_m spanning C_m in $\mathcal{H}^3_+$. Now, by Poincaré's Theorem, the part of $\mathcal{H}^3_+$ remaining after 'scooping out' all the hemispheres is a fundamental domain for the action of $G =< R_1, ..., R_n >$ and the only relations are $R_m^2 = I >$.

Note that the limit set here is S^1, but that if we pull the circles C_m apart the limit set becomes a Cantor set, and that if we perturb the sizes and positions of the circles C_m, but keeping them touching adjacent circles, the limit set becomes a *quasicircle* (a fractal homeomorphic to a circle). Going up in dimension an anologous construction can be used to obtain a group having limit set a wildly embedded circle in S^3.

3. *Schottky groups*

Take $g \geq 1$ pairs of mutually disjoint circles $C_1, C_1', \ldots C_g, C_g'$ in $\mathbb{C}$ with mutually disjoint interiors. For each j choose any Mobius transformation A_j that maps C_j to C_j' and the interior of C_j to the exterior of C_j'. The group G generated by $\{A_j\}_{1 \leq j \leq g}$ is called a *Schottky group of genus g*. Writing D_j and D_j' for the interiors of C_j and C_j' in $\hat{\mathbb{C}}$ bounded by C_j, it is easy to see that $\hat{\mathbb{C}} - (\bigcup_j D_j \cup \bigcup_j D'j)$ is a fundamental domain for G (so in particular G is discrete) that $\Lambda(G)$ is a Cantor set, that G is a free group on the generators $\{A_j\}$ and that $\Omega(G)/G$ is a surface S_g of genus g. The quotient $\mathcal{H}^3_+/G = M_g$ is a *handlebody*, a 3-manifold M_g constructed by adding g handles to a sphere. The boundary of M_g is S_g. *(Observe that in this example the fundamental domain on $\hat{\mathbb{C}}$ is not a Dirichlet domain, indeed $PSL(2,\mathbb{C})$ does not preserve any metric on $\hat{\mathbb{C}}$.)*

7 Quadratic maps and the Mandelbrot Set

7.1 The Mandelbrot set and its connectivity

Proposition 7.1 Every quadratic map $f(z) = \alpha z^2 + \beta z + \gamma$ with $\alpha \neq 0$ is conjugate to $q_c(z) = z^2 + c$ for a unique c.

Proof The conjugacy h must send ∞ to itself, and hence have the form $h(z) = kz + l$.

$$hf(z) = k(\alpha z^2 + \beta z + \gamma) + l \qquad q_c h(z) = (kz + l)^2 + c$$

These are equal (for all z) if and only if $k\alpha = k^2, k\beta = 2kl$ and $k\gamma + l = l^2 + c$. Thus we must have $k = \alpha, l = \beta/2$ and $c = \alpha\gamma + \beta/2 - \beta^2/4$. QED

Another useful parametrisation of the quadratic maps is given by the *logistic family*

$$p_\lambda(z) = \lambda z(1 - z)$$

Clearly p_λ is conjugate to q_c if and only if $c = \lambda/2 - \lambda^2/4$ (by Proposition 7.1).

The q_c parametrisation is more convenient when we are dealing with critical points, and the p_λ parametrisation is more convenient when we are dealing with fixed points and their multipliers. Note that q_c has critical points $0, \infty$, the latter a superattracting fixed point, and p_λ has fixed points 0 and $1 - 1/\lambda$, with multipliers λ and $2 - \lambda$ respectively.

Definition The *Mandelbrot set* is the subset of parameter space defined by

$$M = \{c : J(q_c) \text{ connected}\} \subset \mathbb{C}$$

Theorem 7.2 *M is the set of values of the parameter c such that the orbit $q_c^n(0)$ of the critical point 0 does not tend to the point ∞*

Proof If the orbit of 0 does not tend to ∞ then there is no critical value other than ∞ in the basin of attraction, $B(\infty)$, of ∞, and so there is no obstruction to extending the Böttcher coordinate (Section 5 of these notes) from a neighbourhood of ∞ to the whole of this basin. Hence B_∞ is homeomorphic to the open unit disc and its complement $\hat{\mathbb{C}} \setminus B_\infty$ is therefore connected, as is their common boundary ∂B_∞. But ∂B_∞ is closed and completely invariant, and cannot contain any points of the Fatou set (since any point in ∂B_∞ has bounded orbits, yet arbitrarily close to it are points with orbits going to ∞). So ∂B_∞ is the Julia set $J(q_c)$.

