

Extending holomorphic motions

by

DENNIS P. SULLIVAN and WILLIAM P. THURSTON

*I.H.E.S., Bures-sur-Yvette
France*

*Institute for Advanced Study
Princeton, NJ, U.S.A.*

1. Introduction

The notion of an isotopy of one set within another set is one of the key concepts of topology. Here is one way this concept can be generalized to a holomorphic context:

Definition. If X is a subset of \mathbf{C} , a *holomorphic motion* of X in \mathbf{C} is a map

$$f: T \times X \rightarrow \mathbf{C}$$

defined for some connected open subset $T \subset \mathbf{C}$ containing 0 such that

- (a) for any fixed $x \in X$, $f_t(x) = f(t, x)$ is a holomorphic mapping of T to \mathbf{C} ,
- (b) for any fixed $t \in T$, f_t is an embedding, and
- (c) f_0 is the identity map of X .

We think of t as a kind of complex time parameter. Note that in the definition, there is no requirement of holomorphy in the X -direction. X should be thought of with just its topological structure or its quasiconformal structure, although even continuity doesn't directly enter into the definition; the only restriction is in the t direction. We will see that continuity is a consequence of the hypotheses, by the lambda lemma of Mañé, Sad and Sullivan ([2], Theorem 2).

This definition is applicable in a number of interesting situations. For instance, the limit sets of Kleinian groups often move holomorphically as parameters are varied. Similarly, the Julia sets for iterated rational maps often move holomorphically with the parameters.

In topology, it is important to know whether an isotopy of one space within another can be extended to an ambient isotopy, that is, to an isotopy of the big space which restricts to the given isotopy of the small space. Without additional conditions, an

extension is not generally possible. For example, consider a loosely knotted piece of string within \mathbf{R}^3 . Physically, one can take hold of the ends and pull it tight. Mathematically, this corresponds to an isotopy of the knot in \mathbf{R}^3 which shrinks the knotted part to an arbitrarily small size—and finally, makes it disappear. Such an isotopy certainly cannot be extended to all of \mathbf{R}^3 , since the fundamental group of the complement of the circle is different at the beginning of the isotopy and the end.

It is not hard to arrange the isotopy so that the motion of any point is differentiable in time. The isotopy can even be chosen so that it is real-analytic in the time parameter, if the knot vanishes and then instantly reappears.

There are also examples in two dimensions: there are isotopies of closed sets in the plane which cannot be extended to ambient isotopies. The simplest example is for X a countable set with one accumulation point. If the set is represented as say $\{x_i = 1/i \mid i \in \mathbf{Z}_+\}$ together with $x_\infty = 0$, the isotopy can be choreographed by making each point x_i swing around its two neighbors, two turns about x_{i+1} for each turn about x_{i-1} , faster and faster as $i \rightarrow \infty$. Any extension of the isotopy restricted to the first k points sends the arc $\overline{x_{k-1}x_k}$ at the beginning of the isotopy to an arc which winds around x_i by the end of the isotopy. To do this for all i would violate continuity at 0.

The question of extending an isotopy to an ambient isotopy also has significance in the holomorphic context. Our main result addresses this problem, as follows:

THEOREM 1. *There is some universal constant $a > 0$ such that any holomorphic motion of any set $X \subset \mathbf{C}$ parametrized by the disk $T = D_1$ of radius 1 about 0 can be extended to a holomorphic motion of \mathbf{C} with time parameter the disk D_a of radius a about 0.*

An important correlate of an open set moving holomorphically in \mathbf{C} is quasiconformality:

PROPOSITION. *For any holomorphic motion f of \mathbf{C} , each map f_t is quasiconformal.*

Proof. The cross-ratio of any quadruple of points in \mathbf{C} during the motion is a holomorphic function which omits the values 0, 1, and ∞ . The derivative of a holomorphic function sends a tangent vector in its domain to a vector in its range which cannot be longer, when lengths are measured by the Poincaré metrics for the domain and the range. Therefore, for any particular value of t and for all quadruples of points whose initial cross-ratios are in some compact set K of cross-ratios, the cross-ratios of f_t of the quadruples are uniformly bounded. It readily follows that f_t is quasiconformal.

By the theory of quasiconformal maps, the converse is also true: any quasiconformal orientation-preserving map of \mathbf{C} can be connected to the identity by a holomorphic motion. To find such a holomorphic motion, one simply takes the Beltrami coefficient for the quasiconformal map and multiplies it by a complex variable t .

