Lecture Notes in Mathematics

Kevin M. Pilgrim

Combinations of Complex Dynamical Systems

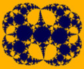

Springer

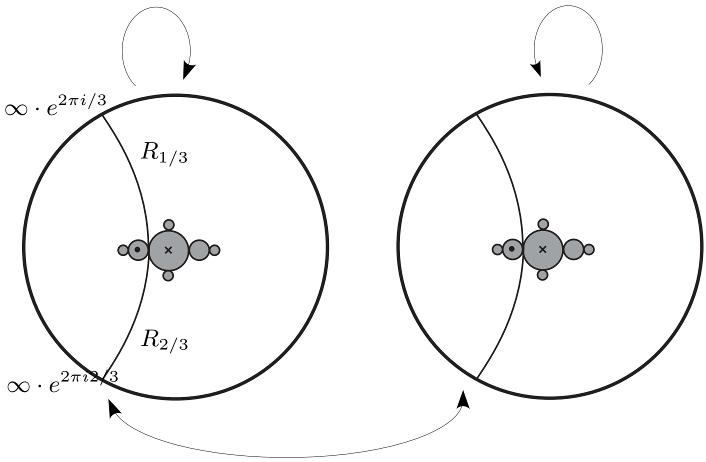

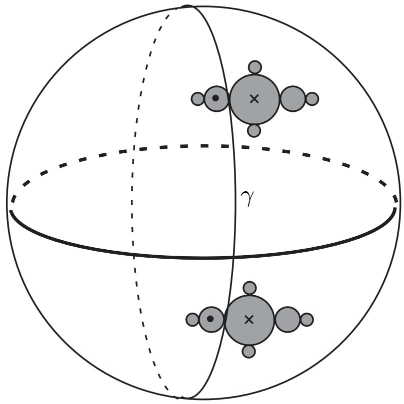

Fig. 1.2. An obstructed mating. Postcritical points are indicated by solid dots and critical points by crosses. The two overlapping crosses and dots correspond to the two period 2 critical points.

Identify $S^2 \times \{x, y\}$ with $\widehat{\mathbb{C}} \times \{x, y\}$ via a homeomorphism so that the postcritical set of $\mathcal{F}$ is $\{0, \infty\} \times \{x, y\}$. With a suitable generalization of the notion of combinatorial equivalence to maps defined on unions of spheres (see §4.2), $\mathcal{F}$ is combinatorially equivalent to the map which sends $(z, y) \to (z^2, x)$ and $(z, x) \to (z, y)$.

1.3.3 An obstructed expanding Thurston map

Here is a general construction. Let $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{GL}_2(\mathbb{R})$ be a matrix with integral coefficients. The linear map $\mathbb{R}^2 \to \mathbb{R}^2$ defined by A preserves the lattice $\mathbb{Z}^2$ and thus descends to an endomorphism $T_A : T^2 \to T^2$ of the torus

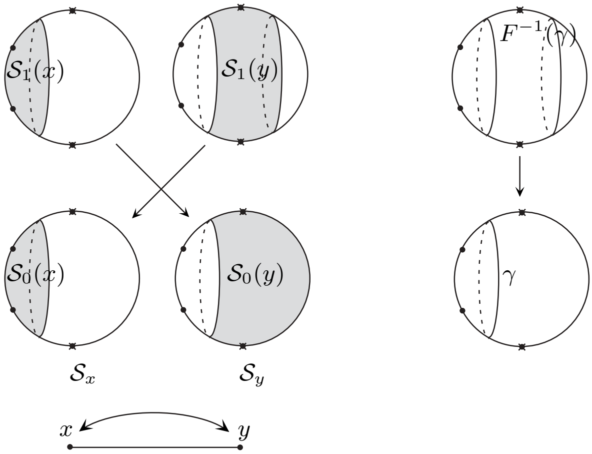

Fig. 1.3. Decomposition of the obstructed mating by cutting along the obstruction. Postcritical points are indicated by solid dots and critical points by crosses. The two overlapping crosses and dots correspond to the two period 2 critical points.

$T^2 = \mathbb{R}^2/\mathbb{Z}^2$. This endomorphism commutes with the involution $\iota : (x, y) \to (-x, -y)$. The quotient space $T^2/(v \sim \iota(v))$ is topologically a sphere S^2 and so T_A descends to a map $F_A : S^2 \to S^2$. The set of critical values of F_A is the image on the sphere of the set of points of order at most two on the torus. Since the endomorphism on the torus must preserve this set of four points, F_A is postcritically finite.

If e.g. $A = \begin{pmatrix} 3 & 0 \\ 0 & 2 \end{pmatrix}$ then F_A is expanding with respect to the orbifold metric inherited from the Euclidean metric on the torus. Let γ be the curve which is the image of the line $x = 1/4$. Then $\Gamma = \{\gamma\}$ is a multicurve whose Thurston matrix is $(1/2+1/2+1/2) = (3/2)$ and is therefore an obstruction; see Figure 1.4 where the metric sphere is represented as a "rectangular pillowcase" i.e. the union of two rectangles along their common boundary.

Using a similar decomposition process as in the previous example, we may produce a map $\mathcal{F} : S^2 \times \{x, y\} \to S^2 \times \{x, y\}$, this time sending each component to itself by a degree two branched covering.

Note that since the components of $F^{-1}(\gamma)$ map by degree two, the extension over the complements of $\mathcal{S}_1(x), \mathcal{S}_1(y)$ is now more complicated. Again, to keep things as simple as possible, we extend so that these complementary components, which are disks, map onto their images (again disks) by a quadratic branched covering which is ramified at a single point (say at z_x, z_y) which we arrange to be fixed points of $\mathcal{F}$.

It turns out that the resulting map $\mathcal{F}$ is combinatorially equivalent to the map of $\widehat{\mathbb{C}} \times \{x, y\}$ to itself given by $(z, x) \mapsto (z^2 - 2, x)$ and $(z, y) \mapsto (z^2 - 2, y)$ (the points z_x, z_y are identified with the point $\infty \in \widehat{\mathbb{C}}$).

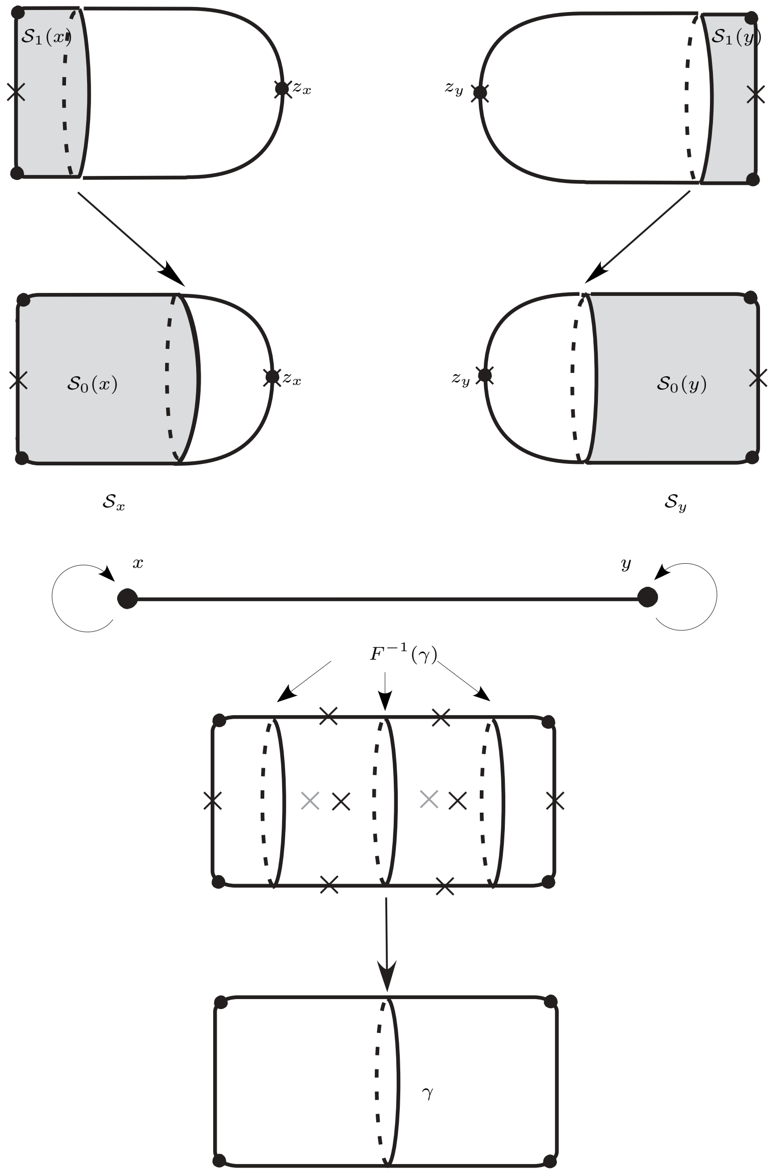

Fig. 1.4. An obstructed expanding map.

Note, however, that a great deal of information is lost in this naive decomposition: the degree of $\mathcal{F}$ is two, whereas the degree of the original map is six. The method of decomposition we will present in §5 will proceed roughly in the same manner presented in the above examples, but with the postcritical set P_F replaced by its full inverse image, $Q_F = F^{-1}(P_F)$. Since we are also interested in recovering the original map F from $\mathcal{F}$ together with some other data, we greatly refine the definition of the "trees" and their dynamics shown in Figures 1.3 and 1.4.

1.3.4 A subdivision rule

Another source of examples, of which the one below is prototypical, comes from *finite oriented subdivision rules with edge-pairings*, as introduced by Canyon, Floyd, Kenyon, and Parry [CFP3]. Again, regard the sphere as the quotient space of two Euclidean squares A and B whose oriented boundaries are identified as shown in Figure 1.5.

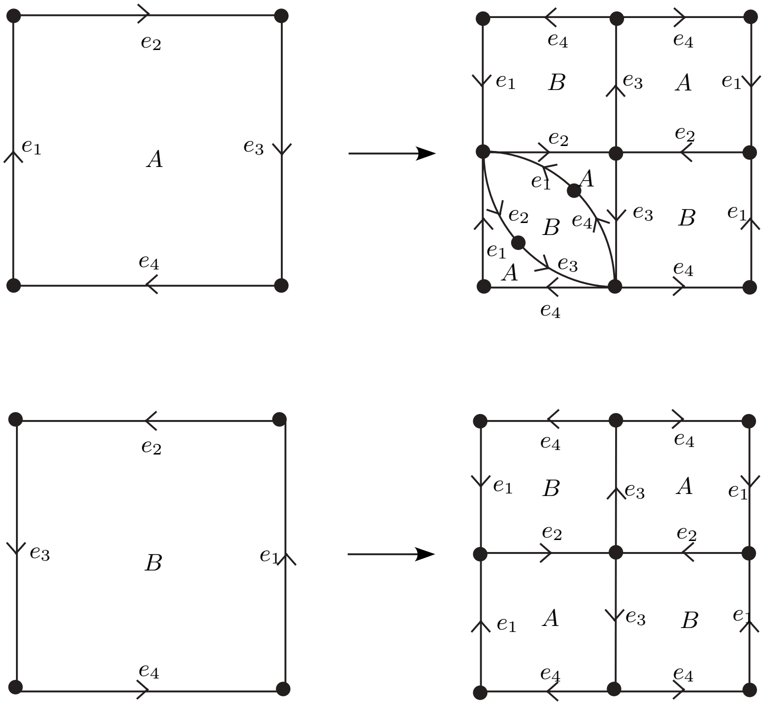

Fig. 1.5. A subdivision rule.