Conversely, if the orbit of 0 does go to ∞ then $J(q_c)$ is totally disconnected (a Cantor set) by the argument sketched earlier for the example $|c|$ large. QED

Definition The *filled Julia set* of q_c is $K(q_c) = \{z : q_c^n(z) \not\to \infty\}$

Note that $\partial K(q_c) = J(q_c)$, and that if $c \notin M$ then $K(q_c) = J(q_c) =$ Cantor set.

Theorem 7.3 (Douady and Hubbard 1982) *The Mandelbrot set M is connected*

Proof In fact Douady and Hubbard proved a much stronger result, that there is a conformal bijection between the complement $\hat{\mathbb{C}} - M$ of the Mandelbrot set and the complement $\hat{\mathbb{C}} - \mathbb{D}$ of the open unit disc. It is an immediate consequence of this that M is connected.

When $c \in M$, the Böttcher coordinate defines a conformal bijection

$$\phi_c : \hat{\mathbb{C}} - K(q_c) \to \hat{\mathbb{C}} - \mathbb{D}$$

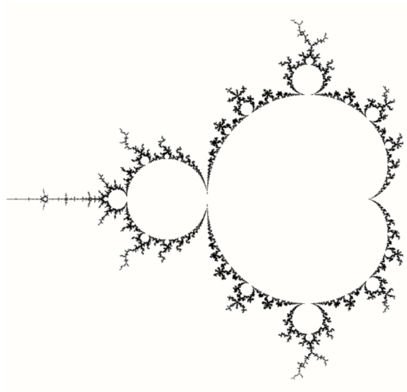

Figure 13: The Mandelbrot set

$$\phi_c(z_0) = z_0(1+\frac{c}{z_0^2})^{1/2}(1+\frac{c}{z_1^2})^{1/4}(1+\frac{c}{z_2^2})^{1/8}...$$

(conjugating q_c to $z \to z^2$). When $c \notin M$ the map ϕ_c, though not defined on the whole of the complement of K_c, is nevertheless defined on a neighbourhood of ∞ and as far as the critical value c of q_c. Define

$$\Psi : \hat{\mathbb{C}} - M \to \hat{\mathbb{C}} - \mathbb{D}$$

$$\Psi(c) = \phi_c(c)$$

This is a conformal bijection (see Douady and Hubbard, Comptes Rendues 1982, for more details). QED

Conjecture ('MLC') *M is locally connected*

A set X is called *locally connected* if every $x \in X$ has arbitrarily small connected open neighbourhoods. If M is locally connected then by a theorem of Carathéodory the map Ψ^{-1} extends to a continuous map from the boundary of $\hat{\mathbb{C}} - \mathbb{D}$ (a circle) onto the boundary ∂M of the Mandelbrot set. This would give us a purely combinatorial description of ∂M and many open questions concerning M would be resolved.

Definition A component of the interior of M is said to be *hyperbolic* if for every c in the component the map q_c has an attracting or superattracting periodic orbit.

Conjecture ('Hyperbolicity is dense') *Every component of the interior of M is hyperbolic*

Douady and Hubbard showed in their 1985 Orsay lecture notes that 'MLC' implies 'Hyperbolicity is dense'.

Both conjectures seem to be very difficult to resolve. Over the past 3 decades there has been a great deal of work on them. The set of points of ∂M at which local connectivity is known to hold has been steadily increased: Yoccoz proved it for 'all but infinitely renormalizable points' and Lyubich extended this to certain of these. Most experts seem to believe that MLC should be true, but it is known that the analogous set for cubics in place of quadratics is *not* locally connected (Lavaurs, Milnor), and that there exist quadratic maps q_c having non-locally-connected Julia sets. As far as 'Hyperbolicity is dense' is concerned, this has been proved for components of M meeting the real axis (Lyubich, McMullen, Swiatek: see McMullen's 1994 book 'Complex Dynamics and Renormalization') but the general question is still unresolved. Shishikura's proved in 1994 that the boundary ∂M of the Mandelbrot set has Hausdorff dimension 2.