COROLLARY. *If f is a holomorphic motion of a set X , then for each time t the map f_t extends to a quasiconformal map of \mathbf{C} .*

Proof. By the preceding two assertions, it follows that the corollary is true when t is sufficiently small. To obtain the result for a general t , walk from 0 to t in small steps and compose the quasiconformal maps for each step to obtain the big quasiconformal map.

In a companion paper, Bers and Royden [1] will give another proof which shows that a can be taken to be $1/3$. We do not know if the constant a of the theorem can be taken to be 1.

In section 6 we will state and prove another theorem which does work globally. This result is obtained from our main statement by changing all holomorphic motions to quasiconformal motions. The definition of a quasiconformal motion of an arbitrary set will be found in that section.

There is a simple but very useful special case of the main theorem which was proven by Mañé, Sad and Sullivan [2], which they dubbed the λ -lemma.

THEOREM 2 (λ -lemma). *A holomorphic motion of a set $X \subset \mathbf{C}$ can be extended to a holomorphic motion of the closure of X , with the same time parameter set T .*

Proof. The idea is that the motion of the points of the set must be equicontinuous because they remain disjoint.

This can be seen more explicitly by first choosing two finite points of X together with infinity, then normalizing by changing coordinates via the unique family of affine transformations sending these three points to $\{0, 1, \infty\}$. The Möbius transformations depend holomorphically on T , so the problem of extending the motion of the new family is the same as that for the old family. (We ignore here the exceptional case that there is only one point of X ; the problem in this case is trivial.)

The points in the new family—except those which are fixed—move in the complement of $\{0, 1, \infty\}$. The map from the parameter space T to the three-punctured sphere takes vectors of unit length in the Poincaré metric of T to vectors of length not exceeding one in the Poincaré metric of the three-punctured sphere.

The motion of points with respect to T is equicontinuous, so the closure of the set of all the maps $t \rightarrow f_t(x)$ is compact in the topology of uniform convergence on compact sets. Consider the closure, which consists of holomorphic functions from T to \mathbb{C} .

If g and h are any two distinct holomorphic functions on an open set T , their graphs intersect iff $g-h$ has a zero of some order at some point in T . This property is stable under perturbation. Therefore, the property of a pair of holomorphic functions that their graphs intersect is open. The opposite property, that the graphs of g and h are disjoint, is closed. The set of graphs of functions in the closure of our family are disjoint, and therefore there is exactly one through each element of X . This defines an extension of the holomorphic motion of X to a holomorphic motion of its closure.

It is easy to extend Theorem 1 to the case that the parameter space T has arbitrary dimension.

It is also easy to extend it to the case of the motion of a set X in an arbitrary Riemann surface, or even a family of Riemann surfaces depending on the parameter space T . For the latter case, one uses the Bers embedding of Teichmüller space to translate the problem into a motion of a set in \mathbb{C} .

A problem closely related to the problem of extending holomorphic motions is the holomorphic axiom of choice: Given a holomorphically moving set, is it possible to choose a holomorphically moving point in the complement of the set? If we strengthen this slightly so that the point for $t=0$ is chosen in advance, then a positive answer, even in the case that X consists of only a finite set of points implies the general extendibility of holomorphic motions to all of \mathbb{C} .

To see this, first observe that the ability to choose an additional point with arbitrary initial value in the complement of a holomorphically moving finite set implies the same ability in the complement of a general holomorphically moving set: if $F \subset X$ is a large finite set of points which comes very close to every point in X , then a choice which starts out not too close to X and remains in the complement of F must also avoid X for most of the time parameter (since the distance to X remains large in the Poincaré metric of the complement of F). Using equicontinuity, one can pass to the limit as F becomes dense in X and obtain a choice, through the given point, which misses X entirely. Thus the ability to choose a moving point in the complement of finite sets implies this ability for general sets. By adding more and more moving points, after a countable number of steps one has a dense set of moving points, so that the motion extends by continuity to a motion of the complex plane.

2. Framework for the proof

The strategy of the proof is to divide and conquer. We may assume that X is closed. First, we will show that for any point p in the complement of X the motion of X can be extended to a motion of some neighborhood of p for some subset of the time parameter space T . We will do this so that the subset T_0 is independent of p . We must do this carefully, so that it will be possible to piece together the choices from different neighborhoods, using a partition of unity for $\mathbb{C}-X$.