We regard this as a CW-structure on the sphere. A subdivision rule, loosely speaking, is a procedure for refining this CW structure to obtain a new CW-structure on the sphere. In Figure 1.5, the arrows indicate this process of refinement. To produce a branched covering F, note that a choice of orientation-preserving maps of cells which sends every oriented 1- and 2- cell on the right to the unique cell on the left having the same label descends to a well-defined degree five cellular map on the sphere which is cellular with respect to the cell structures on the right and left spheres. Differing choices yield combinatorially equivalent maps of the sphere.

This map may be produced from the Lattès example with $A = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}$ by the combinatorial surgery procedure of "blowing up an arc" [PT]; see §1.5.2. I do not know if this F is combinatorially equivalent to a rational map. Presumably, there is a metric on the sphere which is expanded under F. This is clear combinatorially: one application of F refines every 1- and 2-cell.

This example generalizes; we shall discuss motivation for considering branched coverings which arise from subdivision rules in §1.7.3.

1.4 Summary of this work

On the surface, Thurston's characterization theorem [1.1] seems like the end of the story of the classification problem. However, there are still many areas of incomplete understanding:

1. Thurston's characterization is implicit and involves checking *a priori* infinitely many conditions. There is no known general algorithm which decides whether or not a Thurston map is obstructed.
2. There are no known general methods for implementing Thurston's iterative algorithm. Apart from the numerics, the obstruction is the lack of a means for numerically approximating a rational function with prescribed critical values and prescribed combinatorics as a (non-dynamical) branched covering of the sphere. This is a very hard problem, even for polynomials ramified only over zero, one, and infinity–see [BS].
3. There are no known general methods for locating the canonical obstruction Γ_c, if it exists.
4. There is no known way to effectively enumerate postcritically finite rational maps.
5. Since Γ_c is canonical, it seems reasonable that cutting along Γ^c should result in "simpler" maps which might be easier to analyze. However, there is no extant general theory of combinations and decompositions.

The main goal of this work is to provide a solution to Problem (5) above. We shall give:

- a combination procedure (Theorem 3.2), taking as input a list of data consisting of seven objects satisyfing fourteen axioms, and producing as output a well-defined branched mapping F of the sphere to itself;

- an analysis of how the combinatorial class of F depends on the input data (Theorem 4.5), as well as explicit bounds on the number of classes of maps F which can be produced by varying certain portions of the data and keeping others fixed (Corollary 4.6, Theorem 7.1);
- a decomposition procedure, taking as input a branched mapping F and producing as output such a list of input data, in a manner which is inverse to combination (Theorem 5.1);
- an analysis of how the result of decomposition depends on F and some choices used in the decomposition process (Theorem 6.1);
- a structure theorem for postcritically finite branched mappings (Theorem 10.2), informally stated as follows:

  Canonical Decomposition Theorem: *A Thurston map F is, in a* canonical *fashion, decomposable along a multicurve Γ^c into "pieces", each of which is of one of three possible types:*
  1. *(elliptic case) a homeomorphism of spheres,*
  2. *(parabolic case) covered by a homeomorphism of planes, or*
  3. *(hyperbolic, rational case) equivalent to a rational map of spheres.*

  A priori $\Gamma^c \supset \Gamma_c$. Unfortunately, we do not know if $\Gamma^c = \Gamma_c$–at present our arguments require inductively cutting along canonical obstructions in a process that must terminate. We conjecture that only one step is needed.
- applications of our analysis to the structure of (combinatorial) symmetry groups of Thurston mappings (Theorem 8.2).

  For a finite, nonempty set Q in S^2, let $\mathrm{Mod}(S^2, Q)$ denote the *mapping class group* of orientation-preserving homeomorphisms of S^2 to itself which send Q to itself, modulo isotopy through homeomorphisms fixing Q. Given a Thurston map F, let $Q = F^{-1}(P_F)$, and let $\mathrm{Mod}(F)$ denote the subgroup of $\mathrm{Mod}(S^2, Q)$ represented by maps α for which $\alpha \circ F \circ \alpha^{-1}$ is combinatorially equivalent to F.

  Informally, the two main results are the following:

  Theorem: $\mathrm{Mod}(F)$ *reduces along Γ^c. That is, every element α of* $\mathrm{Mod}(F)$ *sends Γ^c to itself, up to isotopy relative to Q.*

  Twist Theorem: *Let F be a Thurston map and Γ an invariant multicurve. If* 1 *is an eigenvalue of the Thurston linear map F_Γ, then* $\mathrm{Mod}(F)$ *contains a free abelian group of rank ≥ 1.*
- an analysis of what happens when two Thurston obstructions intersect (Theorem 8.7);
- examples from complex dynamics (§9) where we generalize existing combination procedures, e.g. mating.

Here is a summary of the remainder of this Introduction.

§1.5 is a survey of known results regarding the combinatorics of complex dynamical systems. I have tried to give as complete a bibliography as possible, as much of this material is unpublished and/or scattered. Often, references are merely listed without further discussion of their contents. I apologize for any omissions.

§1.6 develops some topological aspects of the "dictionary" between rational maps and Kleinian groups as dynamical systems. In particular, we propose to view the Canonical Decomposition Theorem as an analog of the JSJ decomposition of a closed irreducible three-manifold.

§1.7 discusses connections between the analysis of postcritically finite rational maps and other, non-dynamical topics (e.g. geometric Galois theory; groups of intermediate growth; Cannon's conjecture on hyperbolic groups with two-sphere boundary).

§1.8 discusses the combinatorial subtleties which necessarily arise when trying to glue together noninvertible maps of the sphere. It is an important preamble to the body of this work and should be read before continuing to §2, since terminology and notation used throughout this work is introduced.

§1.9 discusses regularity issues in the definition of combinatorial equivalence. The decomposition and combination procedures in this work are developed for non-postcritically finite maps as well. For such maps, however, there are competing notions for combinatorial equivalence.

1.5 Survey of previous results

In this subsection, we attempt to give a fairly complete survey of results to date concerning the combinatorial aspects of the dynamics of rational maps, focusing on those aspects pertaining to combinations, decompositions, and structure of maps which are nice, e.g. postcritically finite, geometrically finite, or hyperbolic (see below for definitions). We assume some familiarity with one-variable complex dynamics; see e.g. the text by Milnor [Mil4]. In many places we simply state the flavor of the results and give references.

1.5.1 Enumeration

The rigidity portion of Thurston's characterization implies that in principle it should be possible to enumerate postcritically finite rational maps by enumerating the corresponding combinatorial objects (branched coverings). However, no general reasonable enumeration of postcritically finite rational functions or Thurston maps is known, mainly due to the immense combinatorial complexity of the set of such maps. Partial and related results include the following.

Polynomials. In the restricted setting of postcritically finite polynomials, such an enumeration is possible. Let $\Delta \subset \mathbb{C}$ denote the open unit disk, S^1 its boundary, and identify S^1 with $\mathbb{R}/\mathbb{Z}$ via the map $t \mapsto \exp(2\pi i t)$.

Definition 1.4 (Lamination). *A* lamination *is an equivalence relation on S^1 such that the convex hulls of equivalence classes are disjoint.*

Now let $K \subset \mathbb{C}$ be a nondegenerate (i.e. contains more than one point) continuum whose complement is connected. There is a unique Riemann map $\phi : \widehat{\mathbb{C}} - \Delta \to \widehat{\mathbb{C}} - K$ such that $\phi(\infty) = \infty$ and $\phi(z)/z \to \lambda > 0$ as $z \to \infty$. The set $R_t = \{\phi(r \exp(2\pi i t)|1 < r < \infty\}$ is called an *external ray of angle* t and the ray R_t is said to *land* at a point in $z \in \partial K$ if $\lim_{r\downarrow 1} \phi(r \exp(2\pi i t)) = z$. From classical theorems of complex analysis, it is known that almost every ray (with respect to Lebesgue measure on $\mathbb{R}/\mathbb{Z}$) lands, and that K is locally connected if and only if ϕ extends continuously to $\widehat{\mathbb{C}} - \Delta$. The *lamination* associated to K is defined by $s \sim t$ if and only if the external rays R_s, R_t land at the same point. This is indeed a lamination, since distinct external rays cannot intersect and two simple closed curves on the sphere cannot cross at a single point.

Now let f be a monic polynomial, and let $K_f = \{z \in \mathbb{C}| f^{\circ n}(z) \not\to \infty\}$ denote the *filled-in Julia set* of f. It is known that K_f is connected if and only if every finite critical point belongs to K_f. At this point, we digress to define

Definition 1.5 (Mandelbrot set). *The* Mandelbrot set M *is the set of those* $c \in \mathbb{C}$ *for which the filled-in Julia set* K_c *of* $z^2 + c$ *is connected.*

If the filled-in Julia set of a monic polynomial f is connected, then we may apply the above construction to speak of external rays, etc. as so define the *lamination* Λ_f associated to K_f. The *rational lamination* $\Lambda_f^{\mathbb{Q}}$ is the restriction of Λ_f to $\mathbb{Q}/\mathbb{Z}$. Obviously, the lamination $\Lambda_f^{\mathbb{Q}}$ must satisfy certain invariance conditions since it comes from a polynomial. The landing points of rational rays are necessarily periodic or preperiodic points which are either repelling or parabolic; conversely, every point which iterates onto a repelling or parabolic cycle is the landing point of some rational ray; see [Mil4] and the references therein. Kiwi [Kiw] has characterized those rational laminations which arise from polynomials; the analysis of postcritically finite maps plays a key role in the proof. For more on these kinds of laminations, see also [BL1], [BL2], [Kel], [Ree3], [Thu2].