7.2 The geography of the Mandelbrot set

We examine some of the more prominent features of M (Figure 13).

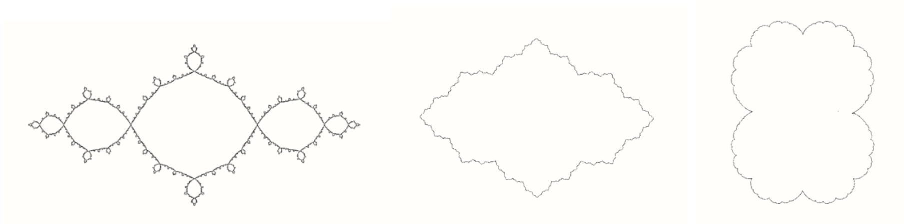

Figure 14: Julia sets for c=-1, c=-0.5 and c=+0.25

Let

$M_0 = \{c : q_c \textit{ has an attracting (or superattracting) fixed point}\}$

$= \{c : J(q_c) \textit{ is a (topological) circle}\}$

Lemma 7.4 $M_0 = \{c : c = \lambda/2 - \lambda^2/4 \textit{ for some } \lambda \textit{ with } |\lambda| < 1\}$

Proof Consider the logistic map p_λ. The multipliers of its fixed points are $\lambda, 2 - \lambda$. Hence

$M_0 = \{c : c = \lambda/2 - \lambda^2/4 \textit{ for some } \lambda \textit{ with } |\lambda| < 1 \textit{ or } |2 - \lambda| < 1\}$

But $\lambda/2 - \lambda^2/4 = (2 - \lambda)/2 - (2 - \lambda)^2/4$. QED

Thus M_0 is a *cardioid* (with a boundary that is smooth except at the cusp $c = 1/4$). Note that there is a bijection between points of M_0 and values of λ such that $|\lambda| < 1$. Thus M_0 is parametrised by the multiplier of the fixed point of q_c. The maps q_c with $c \in M_0 \setminus \{0\}$ are topologically conjugate to one another, indeed they are *quasiconformally* conjugate to one another (as we shall see in Section 8). Note that they cannot be *conformally* conjugate to one another as they have different multipliers at the fixed point. Note also that none of them can be topologically conjugate to $q_0 : z \to z^2$, since for q_0 the critical point 0 is a fixed point, and this is not true for any q_c with $c \neq 0$.

7.2 The intersection of M with the real axis

We consider how the behaviour of q_c varies as we vary the parameter c along the real axis. See Figure 14.

For $c > 1/4$, $J(q_c)$ is a Cantor set (it is an easy exercise to show that the orbit of 0 under q_c tends to ∞).

At $c = 1/4$, there is a neutral fixed point $z = 1/2$, with multiplier 1.

For $-3/4 < c < 1/4$, q_c has an attracting fixed point and $J(q_c)$ is a (topological, indeed *quasi-conformal*) circle, with dynamics conjugate to that of the shift. In particular $J(q_c)$ contains a dense set of repelling periodic orbits.

At $c = -3/4$, both points on the repelling period 2 orbit collide with the attracting fixed point, at a neutral fixed point (which has multiplier -1):

For $-5/4 < c < -3/4$, q_c has an attracting period 2 orbit, and the topology of $J(q_c)$ is the same as that for the (superattractive) case $c = -1$.