The local choices near some point p are fairly easy to make in the case that the set $X \cup \{\infty\}$ is connected, so we will take care of that case first, in section 3. In fact, the local choices are made by using a solution for the case that X consists of two points. This solution can be transplanted (for a limited subset of $T \times X$) if for each p , we choose two reasonably spaced points from the actual set X . A reasonable spacing is achieved if the distance between the two points and the distance of the two points from p is comparable to the distance of p from X . The hypothesis that $X \cup \{\infty\}$ is connected is used to guarantee the existence of two such points.

To do the general case, we will need to analyze the complement of X in terms of the thick-thin decomposition for its Poincaré metric. The thick part of the complement of X is handled like the connected case. The thin part divides into two new cases: cusps and short geodesics. The analysis of cusps again reduces to the case that X has only two points. We will construct a solution for the case of a thin part of X which contains a short geodesic by transplanting a solution for the case that X has exactly three points. This corresponds to the quadruply-punctured sphere, the simplest Riemann surface which can have short geodesics.

To see that local solutions can be pieced together, it helps to picture the trace or graph of the motion, in $T \times \mathbb{C}$. We begin with the motion of X , which gives a foliation in a subset of $T \times \mathbb{C}$ above T . For each p , we find some rule to fill in a foliation in some neighborhood of p which is disjoint from the leaves of X . We must do this so that each leaf projects to a uniform size neighborhood in T .

To pass between choices made for different neighborhoods which overlap (at time $t=0$), we average them (considered as maps of T to \mathbb{C}) using a partition of unity. To see what happens in any region of overlap, we make a t -dependent affine change of coordinates to fix two points of X at 0 and 1. These two points are picked so that they are near to the region in question, compared to their distance from each other and the distance of the region from X . There is no overlapping of different local choices in the region of influence of a short geodesic or a cusp—such a region is contained in a single

coordinate patch, so it is always possible to find two points as above. Note that the process of averaging is unaffected by affine changes of coordinates.

In this coordinate system, no matter what the overlapping choices of holomorphic motions, they can move at speed at most 1 in the Poincaré metric for the thrice-punctured sphere. Since the region at time $t=0$ has a distance from X (following the affine renormalization) which is bounded above zero uniformly for all problems under consideration, convex combinations of the moving choices remain disjoint from X for a uniformly bounded time.

When we understand how to make local extensions of a holomorphic motion, we must still make sure that after averaging by a partition of unity, the motions of different points remain disjoint from each other.

There are two effects to take care of. First, if the diffeomorphisms from one fiber to another are not near each other in the C^1 topology, then convex combinations of them might have derivative zero. In order to take care of this, we will make sure that (after the affine renormalization) the maps from one copy of \mathbb{C} to another all have derivatives near the identity. This precaution would suffice if we were taking convex combinations of the different choices with constant coefficients.

The second problem is that if the derivatives of the partitioning functions (with respect to the fiber variable) are too large, then convex combinations made with them will oscillate wildly. To take care of this problem, we will make sure that we define local motions on neighborhoods of uniform size (with respect to the coordinates as above) and that these neighborhoods have ample overlap, so that the fiber partitioning functions can be chosen with uniformly bounded derivatives.

To convince oneself that these two conditions are exactly what is needed without actually writing down the formula, one can think of the foliation picture: picture a small disk in T , and enlarge it until it has size 1. Then the various local choices automatically become uniformly near to horizontal, so that any effect due to averaging by a controlled partition of unity is absorbed.

3. Extending the motion of a connected set

Let X be a holomorphically moving set in \mathbb{C} such that $X \cup \{\infty\}$ is connected. Consider any point $p \in \mathbb{C}$, disjoint from X at time $t=0$. We will extend the motion of X to include a motion of a neighborhood of p .

Let r be the distance of p from X . Let x_0 be a point in X which has distance r from p . The circle of radius $2r$ around p intersects X , since X is connected and unbounded. Let x_1 be a point of intersection of this circle with X .

The points x_0 and x_1 are separated by at least r and at most $3r$. Normalize the motion of X by t -dependent affine transformations so that the points x_0 and x_1 remain fixed at 0 and 1. After this normalization, the distance of p from 0 at time $t=0$ is at least $1/3$, and the distance from the point 1 is at least $2/3$. Define a holomorphic motions of the neighborhood U of radius $1/6$ about p by fixing it, in the new coordinates.