Postcritically finite polynomials. The Julia set of any postcritically finite rational map is connected and locally connected ([Mil4], Thm. 19.7). Now suppose f is a postcritically finite polynomial. Then the Riemann map ϕ to the complement of K_f extends continuously to the closed disk. It follows that $\Lambda_f^{\mathbb{Q}}$ determines the entire lamination Λ_f. This in turn permits one to reconstruct a branched covering equivalent to f. Thurston rigidity then implies that the rational lamination determines f as long as f is postcritically finite.

In fact, in the setting of postcritically finite polynomials, one can encode f using far less data. Bielefeld, Fisher, and Hubbard [BFH] gave a precise combinatorial enumeration of critically *pre* periodic polynomials in terms of angle conditions on external rays landing at critical values, i.e. by considering a subset of the rational lamination. Milnor and Goldberg [GM] used angles of rays landing at fixed points to develop a conjectural description of all polynomials which was completed by Poirier in [Poi1]. Poirier then gave an equivalent classification of arbitrary postcritically finite polynomials using *Hubbard trees* as combinatorial objects [Poi2]. Hubbard trees are certain planar trees equipped with self-maps satisfying certain rather natural expansivity, minimality, and topological criteria which allow them to be good mimics of the dynamics of such a polynomial. For more on Hubbard trees, see also [AF], [Dou].

In summary, it seems fair to say that the enumeration problem for postcritically finite polynomials is solved, either using laminations, portraits, or Hubbard trees. In particular, for quadratic postcritically finite polynomials of the form $z \mapsto z^2 + c$, the rational lamination, and hence the polynomial itself, is faithfully encoded by a single rational number $\mu \in \mathbb{Q}/\mathbb{Z}$–e.g. if the critical point at the origin is preperiodic, then it lies in the Julia set, and μ is the smallest rational number in $[0, 1)$ such that R_μ lands at c. There is even an explicit algorithm to reconstruct the lamination from μ; see [Dou], [DH1], [Kel], [Lav], [Thu2].

Quadratic postcritically finite rational maps. For postcritically finite quadratic branched coverings, M. Rees [Ree3], [Ree4] developed a sophisticated program for describing such maps in terms of polynomials. A difficulty is that a priori a given quadratic rational function might admit many such descriptions. That is, unlike the case for polynomials, it is unclear how to associate to a general quadratic Thurston map or rational map a normal form, i.e. minimal set of combinatorial data, necessary to determine the map.

General postcritically finite rational maps. Invariants. Indeed, enumerating even simple postcritically finite rational maps is hard. For example, tabulating just the hyperbolic, non-polynomial rational functions of, say, low degree (2 or 3) and small postcritical set (2,3, or 4 points) was a fairly formidable task [BBL$^+$]. For fixed degree and size of postcritical set, it is shown that there can exist infinitely many combinatorially inequivalent branched maps, of which at most finitely many can be equivalent to rational functions.

It is therefore natural to seek combinatorial invariants of Thurston maps. An algebraic formulation of combinatorial equivalence has been developed by Kameyama [Kam2] (cf. also [Pil4]) and the author [Pil3]. This is a somewhat promising development, as the problem of deciding when two branched coverings are combinatorially equivalent is reduced to a computational problem in group theory.

A natural class of maps lying between rational and general Thurston maps are those Thurston maps which are expanding. For fairly general reasons, a degree d expanding map is a quotient via a "coding map" of a one-sided shift on d symbols, and the equivalence relation determining the fibers is also a subshift of finite type (see [Fri] and also the chapter on semi-Markovian spaces in [CP]). In [Kam3], [Kam5] coding maps and the structure of the set of coding maps are investigated.

1.5.2 Combinations and decompositions

No general theory of combinations and decompositions of Thurston mappings, in the combinatorial category, had been developed. For Thurston maps, there is no combination procedure in which a precise analysis of the dependency of the output on the input data is given. At present, our discussion focuses on combinations in the topological category: starting with e.g. two rational maps, a topological combination procedure results in a Thurston map F. One then asks whether F is combinatorially equivalent to a rational map R. Obtaining effective control on the location and structure of potential obstructions is crucial. Note that in this category, the topological dynamics of F is not typically relevant. Other combination procedures which proceed by producing the topological dynamics of R directly using conformal or quasiconformal surgery are briefly mentioned in §1.5.4.

Results relating to existing combination theorems include the following.

Quadratic Matings. The definition of formal mating has already been given in §1.3. There are other notions of mating, in which two polynomials f, g are called mateable if the branched covering $F = f \sqcup g$ has at worst certain "removable" obstructions. Milnor [Mil5], Shishikura [Shi4], Tan [Tan2] and Wittner [Wit] discuss relationships between the different notions. These other notions of mating apply to polynomials which are not necessarily postcritically finite.

More precisely: let $S^2_{f,g}/\sim$ be the quotient space of the sphere on which F is defined obtained by collapsing the closures of external rays to points. One says that f, g are *matable* if (i) this quotient space is a sphere, and (ii) the map of this space to itself induced by F is conjugate to a rational map via a topological conjugacy which is holomorphic on the interiors of the filled-in Julia sets of f and g. For examples and more details, see [Mil5], [YZ], [Luo]. If a rational map of degree d is a mating, the dynamics on its Julia set is a quotient of $z \mapsto z^d$ on the unit circle. The converse is nearly true [Kam4]. Since our focus is on those aspects of the combinatorial theory which admit generalizations beyond polynomials, we will not further discuss these notions and will instead focus on the postcritically finite case here.

A quadratic polynomial $f(z) = z^2+c$ with connected Julia set has typically two fixed points. The landing point of the zero angle external ray is the *β-fixed point*; the other is the α fixed point. Suppose q rays land at the α-fixed point. Then the set of q such rays is cyclically permuted, with rotation number p/q, for some $1 < p < q$. In this case, we say c is in the *p/q-limb* of the Mandelbrot set. The *conjugate* of the p/q limb is the $1 - p/q$ limb.

The investigations of Douady and Hubbard led to formulation of the *Quadratic Mating Conjecture* for quadratic polynomials, now a theorem:

Theorem 1.6 (Quadratic Mating Theorem). *Two postcritically finite quadratic polynomials f, g are matable if and only if f, g do not belong to conjugate limbs of the Mandelbrot set.*

Note that the hypothesis is purely combinatorial. Necessity is clear: if f, g are in conjugate limbs, then the rays landing at the α fixed point join together, and the quotient space $S^2_{f,g}/\sim$ is not a sphere (cf. Figure 1.2). The sufficiency is more subtle. Levy [Lev] used Thurston's Characterization Theorem [1.1] to reduce the proof to ruling out the existence of very special obstructions, now called *Levy cycles*. A *Levy cycle* is a multicurve $\Gamma = \{\gamma_0, \gamma_1, \ldots, \gamma_{n-1}\}$ such that each γ_i has an inverse image mapping by degree one and which is homotopic to γ_{i+1} (subscripts modulo n). Like the obstruction in the example in §1.3.2, Levy cycles are obstructions to the existence of a metric which is expanding for F.

Other partial results were later obtained by Tan [Tan1]. The first complete proof was given by Rees [Ree1]. Tan [Tan2] gave a simpler proof which generalized to maps with two critical points. Rees [Ree1] and Shishikura [Shi4] refined the analysis to show that a natural quotient of the mating of two polynomials f and g, when it exists, is actually topologically conjugate to the equivalent rational map.

Wittner [Wit] also gave a detailed analysis of the phenomena of *shared matings*, i.e. rational functions expressible as matings in essentially distinct ways. This is a fascinating topic which seems to be related to the geometry of parameter space [Eps1]. The current record-holder is the Lattès map

$$R(z) = z\frac{z+\eta^2}{\eta^2 z+1}, \quad \eta^2 = (3 \pm i\sqrt{7})/2$$

which is expressible as a mating in four essentially distinct ways ([Mil5], B.9).

Higher degree matings. Shishikura and Tan [SL] analyzed matings of certain postcritically finite cubic polynomials, and found that the question of determining when the mating is equivalent to a rational map is much more subtle than in the quadratic case–the obstructions need not be of the special, Levy cycle kind, and are much more difficult to control.

Tunings in the quadratic rational family. The concept of tuning was first given for interval maps by Milnor and Thurston. In the complex quadratic setting, it may be described as follows. Given a quadratic polynomial f (or, more generally, a rational map or branched covering) with a periodic simple critical point c whose orbit contains no other critical points, and given quadratic polynomial p, the *tuning of f by p* is the branched covering F obtained roughly as follows. Let c have period $m \geq 2$. Write the cycle as $0 = c_0 \mapsto c = c_1 \mapsto c_2 \mapsto \ldots \mapsto c_{m-1} \mapsto c_0 = 0$. Let D_i be the closure of the immediate basin of c_i. Then each D_i is a disk which we may identify with a copy of the compactified complex plane $\tilde{\mathbb{C}} \times \{i\}$ by adding the circle at infinity to $\mathbb{C}$. Send $\tilde{\mathbb{C}} \times \{0\}$ to $\tilde{\mathbb{C}} \times \{1\}$ by p and $\tilde{\mathbb{C}} \times \{i\}$ to $\tilde{\mathbb{C}} \times \{i+1\}$ by the identity map. The result descends to a well-defined branched covering F. See [Mil1].

Rees [Ree3], [Ree4] generalized the notion of tuning from polynomials to arbitrary (quadratic) branched coverings. She proved that there are no obstructions to realizing the tuning of a rational map by a quadratic polynomial. D. Ahmadi [Ahm] considered generalizations of tuning to the case when one has two critical points in the same cycle.

There appears to be some ambiguity in the definition of tuning given in [Ree3], §1.20, when "type III" maps, which have a preperiodic critical point, are considered–it is not evident that the combinatorial class (see §1.5.1) of the output is well-defined, given the input data. The reason is as follows. Suppose c is a periodic critical point of f along whose orbit the gluing is performed, and suppose d is a preperiodic critical point of f mapping onto c under iteration. After gluing, the forward orbit of d under the new map F is not well-defined unless some additional data is prescribed. This potential ambiguity causes a number of complications, is discussed in more detail in §3.2, and is dealt with by our addition of "critical gluing data" to the input data for a combination procedure.

Tunings in arbitrary degree. In [Pil5] a notion of tuning wherein a general branched covering f is tuned by a family $\mathcal{P}$ of polynomials is given. If f is a rational map satisfying a certain condition ("acylindrical"; compare §1.6) and if $\mathcal{P}$ is "starlike", then it is proven that the tuning of f by $\mathcal{P}$ is equivalent to a rational map. A major ingredient is a tool allowing control of potential obstructions which evolved into the *Arcs intersecting obstructions theorem*, ([PT], Thm. 3.2).

Generalized matings and tunings. Using the Arcs intersecting obstructions theorem, Tan and the author [PT] gave mild generalizations of the mating construction and gave examples of conditions under which the procedure produces branched coverings equivalent to rational maps. In §9.3 we will give a significant generalization of mating, though we do not give conditions for such matings to be equivalent to rational maps.