We digress briefly to justify the bounds $-5/4 < c < -3/4$:

Lemma 7.5 q_c *has an attracting period* 2 *orbit if and only if* $|1 + c| < 1/4$

Proof The points of period 1 or 2 are the solutions of $q_c^2(z) = z$. Expanding $q_c^2(z) - z$ we have

$$q_c(q_c(z)) - z = ((z^2 + c)^2 + c - z = (z^2 - z + c)(z^2 + z + 1 + c) = (z - \alpha)(z - \beta)(z - u)(z - v)$$

where α, β are the fixed points and u, v is the period 2 cycle. The multiplier of the period 2 cycle is $q_c'(u)q_c'(v) = 4uv = 4(1 + c)$. The period 2 cycle is attracting if and only if this has modulus less than 1. QED.

Returning to our journey in parameter space along the real axis:

For $-2 < c < -5/4$, as c decreases through this range, we have a sequence of period doublings until we reach the Feigenbaum point (the 'period-doubling limit point'. This is followed by the whole Milnor/Thurston sequence of periods for real unimodal maps, familiar to dynamicists (in particular this contains all the natural numbers in the Sarkovskii order). The most prominent component of $int(M)$ along the axis after the period-doubling limit is one corresponding to a period three attracting orbit, and we finish at $c = -2$ where the Julia set is the real interval $[-2, +2]$ (and q_c is semi-conjugate to $z \to z^2$: see the exercise early on in these notes).

For $c < -2$, it is again easily proved that the orbit of the critical point 0 tends to ∞ and hence that the Julia set is again a Cantor set.

The behaviour for c at different points along the real axis is is no surprise to real dynamicists since the quadratic family is conjugate to the logistic family. However with c complex we can now leave the main cardioid M_0 at other points than just $c = -3/4$. When c is on the boundary of M_0 at the point where $\lambda = e^{2\pi ip/q}$, q_c has a neutral periodic point with this as multiplier, and when c passes into the adjoining component q_c has an attracting period q orbit. There are then further bifurcations as we pass along a path through different components of $int(M)$. In the next subsection we shall find that studying the combinatorics of 'external rays' can give us a great deal of information about the overall structure of M

Exercise Compute the values of c where q_c has a superattractive period three orbit (that is, where the point 0 has period three).

7.3 Internal and external rays: the 'devil's staircase'

When $c \in M$, for any $\theta \in [0, 1)$, the radial line $arg(z) = 2\pi\theta$ on $\hat{\mathbb{C}} - \mathbb{D}$ (where $\mathbb{D}$ is the unit disc) maps under the inverse ϕ_c^{-1} of the Böttcher map to the *external ray* $\mathcal{R}_\theta$ of argument $2\pi\theta$ on $\hat{\mathbb{C}} - K(q_c)$.

Similarly, in the parameter plane, the radial line $arg(z) = 2\pi\theta$ on $\hat{\mathbb{C}} - \mathbb{D}$ maps under the inverse Ψ^{-1} of the Douady-Hubbard map to the *external ray* $\mathcal{R}_\theta$ of argument $2\pi\theta$ on $\hat{\mathbb{C}} - M$.

A ray is said to *land*, if it accumulates at a unique point of $J(q_c)$ (in the dynamical case) or ∂M (in the parameter case). If $J(q_c)$ (or ∂M respectively) is locally connected then all external rays land (by Carathéodory's criterion). Unfortunately there are examples where $J(q_c)$ is known not to be locally connected, and where certain external rays do not land; moreover the conjecture 'MLC' is still unproved so we cannot be sure that all external rays in the parameter space land.

An outline proof of the following theorem can be found in Carleson and Gamelin, and a full proof can be found in the Douady-Hubbard Orsay notes.

Theorem 7.6 (Douady and Hubbard) *Every parameter space external ray with rational angle θ lands at a point c of ∂M. If θ is a rational with odd denominator then q_c has a parabolic cycle. If θ is a rational with even denominator then the critical point 0 of q_c is strictly preperiodic.*

We first consider the rational external rays which land on the boundary of the main cardioid, M_0. Recall that M_0 is itself parametrised by the unit disc and we can therfore define *internal rays* inside M_0. The *internal ray* of argument ν is the set of values of $c \in M_0$ for which the multiplier of the fixed point of q_c has argument $2\pi\nu$. Consider the end point on ∂M_0 of the internal ray of argument $\nu = 1/3$. This is the value of c for which the fixed point α of q_c has multiplier $e^{2\pi i/3}$ (this c lies at the top of the cardioid: it is where the first period-tripling occurs). The external rays $1/7, 2/7, 4/7$ in the dynamical plane landing at α are as shown in Figure 15.