At time 0, the distance of X from U as measured in the Poincaré metric of the thrice-punctured sphere $\mathbb{C}-\{0, 1\}$ is bounded from zero. Since the set X moves with speed at most 1 in the Poincaré metric of the thrice-punctured sphere, this prescription gives a holomorphic motion of the neighborhood for a uniform disk in the time parameter T .

These choices can be combined, using a partition of unity, to give a motion of \mathbb{C} defined for a disk in T of a size which is uniform with respect to different problems of this form. In fact, consider the motion of another neighborhood V defined in this way where V and U intersect at time 0. In the coordinates for which points in U are fixed, the motion of V is given by a t -dependent affine transformation. The motion of V is governed by two points x_3 and x_4 of X , which at time $t=0$ are uniformly bounded away from each other, from infinity, from U and from V .

The points x_3 and x_4 move at speed at most 1 in the Poincaré metric of the thrice-punctured sphere $\mathbb{C}-\{0, 1\}$. In any bounded subset of \mathbb{C} , the unit circles for this Poincaré metric have bounded diameter. Consequently, x_3 and x_4 are moving at a bounded speed also in the Euclidean metric. Therefore we can find a uniform disk in T such that the derivatives of the comparison between the U motion and the V motion are close to the identity.

Since neighborhoods of the form of U can be chosen to overlap uniformly, this completes the proof in the case that X is connected.

4. Solution for the quadruply-punctured sphere

The method above does not carry over to the case that X is not connected, even if it contains only three points.

For example, let C be any constant, and define a holomorphic motion of a set $x_0=0$, $x_1=1$ and $x_2=e^{-2C}$ by the formula

$$\begin{aligned} f_t(x_0) &= 0, \\ f_t(x_1) &= 1, \\ f_t(x_2) &= e^{C(t-2)}. \end{aligned}$$

We can restrict attention to the unit disk. Suppose that C is very large. The motion of x_2 wraps it many times around the origin, while keeping it quite close to $x_0=0$. If we use a local prescriptions for extending this motion based on choosing subsets of X consisting of two points, we see that the only reasonable choice for x very near 0 is based on x_0 and x_2 , while the only reasonable choice for x near, say, -7 is based on x_1 and either of the other two points. These choices clash fiercely with each other wherever they overlap. There is no reasonable way to average them and get a choice which works over the entire quadruply-punctured sphere.

Luckily, there is an alternate method which gives a global solution in the case of the quadruply-punctured sphere, provided the set T is contractible. When this is the case, there is a well-defined homotopy class of maps from the complement of X at time 0 to the complement of X at time t . Given this homotopy class, there is a well-defined element of the Teichmüller space for a quadruply-punctured sphere associated with each t , and the element of Teichmüller space depends holomorphically on t .

To extend the motion of X , we simply use the Teichmüller maps between the surfaces. The Teichmüller space for a quadruply-punctured sphere is the Poincaré disk, and the Teichmüller maps depend holomorphically on Teichmüller space for the quadruply-punctured sphere. This holomorphic dependence of the Teichmüller maps is a special feature of Teichmüller spaces of complex dimension 1. The holomorphic motion can be seen more explicitly by using the fact that the universal 2-fold-branching cover of the quadruply-punctured sphere is isomorphic to \mathbb{C} , with the isomorphism classically given by elliptic integrals. A holomorphically moving configuration of the three points together with ∞ transforms to a holomorphically moving lattice in \mathbb{C} , which has an extension to a motion of all of \mathbb{C} in the form of t -dependent real affine transformations.

5. The general case

Now we are prepared to consider a general holomorphically moving subset $X \subset \mathbb{C}$. We will use local solutions usually based on the thrice-punctured sphere, but we will switch to local solutions based on the quadruply-punctured sphere in neighborhoods where X appears to be split into two small but nontrivial pieces.

One way to express the choice between the two kinds of local solutions is based on the thick-thin decomposition of the Riemann surface, considered with its hyperbolic metric. In the thick part of the surface (where the injectivity radius is greater than some fixed constant $\varepsilon > 0$), we can use the thrice-punctured sphere. Each component of the

thin part of the surface is either parabolic (coming from a cusp) or hyperbolic. In the hyperbolic thin case, we use the quadruply-punctured sphere, while in the parabolic case we again use the thrice-punctured sphere.

For simplicity, instead of this hyperbolic description of choices for local solutions, we will use an approximately equivalent but more direct description in terms of the Euclidean geometry of \mathbb{C} .