Remarks. All of the combination procedures mentioned above have one feature in common: they produce maps with an invariant multicurve. Moreover, apart from generalized matings and tunings, the degree of the output is equal to the degree of the input.

When the input data for the above combination theorems consists of rational maps (as opposed to branched coverings), the geometry and dynamics of these maps implies that the combinations can be defined in a totally unambiguous fashion. That is, the resulting branched covering of the sphere is well-defined. However, while it is easy to formulate analogs of these procedures for branched coverings, these combination procedures inevitably depend on a variety of choices, and there has appeared yet no explicit discussion of the dependence of the result of combination on these choices.

Blowing up an arc. The operation of *blowing up an arc* [PT] increases the degree. Let α be an arc whose interior lies in $S^2 - P_F$ and whose endpoints lie in P_F. Suppose $F|_\alpha$ is a homeomorphism onto its image. To blow up α, cut the sphere open along α and separate the edges of the cut. Send the complement of the incision via F, and the newly created disk by a homeomorphism onto the exterior of $F(\alpha)$. Under suitable invariance assumptions on α, the result is equivalent to a rational map. A highly useful technical device, the *Arcs intersecting obstructions theorem*, is presented there as a means of controlling the locations of potential obstructions.

Captures. Luo [Luo], Rees [Ree3], and Wittner [Wit] also considered operations called *captures* in the quadratic family. A *capture* is a branched covering $F : S^2 \to S^2$ obtained, for example, as follows. Let $p : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ be a quadratic polynomial with a periodic critical point at $z = 0$. Let x be a preperiodic point in the backward orbit of 0, and let $\beta : [a, b] \to \widehat{\mathbb{C}} = S^2$ be a path running from ∞ to x and avoiding the forward orbit of x. Let $\sigma_\beta : S^2 \to S^2$ be a homeomorphism which is the identity off a small neighborhood of the image of β such that $\sigma(\infty) = x$. The map $F = \sigma_\beta \circ p$ is called a *capture.*

Miscellaneous. While not strictly speaking a combination procedure, Kameyama [Kam1] gives conditions under which a self-similar subset of the sphere is homeomorphic to the Julia set of a rational map.

1.5.3 Parameter space

A major motivation for the development of a combinatorial analysis of postcritically finite rational maps was the desire to understand the rich combinatorial structure seen in pictures of parameter spaces. Studies of this kind are too numerous to be comprehensively listed here, so our account below is necessarily selective. In particular we do not mention the large literature related to real maps.

Definition 1.7 (Hyperbolic map). *A rational map $f : \widehat{\mathbb{C}} \to \widehat{\mathbb{C}}$ is* hyperbolic *if every critical point converges to an attracting cycle under iteration.*

Equivalently: f is expanding on a neighborhood of its Julia set with respect to the Poincaré metric on the complement of the postcritical set. The condition of being hyperbolic is an open condition in the complex manifold Rat_d of rational maps of a given degree; a connected component of the set of hyperbolic maps is called a *hyperbolic component* in parameter space. It is natural to ask how hyperbolic components are deployed in parameter space; when their closures intersect, and how; how these connections are related to the combinatorics of the maps involved; when their closures in the moduli space $\mathrm{Rat}_d/\mathrm{Aut}(\widehat{\mathbb{C}})$ are compact, etc. Conjecturally, hyperbolic maps are dense in Rat_d; see e.g. [Lyu], [McM6]. Any two maps in the same hyperbolic component are conjugate on a neighborhood of their respective Julia sets ([Kam3], Thm. 4.7). A hyperbolic component whose elements have connected Julia sets has a preferred unique "center point" which is postcritically finite [McM1]. Expanding components whose elements have connected Julia sets are essentially polydisks (see [Mil2], [Ree2]) while those corresponding to maps with disconnected Julia set may have more complicated topology (see [Mak1], [McM2]).

Quadratic polynomials. The combinatorial structure underlying the deployment of hyperbolic components within the Mandelbrot set is now very well understood. A few sources are [DR], [Dev], [DH1], [Kel], [Sch], [Thu2], [Kau].

Quadratic rational maps. It turns out that the moduli space of Möbius conjugacy classes of quadratic rational maps is biholomorphic to $\mathbb{C}^2$ [Mil3]. Rees' program [Ree3], [Ree4], [Ree5] yields detailed results on the combinatorial structure of the one-complex dimensional loci $V_n \subset \mathbb{C}^2$ of maps with a critical point of period n. The approach is by analogy with the theory for quadratic polynomials and proceeds using laminations. The last cited work also relates the homotopy types of various spaces of rational maps and their combinatorial analogs. Stimson [Sti] and Rees study the algebraic geometry (in particular, the singularities) of the loci V_n and its relation to combinatorial properties quadratic rational maps.

Other rational families. The combinatorics of certain families of rational functions have also been investigated. Head [Hea] and Tan [Tan3] consider the structure of the cubic rational maps obtained by applying Newton's method to cubic polynomials. Bernard [Ber] has investigated the combinatorial structure of rational maps near a Lattès example. Inninger [IP] gives an explicit real family of examples whose Julia sets are the whole sphere. Barański [Bar2], [Bar1] studies maps with two superattracting fixed points and finds examples of maps with fully invariant Fatou components which are not perturbations of polynomials.

Compactness properties. There have also been investigations into the relationship between the success and failure of combination theorems and the structure of parameter spaces, especially compactness properties of and tangencies between hyperbolic components. Rees [Ree2] gives compactness results for certain real one-parameter subsets of hyperbolic components, independent of the fine combinatorial details of the map. Petersen [Pet] relates the failure of mating to noncompactness of hyperbolic components in the quadratic rational family. Makienko [Mak2] gives sufficient conditions for the noncompactness of hyperbolic components; work of Tan [Tan5] gives simpler and more complete arguments for the same results. Epstein [Eps2] provides the first compactness results for a hyperbolic component, for quadratic rational maps in which both critical points lie in distinct periodic cycles of periods ≥ 2. In his thesis the author [Pil5] proposes a series of conjectures relating the combinatorics of rational maps, the topology of their Julia sets, the geometric realizability of topological combination theorems, and compactness properties of hyperbolic components; see also [McM5] and §1.6 for a survey of this topic.

To bring this down to earth, here is a concrete question which is still open:

Question: Let $f(z) = \frac{\sqrt{3}}{2}i(z + \frac{1}{z})$. The two critical points at ± 1 lie in the same cycle of period four. Does the hyperbolic component $H(f)$ have compact closure in the space of quadratic rational maps modulo Möbius conjugacy?

Geometry of combinations. Epstein shows by giving an explicit example that the operation of mating, when extended to hyperbolic and parabolic maps, is not continuous [Eps1]; the discontinuity phenomena is similar to that present in the intertwining surgery of Epstein and Yampolsky [EY] .

1.5.4 Combinations via quasiconformal surgery

In many cases, it is desirable to have a combination procedure which does not proceed via Thurston's characterization theorem [1.1]. Typically, this procedure is given by the process of *quasiconformal surgery*, which we now describe. The input is one or more rational functions $f, g, \ldots$ and some combinatorial data describing the surgery. The output is a rational function R. The surgery proceeds by first constructing a *K-quasiregular map* $\tilde{R}$, i.e. a map which is locally a rational map followed by a K-quasiconformal homeomorphism. Since quasiconformal maps are quite flexible, this step is often not especially difficult. A common method for producing R from $\tilde{R}$ is to check that the iterates of $\tilde{R}$ are uniformly K'-quasiregular (often, a very delicate and technical step) for some constant K', and then apply a theorem of Sullivan [Sul] (sometimes referred to as the "Shishikura principle") which asserts that under these assumptions, $\tilde{R}$ is quasiconformally conjugate to a rational map R.

Often, the dependence of quasiconformal surgery on the input maps can be explicitly controlled, yielding results about the geometry of parameter spaces. Examples of this are too numerous to list comprehensively but we point out here a few prototypical examples drawn from the themes of combination, decomposition, and applications to parameter spaces of rational maps. First is Douady and Hubbard's seminal paper on polynomial-like maps [DH2] where the technique is first introduced. Branner and Douady [BD] relate portions of the cubic and quadratic polynomial parameter spaces, while Branner and Fagella [BF] relate different limbs of the Mandelbrot set. Haïssinsky [Haï] interprets tuning via quasiconformal surgery. Epstein and Yampolsky [EY] start with a pair of quadratic polynomials and produce a cubic polynomial via "intertwining" surgery.

We note that surgery has played a prominent role in the analysis of maps with Siegel disks; cf. work of Ghys, Douady, Hermann, and Świątek as well as Petersen and Zakeri. The basic idea is to start with a generalized Blaschke product which sends the unit circle onto itself by a homeomorphism with some given rotation number, and to do surgery to produce a polynomial; the unit circle turns into the boundary of a Siegel disk after surgery. Also, Shishikura [Shi1], [Shi2] applies surgery profitably to the study of e.g. maps with Herman rings and Siegel disks; see also [Shi3] for trees associated with configurations of Herman rings on the sphere.

1.5.5 From p.f. to geometrically finite and beyond

It is tempting to speculate on the extent to which Thurston's characterization and rigidity theorem [1.1] generalizes beyond the postcritically finite setting.

Rational laminations were conjectured to be good combinatorial objects with which to classify polynomials. It was conjectured that a polynomial p with connected Julia set and all cycles repelling is uniquely determined by its rational lamination. (The condition, all cycles repelling, is present to rule out deformations supported on the Fatou set, which are uninteresting in this context.) In the quadratic family, this is equivalent to the famous "MLC" conjecture asserting that the Mandelbrot set is locally connected. If answered in the affirmative, it implies that the dynamics of every point in the Mandelbrot set is essentially faithfully encoded by a single number in $\mathbb{R}/\mathbb{Z}$. Partial results in support of this conjecture were begun by Yoccoz ([Hub]), who proved that the Mandelbrot set is locally connected at c under the assumption that z^2+c is *non-renormalizable* . Since then there have been a number of further improvements in which the hypothesis "non-renormalizable" has been successively weakened. For a good summary, see [Lyu] and the references therein.

A recent theorem of C. Henriksen [Hen], however, asserts that in the cubic family, this rigidity fails. A pair of cubic polynomials having the same rational lamination and having set-theoretically distinct dynamics on the forward orbits of their critical points is constructed; the proof uses the intertwining surgery of Epstein and Yampolsky. Thus the combinatorial classification of polynomials of degree greater than two is apt to be significantly more complicated.