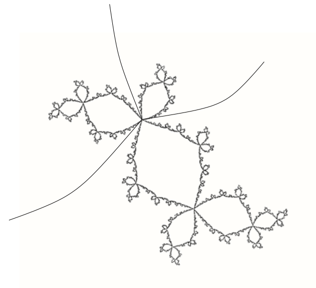

Figure 15: External rays landing at the α-fixed point of Douady's rabbit

Note that we can pick out two particular rays, which together enclose the component of $int(K(q_c))$ containing the critical value. These we have labelled $\theta_-(1/3)$ and $\theta_+(1/3)$. It can be shown that in the parameter space the corresponding external rays with the same arguments, $\theta_-(1/3)$ and $\theta_+(1/3)$, land at c (an example what Douady called 'ploughing in the dynamical plane but harvesting in the parameter plane').

More generally, for c at the end of each internal ray in M_0 of rational argument p/q, the map q_c has a (neutral) fixed point α of rotation number p/q and we can pick out the pair of external rays enclosing the component of $int(K(q_c))$ containing the critical value c. How do we compute the values of $\theta_-(p/q)$ and $\theta_+(p/q)$? Since α is a fixed point of rotation number p/q there are necessarily q external rays landing at α and the effect of q_c on these rays is to permute them in cyclic order. But the action of q_c on arguments of rays is simply that of $t \to 2t \pmod{\mathbb{Z}}$, so our search for candidates for $\theta_\pm(p/q)$ is reduced to a search for finite orbits of $t \to 2t$ on the unit circle $\mathbb{R}/\mathbb{Z}$, arranged in the same order around the circle as an orbit of a rigid rotation through $2\pi p/q$. This is a purely combinatorial question and was answered (though in a slightly different context) by Morse and Hedlund in their pioneering work on symbolic dynamics in the 1930's:

Theorem 7.7 *For each rational p/q there is a unique finite forward invariant orbit $A_{p/q}$ of $t \to 2t$ of rotation number p/q on the circle $\mathbb{R}/\mathbb{Z}$.*

(For a proof of this and other results concerning order-preserving orbits of the shift, see Bullett and Sentenac, Math. Proc. Cam. Phil. Soc. 1994.)

But supposing we have found this orbit $A_{p/q}$, how are we to know which of its points are the special points $\theta_\pm(p/q)$? This turns out to be very straightforward.

Lemma 7.8 *Any ordered orbit of $t \to 2t$ on the circle $\mathbb{R}/\mathbb{Z}$ is contained in a semicircle*

Proof Since $t \to 2t$ doubles distance, any three points on the circle have images in the same order around the circle if and only the three original points lie in a common semi-circle. QED

As a consequence it makes sense to refer to the *least* and *greatest* points of the orbit $A_{p/q}$. We identify the points $\theta_\pm(p/q)$ by observing that the dynamical picure requires that the least point of $A_{p/q}$ be $(\theta_+(p/q))/2$ and the greatest be $(\theta_-(p/q))/2 + 1/2$ (see the picture above for the case $p/q = 1/3$: the inverse image of the component of $int(K(q_c))$ containing the critical value c is that containing the critical point 0).

Algorithm for $\theta_\pm(p/q)$

There is a simple algorithm constructing the binary sequence of each of $\theta_+(p/q)$ and $\theta_-(p/q)$:

Draw a line of slope p/q, through the origin in $\mathbb{R}^2$. To construct $\theta_-(p/q)$, take the integer 'staircase' lying just below this line, but not touching it, and starting at the point $(1,0)$ write 1 for each horizontal step which is followed by a vertical step, 0 for a horizontal step followed by another horizontal one. To construct $\theta_+(p/q)$ do the same with the staircase touching the line.