Let p be any point in $\mathbb{C} - X$. Choose a point x_0 which is at least as close to p as any other point of X . Change coordinates so that p is at 1, x_0 is at 0, and normalize the motion (via a t -dependent translation of \mathbb{C}) to keep x_0 fixed at the origin. Choose a small number ε (fixed throughout this case).

Case (i). (The thick case.) Suppose that there is some other point q in X such that $\varepsilon \leq |q| \leq 1/\varepsilon$. Then choose x_1 to be an element of X which minimizes $|\log|x_1||$. Define a local motion around p to be the affine motion governed by x_0 and x_1 .

Case (ii). (The cusp case.) Suppose that (i) fails, and either there is no point in X with modulus greater than that of p , or there is no point except 0 in X with modulus less than p . By inverting if necessary, we can assume that there is no point of X inside the unit disk except for the origin. Let x_1 be a point of X with minimal modulus m . Define a motion in the disk of radius $2\varepsilon m$ about 0 to be the affine motion governed by x_0 and x_1 . We use only one such choice for this entire disk, governed by x_0 and some point of X of minimal modulus.

Case (iii). (The short geodesic case.) Suppose that (i) and (ii) fail, so that p is contained in an annulus bounded by circles of radii $0 < r < \varepsilon < 1/\varepsilon < R < \infty$ which touches X on its two boundary components, but is otherwise disjoint from X . Let x_1 be an element of X of modulus r , and let x_2 be an element of modulus R . Define a motion of the subannulus of inner and outer radii $s = r/2\varepsilon < 1 < 2R\varepsilon = S$ by restricting the solution for the quadruply-punctured sphere. Make only one such choice for this annulus.

These solutions have been chosen so that they are defined on sets which have ample overlaps. Overlapping can occur only

between cases (i) and (i),
between cases (i) and (ii), or
between cases (i) and (iii).

Overlapping between a pair of thin cases ((ii) or (iii)) is not possible.

The first kind of overlap is exactly like the kind we have already considered for X connected; only the constants are different.

The second kind of overlap is just like the first kind of overlap—it is an overlap of two thrice-punctured sphere coordinates in a region where both choices are reasonable.

The third kind of overlap is somewhat different. To understand it, we must first understand the behaviour of the solution for a quadruply-punctured sphere which has a short geodesic, in the vicinity of its thick part.

Any quadruply-punctured sphere with a short geodesic can be arranged so that its puncture are $x_0=0$, x_1 , x_2 , and $x_3=\infty$, where the ratio x_2/x_1 is very large. The short geodesic separates x_0 and x_1 from x_2 and x_3 . By symmetry, we can focus on the thick part of the quadruply-punctured sphere which is in the half containing x_0 and x_1 . To get a convenient coordinate system, normalize so that x_1 becomes 1.

LEMMA. In the coordinates above, the derivatives of the maps from fiber to fiber given by thin quadruply-punctured sphere solutions are uniformly equicontinuous in the region $P \subset \mathbb{C}$ obtained by removing ε -disks about 0, 1, and ∞ . (An ε -disk about ∞ means the complement of a disk of radius $1/\varepsilon$ about 0.)

Proof. As long as the shape of the quadruply-punctured sphere remains in a compact region of the modular space of the quadruply-punctured sphere, the derivatives automatically remain equicontinuous in P . Thus, the problem is to show that as the quadruply-punctured sphere gets thinner and thinner, these derivatives remain reasonable.

The universal 2-fold branching cover of any sphere over four points is \mathbb{C} . We will transfer the question to \mathbb{C} by taking these universal two-fold branching covers.

The preimage of the branch points forms a lattice in \mathbb{C} . We can normalize the lattices in \mathbb{C} so that the origin is a preimage of 0, and so that is a preimage of 1 nearest to the origin. Thus, the lattices all contain the integers.

The Poincaré metric of the plane minus the lattice maps isometrically to the Poincaré metric of the quadruply-punctured sphere, so as the length of the short geodesic goes to zero, the distance between the two thick parts of the quadruply-punctured sphere tends to infinity, and the distance from the row of lattice points which are integers the next row also tend to infinity. In the limit, the branched covering from \mathbb{C} to $\hat{\mathbb{C}}$ converges to a branched covering map of \mathbb{C} to \mathbb{C} , with 2-fold branching above $0=x_0$ and $1=x_1$.