Even for polynomial maps with totally disconnected Julia sets, such a classification may be very difficult. Emerson [Eme] associates to such a map an infinite tree, equipped with a self-map, which imitates the deployment of annuli in the complement of the Julia set bounded by level sets of the Green's function; in particular he shows that uncountably many combinatorially inequivalent such trees can arise among maps of a given degree.

Less ambitious is to seek first a generalization of Thurston's characterization to the setting of *geometrically finite* maps, i.e. to rational maps for which the intersection of the Julia and postcritical sets is finite. This is discussed briefly by Thurston [Thu2]. In the thesis of David Brown [Bro], an implementation of Thurston's algorithm for non-postcritically finite quadratic polynomials with connected Julia set is introduced. Turning to rational maps, the combinatorics of geometrically finite maps with disconnected Julia sets is discussed in [PL] and [Yin].

A major advance in the understanding of geometrically finite maps has been announced by Cui Guizhen [Cui]. We briefly summarize here the results of this work in progress. A *geometrically finite branched covering map* is a quasiregular map F with the following property. Let P'_F denote the set of accumulation points of the postcritical set P_F; note that this is forward-invariant under F. Then (i) P'_F is finite; (ii) F is holomorphic on a neighborhood of P'_F; (iii) each periodic point of P'_F is either attracting, superattracting, or parabolic. Two geometrically finite branched covering maps are *combinatorially equivalent* if there is a combinatorial equivalence h_0, h_1 between them such that (i) h_0, h_1 are quasiconformal; (ii) $h_0 = h_1$ on a neighborhood of attracting periodic points in P'_F and on attracting parabolic petals of parabolic points in P'_F. Starting with these definitions, he first proves a direct translation of Thurston's characterization and rigidity theorem for hyperbolic rational maps by showing that each such map can be obtained via quasiconformal surgery from a collection of postcritically finite maps and maps with disconnected Julia set. The combination procedure used and the issues involved are quite similar to those presented here.

To extend to maps with parabolics, Cui first shows that each hyperbolic map can be "pinched" to a map with parabolics without changing the dynamics on the Julia set. The proof uses novel distortion estimates, and represents the first instance of obtaining limit points of quasiconformal deformations via purely intrinsic methods. This result is then used to give a characterization of geometrically finite rational maps with parabolics among geometrically finite branched covers. An extra condition, "no connecting arcs" is needed: by "mating" of $z^2+1/4$ with itself one can obtain a geometrically finite branched covering with no Thurston obstructions but which is nonetheless not equivalent to a rational map–an arc, invariant up to isotopy relative to the postcritical set, joins the two parabolic fixed points. The corresponding rational map, however, has a degenerate parabolic point, and this arc should be collapsed to a point.

Having in hand now a generalization of Thurston's theorem, Cui proceeds to establish the existence of limits of various kinds of quasiconformal deformations of geometrically finite rational maps. The technique employed is to identify the limit as a geometrically finite branched cover, verify that it is unobstructed, and then show that the corresponding rational map can be perturbed back to recover the original path of deformations.

1.6 Analogy with three-manifolds

Here, we give some motivation from the topological aspects of the dictionary between the theories of rational maps and Kleinian groups as holomorphic dynamical systems on the Riemann sphere. For a comprehensive survey, see [McM5] and [McM4]. Note that regarding $\widehat{\mathbb{C}}$ as the boundary at infinity of hyperbolic three-space $\mathbb{H}^3$ gives a bijection between orientation-preserving isometries of $\mathbb{H}^3$ and the group $\mathrm{Aut}(\widehat{\mathbb{C}})$ of Möbius transformations.

Let M be a compact, oriented, irreducible (every embedded two-sphere bounds a three-ball) three-manifold. A connected, properly embedded, twosided surface $S \subset M$ (which is neither a sphere, projective plane, or disc isotopic into ∂M) is *incompressible* if the inclusion $S \hookrightarrow M$ induces an injection on fundamental groups. S is said to be *peripheral* if it is isotopic into ∂M. If M contains a nonperipheral incompressible surface S then M is called *Haken*; M is *toroidal* if it contains a nonperipheral incompressible torus. If ∂M is nonempty and incompressible, a *cylinder* in M is a nonperipheral incompressible annulus. The prototypical example of a cylindrical manifold is $M = S \times [-1, 1]$, where S is a closed surface. Nonperipheral tori are topological obstructions to finding a hyperbolic structure on M, and cylinders also play an important role in several respects:

Theorem 1.8. *Let N be a convex compact geometrically finite hyperbolic three-manifold with nonempty incompressible boundary. Then the following are equivalent.*

1. *N is acylindrical.*
2. *$\pi_1(N)$ does not split (as free product with amalgamation or HNN extension) over $\mathbb{Z}$.*
3. *Given any orientation-reversing "gluing" homeomorphism $h : \partial N \to \partial N$, the quotient manifold N/h admits a hyperbolic structure.*
4. *The deformation space of N has compact closure in the space of all hyperbolic structures on M.*
5. *The limit set of the fundamental group of N, regarded as a Kleinian group, is a Sierpinski carpet.*

Alternatively, one might replace condition (3) with the following: the limit of any deformation of N corresponding to pinching a finite set of disjoint simple closed curves exists.

The equivalence of (1) and (2) follows from standard arguments and the fact that if N is cylindrical, then there exists and embedded cylinder. That (3) implies (1) is clear. If a cylinder exists, any gluing map h which identifies the ends of the cylinder yields a torus in N/h. That (4) implies (1) may be proved as follows. Pinching the ends of the cylinder (we may assume they are disjoint simple curves) yields a sequence of deformations whose limit does not exist. This is well-known. That (1) implies (3) is part of Thurston's geometrization theorem; it proceeds by first proving (1) implies (4) [Thu1]. An alternative proof which does not take this route may be found in [McM3]. That a cylinder in a geometrically finite hyperbolic three-manifold causes closures of the domain of discontinuity to intersect follows easily by considering the lifts of geodesics representing ends of the cylinder to the Riemann sphere under the projection map from the universal cover. That (1) implies (5) seems to be well-known, but I have been unable to locate the original reference. It is sometimes attributed to Maskit.

One aspect of the utility of incompressible surfaces stems from the fact that, after cutting M along S, the resulting pieces have homotopy-theoretic properties which are highly representative of those of M: the fundamental group $\pi_1(M)$ splits (as free product with amalgamation or HNN extension) over $\pi_1(S)$. Another is the following theorem, due to Jaco and Shalen and, independently, Johannson [JS]:

Theorem 1.9 (Torus decomposition). *Let M be a closed, orientable, irreducible three-manifold. Then up to isotopy, there is a (possibly empty)* canonical *family of incompressible tori $\mathcal{T}$ such that each component of $M - \mathcal{T}$ is either atoroidal or Seifert fibered.*

With the same notation, Thurston conjectured the following, which includes as a special case the Poincaré conjecture:

Conjecture: (Geometrization conjecture) *The components of $M - \mathcal{T}$ admit exactly one of eight possible geometric structures.*

We propose to view a Thurston map $F : S^2 \to S^2$ as the analog of a compact, orientable, irreducible three-manifold and an invariant multicurve (Definition 1.13) as the analog of a nonperipheral, incompressible surface disjoint from ∂M. With this in mind, our Canonical Decomposition Theorem (Theorem 10.2) can be viewed as an analog of both the Torus Decomposition Theorem and the Geometrization Conjecture for three-manifolds; the canonical nature of the Thurston obstruction follows from the Canonical Obstructions Theorem (Thm. 10.4), a generalization of Theorem 1.3.

The mapping class group of a three-manifold is its group of self-homeomorphisms, modulo those isotopic to the identity. Johannson [Joh] proved

Theorem 1.10 (Finiteness of mapping class groups). *If M is a compact, orientable, irreducible, Haken, acylindrical, atoroidal manifold, then the mapping class group of M is finite.*

Moreover, the conclusion can fail if e.g. the hypothesis of atoroidal is dropped. Our Twist Theorem (Thm. 8.2) asserts that the mapping class group (Definition 8.1) of a branched mapping is infinite if there is an invariant multicurve of a certain kind, and is therefore a partial analog of the converse of Johannson's theorem.

We summarize this analogy in the following table.

Manifolds	**Branched maps**
Cpt., oriented, irreducible 3-mfd. M	Branched map $F : S^2 \to S^2$
Nonperipheral (incompressible) surface	Invariant multicurve
Nonperipheral incompressible torus	Thurston obstruction
Canonical decomposing tori	Canonical obstruction (§10)
Mapping class group	Mapping class group (§8)
Haken	Having an invariant multicurve
Gluing along boundary components	Combination Thm. (§§3, 4)
Cutting along surfaces	Decomposition Theorem (§§5, 6)
Torus Decomposition Thm.	Canonical Decomposition Thm. (§10)

We remark that there is an essential difference between the two sides: for manifolds/Kleinian groups the *Klein-Maskit combinations* (see [Mar], [Mas]) give geometric realizations in the setting of Kleinian groups for the topological operation of gluing along surfaces. For Thurston maps such geometric realizations, in which one finds quasiconformally distorted copies of the original groups which are "glued", do not usually exist. Figure 1.1 makes this clear– there is no qc embedded copy of the Julia set of $z^2 - 1$ in the Julia set of the mating.

1.7 Connections

We mention here briefly some connections between the dynamics of postcritically finite rational maps and other areas of mathematics.

1.7.1 Geometric Galois theory

Recently there has been an attempt to gain an understanding of the structure of the *absolute Galois group* $\mathcal{G} = \mathrm{Gal}(\overline{\mathbb{Q}}/\mathbb{Q})$ by exploiting the remarkable fact that there is a faithful action of $\mathcal{G}$ on a certain infinite set of finite, planar trees, called *dessins*. These dessins are combinatorial objects which classify planar covering spaces $X \xrightarrow{f} \mathbb{C} - \{0, 1\}$ given by polynomial maps f unramified above $\{0, 1\}$. The action of $\mathcal{G}$ on the set of dessins is obtained by letting $\mathcal{G}$ act on the coefficients of f, which one may take to be algebraic.

The main result of [Pil1] is that there is also a faithful action of $\mathcal{G}$ on the infinite set of *Hubbard trees.* Recall from §1.5 that Hubbard trees are finite planar trees equipped with self-maps which classify *postcritically finite* polynomials $f : \mathbb{C} \to \mathbb{C}$ as dynamical systems. Again, one may take the coefficients of such a map to be algebraic, and the action of $\mathcal{G}$ is obtained by letting $\mathcal{G}$ act on the coefficients of f. In fact, it is proved that $\mathcal{G}$ acts faithfully on a highly restricted subset DBP ("dynamical Belyi polynomials") consisting of postcritically finite polynomials f whose iterates are all unramified over $\{0, 1\}$ and whose Hubbard tree is uniquely determined by the dessin associated to f as a covering space, plus a small amount of additional data.