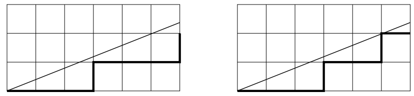

Figure 16: $\theta_-(2/5) = .\overline{01001} = 9/31$ $\theta_+(2/5) = .\overline{01010} = 10/31$

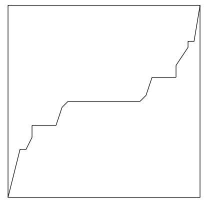

Figure 17: The devil's staircase assigning internal angle to external angles around the cardioid M_0.

Example $p/q = 2/5$: See Figure 16.

Every point on ∂M_0 at the end of an internal ray of *irrational* argument ν corresponds to $\theta_\nu = lim_{p/q \to \nu} \theta_\pm(p/q)$. The assigment of internal angles to external angles as we make a circuit of the boundary of the cardioid M_0 has graph a 'devil's staircase' (see Figure 17). We draw the graph this way round (rather than that assigning external angles to internal angles) in order to have a *continuous* function. It is not difficult to prove that the horizontal steps in the graph above have total length 1. A 'devil's staircase' is the graph of a continuous function that is constant on a set of full measure without being globally constant. This particular devil's staircase another interesting property: it is a theorem due to Douady that every irrational ν corresponds to a *transcendental* θ_ν (see Bullett and Sentenac, Theorem 4).

7.4 External rays landing at points outside the main cardioid

Rational external rays can be used to give us an overall picture of the geography of M. The next step is to consider those landing on the boundary of a component of $int(M)$ immediately adjacent to M_0, say that corresponding to rotation number p/q. This component (which we shall label $M_{p/q}$) has the property that corresponding maps q_c each have an attractive period q orbit. We can parametrise $M_{p/q}$ by the multiplier of this orbit and hence define internal rays inside $M_{p/q}$ in just the same way as we did for M_0. The r/s internal ray in $M_{p/q}$ is the landing point of external rays $\theta_\pm(p/q, r/s)$ obtained from $\theta_\pm(r/s)$ by replacing the digit 0 by the repeating block (of length q) from $\theta_-(p/q)$ and the digit 1 by the repeating block from $\theta_+(p/q)$.

Example

$$\theta_-(1/3, 1/2) = .\overline{001010} \qquad \theta_+(1/3, 1/2) = .\overline{010001}$$

By repeating the same process (which is known as 'tuning') we can compute the arguments of external rays landing on the boundary of any component which is accessible from M_0 by a finite number of boundary crossings. But there are of course components of $int(M)$ which are much further away than this from M_0: for example all components beyond the Feigenbaum point on the real axis are an infinite number of boundary crossings away from M_0. Methods of assigning 'internal addresses' to all hyperbolic components, and algorithms relating these addresses to 'kneading sequences' associated to external rays landing on the components, were developed by Penrose (1990) and by Lau and Schleicher (1994).

7.5 The combinatorial Mandelbrot set

We sketch an algorithm due to Lavaurs (Comptes Rendues 1986) which, if ∂M is locally connected, gives M as the quotient of the unit disc by an equivalence relation defined via a *lamination*.

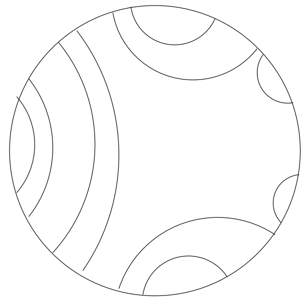

Figure 18: Lavaurs' algorithm

Lavaurs' Algorithm

Write every rational which has odd denominator in the form $p/(2^k - 1)$ with k as small as possible.

1. Connect $1/3$ to $2/3$ (on $\partial\mathbb{D}$) by an arc in $\mathbb{D}$.

2. Assuming all rationals of form $p/(2^{k-1} - 1)$ have been connected in pairs, connect pairs of form $p/(2^k - 1)$, starting with the smallest number not yet connected, and connecting it to the next smallest one possible without crossing arcs already constructed (see Figure 18 for the construction up to and including $k = 4$). The (combinatorial) Mandelbrot set is now obtained by shrinking each of the arcs to a point.