The advantage of passing to the branched covers, \mathbb{C} , is that the Teichmüller maps between quadruply-punctured spheres lift to \mathbb{R} -linear maps of \mathbb{C} to itself. They are represented by matrices of the form

$$A = \begin{pmatrix} 1 & x \\ 0 & y \end{pmatrix}.$$

The distance between two of the lattices in Teichmüller space depends only on the matrix A which transforms one into the other, not on particular lattices. Therefore, a sequence of Teichmüller maps between pairs of quadruply-punctured spheres with bounded Teichmüller separation always has a geometrically convergent subsequence, as one passes to the limit, which is a thrice-punctured sphere. The limiting thrice-punctured sphere is identified with the thrice-punctured sphere defined by $0=x_0$, $1=x_1$, and $\infty=x_3$. The geometric limit map, lifted to the branched coverings C of the thrice-punctured sphere, is also a \mathbf{R} -linear map, described by a matrix A of exactly the same form.

Since the Teichmüller space for the quadruply-punctured sphere is equivalent to unit disk, and the Teichmüller metric is its Poincaré metric it follows that, in a holomorphically moving family, the Teichmüller distance between quadruply-punctured spheres defined by the motion of four points does not exceed the Poincaré distance in the parameter space.

The lemma follows.

With the aid of the lemma, we can take care of overlaps between coordinate neighborhoods of types (i) and type (iii), thus completing the proof of the main theorem.

6. Extending quasiconformal motions

What is the right definition for a quasiconformal map between two arbitrary subsets of $\hat{\mathbf{C}}$? It does not work to study only what happens to the sphere of a given radius about a point in the set: such a condition would say absolutely nothing about any of the geometric Cantor sets contained on the line which are obtained by iteratively removing the middle α subinterval, when $\alpha < 1/3$, since the spheres about points in such a Cantor set intersect the Cantor set in at most one point! Similarly, it does not work to look just at what happens to quadruplets of points whose cross-ratios lie in any given compact set, because one can construct Cantor sets such that the cross-ratios of points in that set avoid the given compact set of cross-ratios.

Definition. A quasiconformal homeomorphism $f: X \rightarrow Y$ between subsets of $\hat{\mathbf{C}}$ is a homeomorphism such that the cross-ratio of any quadruple of points in X has a bounded distance from the cross-ratio of the image points in Y . The distance between

cross-ratios is measured using the Poincaré metric on the space of values of the cross-ratio function, namely the thrice-punctured sphere $\mathbb{C} - \{0, 1, \infty\}$. A *quasiconformal motion* of a subset $X \subset \mathbb{C}$ is a map from $T \times X$ to $\hat{\mathbb{C}}$, where T is a connected space (frequently an interval of time), such that

- (i) for some basepoint $t_0 \in T$, $f_{t_0} = \text{id}$, and
- (ii) for any $t \in T$ and for any $\varepsilon > 0$ there is a neighborhood of t such for any quadruple of points in X , cross-ratios of the various images of the quadruple for time parameters in the neighborhood all lie within an ε -ball in the Poincaré metric of the thrice-punctured sphere.

This definition precisely captures the property of holomorphic motions which we used to make the definitions for local holomorphic motions. The definitions immediately suggest the following statement:

THEOREM 3. *For any quasiconformal motion of any subset X in \mathbb{C} whose time parameter space is an interval I , there is an extension of the motion to all of \mathbb{C} defined over all of I .*

Proof. We may suppose that two finite points of X are 0 and 1, and that these two points do not move.

First we extend the motion to the closure of X , by passing to limits of sequences of motions of points. The motion of other points is equicontinuous in the Poincaré metric of $\mathbb{C} - \{0, 1\}$ (since the cross-ratio of 0, 1, ∞ and x changes at a controlled rate) so the closure of the set of functions $f_*(x)$ which define the motion of points is compact. Two distinct functions g and h in the closure of this set of functions have disjoint graphs, since the cross-ratio of 0, 1, $g(x)$, $h(x)$ changes at a controlled rate, so it can never reach 0, 1, or ∞ in a finite time.

For small neighborhoods in the T plane, the extension of the motion from the closure of X to all of \mathbb{C} is done just as for the extension of holomorphic motions. These extensions can be pieced together using a partition of unity of I , to obtain a global quasiconformal extension.