There are several intriguing aspects to this dynamical point of view. First, it turns out that the natural class of objects with which to work consists of actual polynomials as opposed to equivalence classes of polynomials. Second, the dynamical theory is richer. In particular, a special class of dynamical Belyi polynomials is introduced, called *extra-clean* DBPs, which is closed under composition, hence under iteration. This allows one to associate a *tower* of invariants to a single given polynomial f, namely the monodromy groups $\mathrm{Mon}(f^{\circ n})$ of its iterates. Finally, the dynamical theory here embeds into the non-dynamical one in the following sense: there is a $\mathcal{G}$-equivariant injection of the set of extra-clean *dynamical* Belyi polynomials into the set of *non-dynamical* isomorphism classes of Belyi polynomials given by $f \mapsto f^{\circ 2}$. From the point of view of dynamics, this is remarkable: the dynamics of such an f, which involves an identification of domain and range, is completely determined by the isomorphism class of $f^{\circ 2}$ as a *covering space*, which does not require such an identification.

1.7.2 Gromov hyperbolic spaces and interesting groups

Let (U_n, u_n), $n = 0, 1, 2, \ldots$ be a sequence of path-connected, pointed topological spaces, and let $f_n : (U_n, u_n) \to (U_{n-1}, u_{n-1})$, $n = 1, 2, 3, \ldots$ be covering spaces which are unramified, at least two- but finite-sheeted, and not necessarily regular. The composition $f^n = f_1 \circ f_2 \circ \ldots \circ f_n : (U_n, u_0) \to (U_0, u_0)$ is an unramified covering. Let the fiber of f^n over u_0 be denoted $f^{-n}(u_0)$.

Let $G = \pi_1(U_0, u_0)$ denote the fundamental group of U_0 based at u_0. By path-lifting, for each n, there is a transitive right action of G on the inverse image $f^{-n}(u_0)$ of u_0 under f^n. The quotient of G by the kernel of this action is the *monodromy group* of f^n, denoted $\mathrm{Mon}(f^n)$. If $v_n \in f^{-n}(u_0)$ and $g \in G$, then clearly $(f_n(v_n))^g = f_n(v_n^g)$. Hence for each $n \geq 2$ the groups $\mathrm{Mon}(f^n)$ are imprimitive, and there is a surjective homomorphism $\mathrm{Mon}(f^n) \to \mathrm{Mon}(f^{n-1})$. The inverse limit $\mathrm{Mon}(f^1) \leftarrow \mathrm{Mon}(f^2) \leftarrow \mathrm{Mon}(f^3) \leftarrow \ldots$ is thus a profinite group; its isomorphism type is independent of the choices of basepoints $\{u_n\}$.

If $U_n = F^{-n}(S^2 - P_F)$ where F is a Thurston map, the result is what is termed in [BGN] the *iterated monodromy group (IMG)* of F. According to [BGN], the IMG of $z^2 + i$ has intermediate growth, i.e. with respect to some generating set, the number of group elements expressible as a word of length

n in the generators grows faster than any polynomial but slower than any exponential function. IMGs are shown to be examples of what are termed "self-similar groups", and when F is expanding with respect to some metric (e.g. when F is rational) then the corresponding IMG satisfies an additional property of being what is termed "contracting". In this case, a general mechanism for identifying the Julia set as the boundary at infinity of a certain Gromov hyperbolic graph is given, and in a manner which permits one to, in principle, algorithmically reconstruct the dynamics of F from purely combinatorial information.

1.7.3 Cannon's conjecture

The investigations of Cannon, Floyd, and Parry [CFP1], [CFP2], [CFP3] into connections between subdivision rules and postcritically finite rational maps are motivated in part by a desire to prove

Cannon's conjecture *Let G be a Gromov hyperbolic group whose boundary at infinity is homeomorphic to the two-sphere. Then G acts discretely, cocompactly, and isometrically on hyperbolic three-space.*

If true, this would be one step toward verification of the Geometrization Conjecture. The idea, very roughly, is as follows. Successively growing the ball of radius n in the Cayley graph starting at the identity yields a sequence of finer and finer "subdivisions" of the boundary at infinity. Cannon [Can] gives conditions for the existence of a G-invariant conformal structure on the boundary at infinity. The existence of such a structure then implies the conjecture for G. In actuality the situation is somewhat more complicated, and verifying these conditions is difficult. However, in many classes of examples having the flavor given in §1.3.4, one can in fact verify these conditions and conclude that the branched covering is indeed equivalent to a rational map. It is hoped that the converse is possible as well, i.e. rationality of the Thurston map induced from the subdivision rule should imply that Cannon's conditions for conformality hold. For more details, see [CFP3] and recent work of M. Bonk and B. Kleiner.

1.8 Discussion of combinatorial subtleties

In this subsection, we first recall for reference the concepts and definitions arising in Thurston's characterization of rational functions. Our results require the introduction of some significant generalizations of these concepts, as well as some slight modifications to these standard definitions. In the following subsections, we motivate these generalizations by informally sketching the process of decomposition, and introduce terminology and notation which will be used in the remainder of this work.

1.8.1 Overview of decomposition and combination

The development of a general theory of combinations and decompositions necessitates the simultaneous introduction of three levels of generalizations of the category whose objects are postcritically finite branched coverings of the two-sphere and whose morphisms are pairs h_0, h_1 of homeomorphisms yielding combinatorial equivalences. Specifically:

- Instead of a Thurston map $F : S^2 \to S^2$, a map of a single sphere to itself, we consider $\mathcal{F}^{\mathcal{S}} : \mathcal{S} \to \mathcal{S}$, a postcritically finite branched covering map of a finite set $\mathcal{S}$ of spheres to itself, not necessarily surjective. For such maps $\mathcal{F}$, critical points and the postcritical set $P_{\mathcal{F}}$ are defined in the same manner as for a single map.
- Instead of using P_F, the postcritical set, in the definitions of combinatorial equivalence and multicurves, we use an *embellishment* $\mathcal{Y} \subset \mathcal{S}$. Here, $\mathcal{Y}$ will be a closed, forward-invariant set containing $\mathcal{F}^{-1}(P_{\mathcal{F}})$; see the next section.
- Instead of considering only postcritically finite maps, we consider non-postcritically finite maps having some tameness properties near accumulation points of the postcritical set.

The last generalization is needed in order to give a decomposition theory for e.g. rational functions with disconnected Julia sets, since Julia sets of postcritically finite rational maps are necessarily connected.

Before a more detailed discussion of the first two generalizations, let us see informally how they naturally arise in decompositions. In §5 we will make the decomposition process precise. Let F be a branched covering with $Q = F^{-1}(P_F)$. Let Γ be a finite invariant multicurve in $S^2 - Q$. Cut the sphere apart along a set of annuli $\mathcal{A}_0$ in $S^2 - Q$ whose core curves are the elements of Γ. For the moment, discard these annuli $\mathcal{A}_0$. For each of the remaining pieces, cap the resulting holes by disks with preferred center points to obtain a collection $\mathcal{S}$ of spheres. Denote the collection of resulting center points of these disks by $\mathcal{Z}$. Then $\mathcal{Z}$ is finite. The set Q yields a corresponding set in this union of spheres which we denote by $\mathcal{Q}$. Set $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$. After a suitable extension, one finds that the original map $F : (S^2, Q) \to (S^2, Q)$ yields a map $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ from a finite set of spheres to itself such that $\mathcal{Y}$ and $\mathcal{Z}$ are forward-invariant under $\mathcal{F}^{\mathcal{S}}$. One subtlety is that $\mathcal{Q}$ need no longer be forward-invariant under $\mathcal{F}^{\mathcal{S}}$. Another subtlety is that it is possible for $\mathcal{Z}$ to be disjoint from the orbits of critical points, e.g. if elements of Γ all map by degree one.

Also, we wish to be able to reconstruct the original map F. Therefore, when formulating a combination procedure, we should expect that we must record the following kinds of data:

- **(Mapping tree, §2.1)** a combinatorial object capturing how the annuli in $\mathcal{A}_0$ are deployed on the sphere, as well as some rudimentary dynamics;

- **(Sphere Maps, §2.2)** a map of pairs $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$, where $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$;
- **(Annulus Maps, §2.3)** a map $\mathcal{F}^{\mathcal{A}} : \mathcal{A}_1 \to \mathcal{A}_0$, recording the dynamics above $\mathcal{A}_0$, which was forgotten in the process of decomposition sketched above, (here, $\mathcal{A}_1 \subset \mathcal{A}_0$);
- **(Topological Gluing, §3.1)** a choice of map gluing the annular pieces and spherical pieces together;
- **(Critical Gluing, §3.2)** a choice of what to do with points in $\mathcal{Q}$ which map to points in $\mathcal{Z}$ under $\mathcal{F}^{\mathcal{S}}$ (see the examples in §9.2);
- **(Missing Pieces, §3.3)** a choice of how to define the new map $F : S^2 \to S^2$ on those regions not accounted for by $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}$. The noninvertible nature of these dynamical systems implies that in all but the simplest cases (matings; see §§1.3 and 9.2) this is a potential source of nontrivial ambiguity in the definition. We will resolve this point in §4 by showing that, provided a suitable normal form is used, the choices made here have no influence on the outcome.

There are several obvious and a few non-obvious compatibility requirements which must be satisfied. In §§2 and 3 we adopt an axiomatic approach to enumerating these requirements which will guarantee that the results of combination and decomposition will be well-defined, and that these processes are inverse to one another.

1.8.2 Embellishments. Technically convenient assumption.

Technically, it is far more convenient to reformulate the set-theoretic conditions in the definition of combinatorial equivalence and multicurves using the *full preimage* of the postcritical set as opposed to the postcritical set itself. One reason is as follows:

Lemma 1.11. *Let γ be a simple closed curve in $S^2 - Q$, where $Q = F^{-1}(P_F)$. If γ is essential in $S^2 - Q$, then no two components of $F^{-1}(\gamma)$ are homotopic in $S^2 - Q$.*

Proof: Suppose the contrary. Then there exist two preimages $\widetilde{\gamma}_1, \widetilde{\gamma}_2$ of γ such that $\widetilde{\gamma}_1 \sqcup \widetilde{\gamma}_2$ is the boundary of an annulus R in $S^2 - F^{-1}(\gamma)$ with $R \cap Q = \emptyset$. The restriction of F to R is a proper unramified covering since Q contains the critical points and by construction R is a component of the preimage of the complement of γ. Hence the image $F(R)$ is a nondegenerate annulus, which is impossible. $\Box$As another consequence, if Γ_0 is a multicurve in $S^2 - Q$, then it can be shown that the annuli in $\mathcal{A}_0$ map by unramified coverings under F. This is convenient, since then in our “mapping tree” caricature of the dynamics, any folding caused by critical points will be concentrated in vertices of this tree. See also Lemma 5.6.

Thus, in general, it is convenient to reformulate the definitions of multicurve and combinatorial equivalence so as to replace the postcritical set P_F

by its full preimage Q_F. However, we have already seen in the previous subsection how the process of decomposition leads naturally to consideration of maps $\mathcal{F}^{\mathcal{S}}$ defined on a set $\mathcal{S}$ of spheres. With this in mind, we formulate

Definition 1.12 (Embellished map of spheres). *An* embellished map of spheres *is a triple* $(\mathcal{F}, \mathcal{S}, \mathcal{Y})$ *where*

- $\mathcal{S}$ *is a finite set of spheres,*
- $\mathcal{Y} \subset \mathcal{S}$ *is closed,*
- $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ *is a (tame) map of pairs such that*
  - $\mathcal{F}$ *is an orientation-preserving branched covering,*
  - $\mathrm{Crit}(\mathcal{F}) \subset \mathcal{Y}$*, where* $\mathrm{Crit}(\mathcal{F})$ *denotes the set of critical points of* $\mathcal{F}$*,*
  - $\mathcal{F}^{-1}(\mathcal{F}(\mathcal{Y})) = \mathcal{Y}$*.*

Two embellished maps $(\mathcal{F}, \mathcal{S}, \mathcal{Y}), (\mathcal{G}, \mathcal{S}', \mathcal{Y}')$ *are called* combinatorially equivalent *if there are orientation-preserving (tame) homeomorphisms of pairs* $\mathcal{H}_0, \mathcal{H}_1 : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}', \mathcal{Y}')$ *such that* $\mathcal{H}_0 \circ \mathcal{F} = \mathcal{G} \circ \mathcal{H}_1$ *and* $\mathcal{H}_0$ *is isotopic to* $\mathcal{H}_1$ *through (tame) homeomorphisms agreeing on* $\mathcal{Y}$*.*

Usually, we will drop the adjective “embellished”.
Remarks:

1. When $\mathcal{S}$ has just one connected component we will usually use Roman fonts F, G, F', F'', Q etc. Also, we will sometimes refer to an embellished map of spheres by e.g. the single symbol $\mathcal{F}$, where it is to be understood that there is a distinguished subset $\mathcal{Y} \subset \mathcal{S}$ associated to $\mathcal{F}$.
2. Here, “tame” is an unspecified set of regularity conditions on the geometry of the sets $\mathcal{Y}, \mathcal{Y}'$, the regularity of the maps $\mathcal{F}$, $\mathcal{G}$, $\mathcal{H}_0$, $\mathcal{H}_1$, and the regularity of the isotopy joining $\mathcal{H}_0$ to $\mathcal{H}_1$. This is discussed in more detail in §1.6.
3. We allow of course the case when $\mathcal{S}$ consists of a single sphere.
4. Kameyama [Kam2] uses the term *furnished* in the case when a single map is involved and $\mathcal{Y}$ is finite. In our setting, when combining maps, the set $\mathcal{Y}$ will be a disjoint union $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$, where $\mathcal{Z}$ is finite and forward-invariant.

1.8.3 Invariant multicurves for embellished map of spheres. Thurston linear map.

A *multicurve* in $\mathcal{S} - \mathcal{Y}$ is a finite collection of simple, closed, essential, nonperipheral curves in $\mathcal{S} - \mathcal{Y}$, no two of which are homotopic in $\mathcal{S} - \mathcal{Y}$. We shall need to require a slightly stronger definition of invariance, however.

Definition 1.13 (Invariant multicurve). *Let* $(\mathcal{F}, \mathcal{S}, \mathcal{Y})$ *be an embellished map of spheres and* Γ *a multicurve in* $\mathcal{S} - \mathcal{Y}$*.* Γ *is* $\mathcal{F}$-invariant *if*

1. *for all* $\gamma \in \Gamma$*, each component of* $\mathcal{F}^{-1}(\gamma)$ *is either peripheral with respect to* $\mathcal{Y}$*, or homotopic in* $\mathcal{S} - \mathcal{Y}$ *to an element of* Γ*, and*

2. *for each $\gamma_i \in \Gamma$, there exists a $\gamma_j \in \Gamma$ such that γ_i is homotopic in $\mathcal{S} - \mathcal{Y}$ to a component of $\mathcal{F}^{-1}(\gamma_j)$.*

The first condition is the same as the usual one; the second is equivalent to the following. Associated to an $\mathcal{F}$-invariant multicurve $\Gamma = \{\gamma_1, ..., \gamma_N\}$ is the *Thurston linear map* $\mathcal{F}_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ defined as in §1.5.1. The second invariance condition is equivalent to the condition that the matrix $(\mathcal{F}_\Gamma)$ has a nonzero entry in each row.

Remark: It is not always the case that a nonempty multicurve satisfying the first condition contains a proper subset which is a sub-multicurve satisfying both conditions–*a priori* it is possible for a multicurve to consist of a collection of curves, each of whose preimages are inessential or peripheral. This happens if and only if the matrix for the Thurston linear map is the zero matrix. However, in all other situations, deleting those rows and columns corresponding to zero rows will result in a multicurve verifying both conditions.

1.9 Tameness assumptions

Bearing in mind the generalizations discussed above, let $(\mathcal{F}, \mathcal{S}, \mathcal{Y})$ be an embellished map of spheres. We might wish to require some tameness assumptions on $(\mathcal{F}, \mathcal{S}, \mathcal{Y})$. In order for these properties to be natural, one must require some additional regularity on

1. the set $\mathcal{Y}$ itself,
2. the map $\mathcal{F}$ near $\mathcal{Y}$, and
3. the maps $\mathcal{H}_0, \mathcal{H}_1$ occurring in the definition of combinatorial equivalence,
4. the isotopy joining $\mathcal{H}_0$ and $\mathcal{H}_1$.

A typical restriction in (1) is that $\mathcal{Y}$ have only finitely many accumulation points. This will be the case if the map $\mathcal{F}$ is e.g. a single *geometrically finite* rational map, i.e. one for which the intersection of the Julia set and postcritical set is finite.

This requirement alone is usually too weak to be used for reasonable topological characterizations of rational functions with infinite postcritical set. For instance, if the maps $\mathcal{H}_0, \mathcal{H}_1$ in the definition of combinatorial equivalence are required only to be homeomorphisms, then the map $\mathcal{F}(z) = z^2 - \frac{1}{10}$ having an attracting fixed point is combinatorially equivalent to the map $\mathcal{G}(z) = z^2 - \frac{3}{4}$ having a parabolic fixed point. Thus, the topological model may fail to distinguish between maps which are geometrically essentially different, and which indeed have non-homeomorphic Julia sets.

Geometrically finite rational maps are also well-behaved near accumulation points of their postcritical set. In contrast, one might e.g. postcompose such a map by an infinite sequence of Dehn twists along curves converging to an attractor to obtain a new branched map with pathological behavior near this attractor. It might therefore be reasonable to pin down the local behavior of the map near accumulation points of $\mathcal{Y}$.

One possible way around this difficulty is to require that near such points, (2) the map $\mathcal{F}$ is locally C^1 conjugate to a rational function, (3) the maps $\mathcal{H}_0, \mathcal{H}_1$ are C^1, and (4) the isotopy joining $\mathcal{H}_0$ to $\mathcal{H}_1$ take place through C^1 maps. This will serve to distinguish between the maps $\mathcal{F}$ and $\mathcal{G}$ above, since the geometry of the two postcritical sets are radically different near the attractor. Alternatively, one could replace C^1 with quasiconformal. One could also strengthen the requirement in (2) to holomorphic conjugacy, but otherwise keep the remaining requirements unmodified. For other alternatives, see e.g. Cui [CJS], Epstein-Keen-Tresser [EKT], and McMullen ([McM7], Appendix A).

In general, then, when the postcritical set is infinite, there may be various choices of regularity on the maps h_0, h_1 near the points of Y which one might wish to require.

To keep our theory of general applicability we leave the details unspecified, referring to whatever restrictions are imposed in (1)-(4) above as *tameness assumptions.* Our combination and decomposition constructions, however, are robust in the following sense. When combining, the new dynamical system is produced by C^0 modifications of the original maps away from the set $\mathcal{Q}$, which contains all accumulation points of $\mathcal{Y}$ since $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$ is closed and $\mathcal{Z}$ is finite. The Uniqueness of Combinations Theorem 4.5 (see §4) asserts roughly that the combinatorial class of the map resulting from a combination depends only on the classes of the input data (suitably defined). The proof of this theorem relies on two steps: gluing together maps yielding equivalences of input data, and C^0 modifications away from the sets $\mathcal{Q}$ and Q. Thus, when combining, we may safely assert that there will be no loss of regularity of equivalences near points of Q. Similarly, when decomposing a branched covering $F : (S^2, Q) \to (S^2, Q)$, we may assume that the resulting embellished map of spheres $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$, where $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$, has the same tameness properties near $\mathcal{Q}$ as does F near Q.

More precisely, we need that whatever tameness definitions are used, the following results should hold:

Tameness Restrictions

1. **Space of tame homeomorphisms is path connected.** If $h : \mathcal{S} \to \mathcal{S}$ is a tame homeomorphism which fixes each component of $\mathcal{S}$, then there is a continuous path $h^t, 0 \leq t \leq 1$, of tame homeomorphisms with $h^0 = \mathrm{id}$ and $h^1 = h$.
2. **Equivalence of definitions of combinatorial equivalence.** Suppose $\mathcal{F}^t : (\mathcal{S}, \mathcal{Y}^t) \to (\mathcal{S}, \mathcal{Y}^t), \ \ 0 \leq t \leq 1$ is a continuous, one-parameter family of tame embellished maps of spheres, such that $\mathcal{Y}^t = h^t(\mathcal{Y}^0)$, where h^t is a path of tame homeomorphisms. We require that $\mathcal{F}^0$ *be combinatorially equivalent to* $\mathcal{F}^1$.
3. **Equivalence of subsurfaces.** If $\mathcal{V} \subset \mathcal{S}$ is a compact subsurface satisfying

a) $\mathcal{V}$ has finite connectivity, i.e. has finitely many connected components, each with finitely many boundary components,

b) ∂V consists of finitely many tame Jordan curves in $\mathcal{S} - \mathcal{Y}$

then we say $\mathcal{V}$ is *admissible.* We say admissible $\mathcal{V}, \mathcal{V}'$ are *ambiently homeomorphic* if there is a homeomorphism $h : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ which is the identity on $\mathcal{Y}$ and sending $\mathcal{V}$ to $\mathcal{V}'$.