7. Naturality

It would be good if we could define an extension of a holomorphic motion which would be completely canonical—in particular, a canonical extension would entail that

- (a) it would be invariant under any change of coordinates by a Möbius transformation, and

(b) it would be independent of choice of origin in the T plane.

Condition (b) seems quite hard. In particular, it would entail that the motion could be extended globally, over the entire T plane. In the companion paper by Bers and Royden, they give a construction of an extension of a holomorphic motion which depends on the choice of origin in the T plane but is otherwise canonical, so that property (a) is satisfied.

Condition (a) does not come as easily from our point of view, since the construction we made involved arbitrary choices of local holomorphic motions, subject to certain inequalities. To get around the fact that our choices are arbitrary, we can average different choices, but another difficulty crops up: the averaging procedure we have used itself depends on an arbitrary choice, the choice of the point at infinity. A probability measure on \mathbb{C} has a well-defined mean, provided the measure is not too dense near ∞ . When the measure is transformed by a Möbius transformation, its mean is not usually the image of its old mean.

To circumvent this problem, we can define a more general mean: if μ and ν are probability measures on the Riemann sphere, we will define the mean of μ with respect to ν .

First, if ν is a measure concentrated on a single point Q , then the mean of μ with respect to ν is the obvious mean, calculated by sending Q to infinity, forming the mean of μ , and transforming back to our original coordinates.

For arbitrary ν and μ , the mean will not always exist, but if there are disjoint round disks containing the support of ν and the support of μ , that is sufficient to guarantee existence. The idea in the general case is that the measure ν defines a transformation T_ν from probability measures on the Riemann sphere to probability measures on the Riemann sphere. The transformation is defined by taking the mean of mean of μ with respect to the various points in the support of ν , and weighting them according to ν .

LEMMA. *If the support of μ is contained in a disk D_1 and the support of ν is contained in a disjoint disk D_2 , then the support of $T_\nu(\mu)$ is contained in a proper subdisk D'_1 of D_1 .*

Proof. This is transparent.

COROLLARY. *If the support of μ and the support of ν are contained in disjoint disks, then the sequence of measures $T_\nu^n(\mu)$ converge to a measure concentrated at a point.*

Proof. The sequence of images of μ by T_ν is supported in a shrinking sequence of disks. If minimal support disks were to converge to disks of finite radius, then there would be a subsequence of the measures which would converge weakly to a measure such that T_ν would violate the previous lemma.

Definition. The mean of one probability measure μ on the Riemann sphere with respect to another measure ν is defined to be the support of the limit of $T_\nu^{on}(\mu)$, provided this limit exists and is supported on a single point.

Now we are prepared to construct an extension of a holomorphic motions of a set X which is equivariant with respect to any group Γ of Möbius transformations which preserves the original motion. The extension to the closure of X is immediate from the lambda lemma, so we may assume that X is closed. We may also suppose that the Γ is a closed subgroup of the group of Möbius transformations—otherwise, form its closure.

To extend the motion equivariantly to the complement of a closed set, the idea is that we pick a collection of triples and quadruples of points in X together with a measure μ_p on this collection for any point p , to govern the motion of p . (The measure here is a generalization of a partition of unity.) The collections of triples and quadruples, and the measures on them can easily be made equivariant. Since thin parts of the complement of X are disjoint, we can also make sure that for each thin part there is exactly one triple or quadruple associated, which will govern the motion on the bulk of this thin part. We can make sure that the measures vary slowly, so that the total measure of any set of triples and quadruples has a bounded derivative with respect to p as measured with the Poincaré metric of the complement of X .

In addition, we can choose for each point $p \in S^2 - X$ a measure ν_p on X , in such a way that ν_p depends equivariantly on p and has bounded derivative with respect to p .

Now we can define the motion of p to be the mean with respect to ν_p of the motions of p weighted according to μ_p . This mean exists for a definite neighborhood of 0 in T , by the corollary. It depends holomorphically on t , since it is obtained by a procedure of averaging holomorphic motions. The motions of different points are disjoint for a definite neighborhood in T , by the same considerations as before.

References

- [1] BERS, L. & ROYDEN, H. L., Holomorphic families of injections. *Acta Math.*, 157 (1986), 259–286.
- [2] MAÑÉ, R., SAD, P. & SULLIVAN, D., On the dynamics of rational maps *Ann. Sci. Ecole Norm. Sup.*, 16 (1983), 193–217.

Received August 20, 1984

Received in revised form August 14, 1985