The requirement is the following: *If $\mathcal{V}, \mathcal{V}'$ are ambiently homeomorphic, then the homeomorphism h can be taken to be tame and in the same isotopy class as h rel $\mathcal{Y}$.* Said loosely another way, the equivalence relation of "ambiently homeomorphic" on the set of admissible subsurfaces should be the same in the tame and topological categories.

Remarks:

1. The converse implication in (2) holds trivially using (1), since a pair of tame homeomorphisms yielding a combinatorial equivalence between $\mathcal{F}$ and $\mathcal{F}'$ defines a path between $\mathcal{F}$ and a conjugate of $\mathcal{F}'$, which in turn is joined to $\mathcal{F}'$ by a path defined by an isotopy joining the conjugating map to the identity.
2. When $\mathcal{Y}$ is finite and $\mathcal{S} = S^2$, (1) follows from ([Bir], Thm. 4.4). That (2) holds in this case was first mentioned explicitly and used by M. Rees [Ree3].

2

Preliminaries

In this section, we formulate axiomatically the combinatorial and dynamical data needed to define our combination and decomposition operations. A central aim is to allow a complicated map to be decomposed into simpler pieces in such a way that the original map is recoverable from the simpler pieces and certain gluing data. In order to achieve this, when defining our combination procedure, we simultaneously consider the pieces and the gluing data throughout. To ensure realizability of the data by an embellished branched covering of the sphere to itself, we prescribe that the data satisfy a lengthy but rather natural set of axioms guaranteeing compatibility of combinatorics, dynamics, and topology.

We begin in Section 2.1 with the combinatorial aspects. A *mapping tree* is a caricature of a single embellished branched covering [1] $F : (S^2, Q) \to (S^2, Q)$ with an invariant multicurve Γ, obtained in the following way. Full details are contained in Section 5, where we prove the Decomposition Theorem [5.1]. We emphasize that the reader pay particular attention to the notation developed below.

- Thicken the elements of the multicurve Γ to form a collection of disjoint, essential, nonparallel closed annuli in $S^2 - Q$ whose union we denote by $\mathcal{A}_0$. Thus, $\mathcal{A}_0$ is a subset of the sphere. A component of $\mathcal{A}_0$ is denoted by $\mathcal{A}_{a_{0,i}}$ and the set of such components is denoted by A_0. Thus,

$$\mathcal{A}_0 = \bigcup_{a_{0,i} \in A_0} \mathcal{A}_{a_{0,i}}.$$

- We arrange by precomposing with a homeomorphism isotopic to the identity so that the map $\mathcal{F}$ is in "standard form" (see Definition 5.2).

[1] Here, the triple (F, S^2, Q) is required only to be an embellished branched covering, i.e. Q is closed, forward-invariant, contains the critical points, and coincides with the preimage of its image.

K.M. Pilgrim: LNM 1827, pp. 37–48, 2003.

- Let $\mathcal{A}_1$ be the union of preimages of elements of $\mathcal{A}_0$ which are essential in $\mathcal{A}_0$. Note that each component of $\mathcal{A}_1$ is an annulus mapping under F as an unramified covering onto its image. Also, $\mathcal{A}_1$ is a subset of the sphere. A component of $\mathcal{A}_1$ is denoted by $\mathcal{A}_{a_{1,j}}$ and the set of such components is denoted by A_1. Thus,
$$\mathcal{A}_1 = \bigcup_{a_{1,j} \in A_1} \mathcal{A}_{a_{1,j}}.$$
The standard form implies that $\mathcal{A}_1 \subset \mathcal{A}_0$ and that $\partial \mathcal{A}_1 \supset \partial \mathcal{A}_0$.
- Let $\mathcal{S}_0$ denote the union of the set of connected components of $\overline{S^2 - \mathcal{A}_0}$. Thus, $\mathcal{S}_0$ is a subset of the sphere. A component of $\mathcal{S}_0$ is denoted $\mathcal{S}_0(x)$ and the set of such components is denoted by X. Thus,
$$\mathcal{S}_0 = \bigcup_{x \in X} \mathcal{S}_0(x).$$
- Within each component $\mathcal{S}_0(x)$ of $\mathcal{S}_0$, it turns out that there is a unique component $\mathcal{S}_1(x)$ of $\mathcal{F}^{-1}(\mathcal{S}_0)$ whose boundary contains $\partial \mathcal{S}_0(x)$. We let $\mathcal{S}_1$ denote the union of such components. Thus,
$$\mathcal{S}_1 = \bigcup_{x \in X} \mathcal{S}_1(x).$$
- Let $\mathcal{C} = \overline{\mathcal{A}_0 - \mathcal{A}_1}$. Thus $\mathcal{C}$ is a subset of the sphere contained in $\mathcal{A}_0$. A component of $\mathcal{C}$ is denoted by $\mathcal{C}_c$ and the set of such components is denoted by C. Thus
$$\mathcal{C} = \bigcup_{c \in C} \mathcal{C}_c.$$
Then each annulus in $\mathcal{A}_0$ is an alternating union of annuli in $\mathcal{A}_1$ and in $\mathcal{C}$, beginning and ending with annuli in $\mathcal{A}_1$. Possibly there are no components of $\mathcal{C}$ in a given component of $\mathcal{A}_0$. Figure 2.2 shows a caricature of how these subannuli are deployed.
- We set $\mathcal{U}$ to be the union of the set of connected components of $\overline{\mathcal{S}_0(x) - \mathcal{S}_1(x)}$, as x ranges over X. A component of $\mathcal{U}$ is denoted by $\mathcal{U}_u$ and the set of such components is denoted by U. Note that the components of $\mathcal{U}$ are disks. See Figure 2.1. Thus,
$$\mathcal{U} = \bigcup_{u \in U} \mathcal{U}_u.$$
- Finally, we let B_1 denote the set of boundary components of $\mathcal{A}_1$. Let B_0 denote the subset of B_1 consisting of boundary components of $\mathcal{A}_0$. Then $B_0 \neq \emptyset$. Note that $\partial \mathcal{A}_0 = \partial \mathcal{S}_0$, since $S^2 = \mathcal{S}_0 \cup \mathcal{A}_0$, and the union is along boundary components.

The mapping tree T is then the following graph. The set of nodes (vertices) is the set $X \sqcup B_1 \sqcup U$. The edges are as follows. A node $u \in U$ is connected to exactly one other node; this node is an x-node, and is the unique x such that $\mathcal{U}_u$ and $\mathcal{S}_1(x)$ share a boundary component. A node in X is adjacent only to nodes in U and to nodes in B_1, and the latter case occurs only when the boundary component corresponding to b_1 is a boundary component of $\mathcal{S}_0(x)$. The remaining edges join pairs b^+, b^- of elements of B_1 bounding a common annulus, which may be either in $\mathcal{A}_1$ or in $\mathcal{C}$.

The mapping tree T comes equipped with a projection π from S^2 to T which collapses

- boundary components of $\mathcal{A}_1$ to B_1,
- annuli in $\mathcal{A}_1$ and $\mathcal{C}$ to edges,
- components of $\mathcal{U}$ to vertices in U,
- subsurfaces of the form $\mathcal{S}_1(x)$ to the subtree spanned by x and all vertices (necessarily in $B_1 \sqcup U$) which are adjacent to x.

At this point the reader is urged to turn to the examples in Section 9.

Note that the definitions of the tree and projection map lead to some inequity within the definitions of nodes and edges. For instance:

- $\pi^{-1}(u), u \in U$ is a disk $\mathcal{U}_u$
- $\pi^{-1}(b), b \in B_1$ is a Jordan curve
- $\pi^{-1}(x)$ is *neither* $\mathcal{S}_0(x)$ nor $\mathcal{S}_1(x)$
- the preimage of an edge joining two vertices in B_1 is an annulus in either $\mathcal{A}$ or $\mathcal{C}$.

Despite this apparent asymmetry, the graph-theoretic "closed neighborhood" of x in T, which we call the "star" of the node x and denote by $S_1(x)$, is a convenient caricature of the surface $\mathcal{S}_1(x)$; see Figure 2.1.

As we shall see, the mapping tree T comes also equipped with a self-map $f : T \to T$ which can be made to respect the projection π in a certain combinatorial fashion. There is also a natural degree function $\deg : \{\text{nodes of } T\} \to \mathbb{N}$ given by the topological degree of F on a suitable subset of the sphere.

Finally, imagine cutting apart the sphere using boundary components of $\mathcal{A}_0$, and "capping" the resulting boundary components of $\mathcal{S}_0$ with disks to form a collection $\mathcal{S}$ of spheres. Then above the mapping tree lie two kinds of combinatorial dynamical systems: one corresponding to the spheres $\mathcal{S}$, and one to the annuli $\mathcal{A}_1, \mathcal{A}_0$.

In the next sections, we give a rather long list of axioms satisfied by this setup in order to formulate precisely our combination theorems.

2.1 Mapping trees

A *mapping tree of degree d* is a triple $\mathbf{T} = (T, f, \deg)$, where

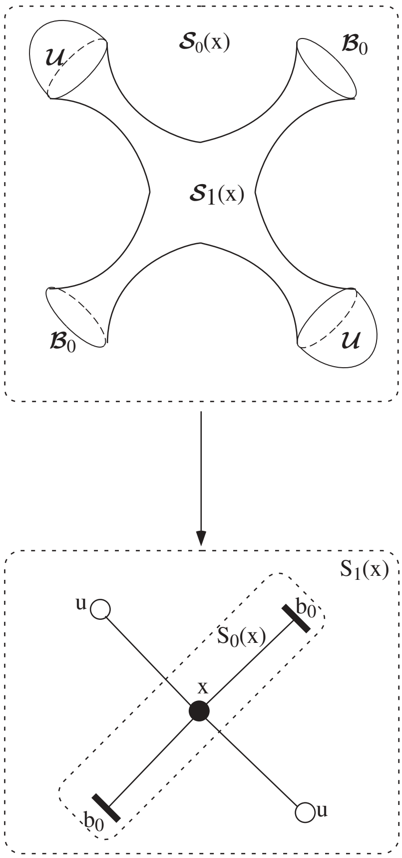

Fig. 2.1. The star $S_0(x)$ is a caricature of $\mathcal{S}_0(x)$.

- T is a finite, connected tree,
- the set of vertices of T is a disjoint union

$$X \sqcup B_1 \sqcup U,$$

- $f : T \to T$ is a continuous self-map
- $\deg : X \sqcup B_1 \sqcup U \to \mathbb{N}$

subject to the following axioms:

A–1 **Valence.**

1. U vertices have valence one.
2. B_1 vertices have valence two.

A–2 **Adjacency.**

1. U vertices are adjacent only to X vertices.

2. B_1 vertices are adjacent either to two vertices in B_1, or to one vertex in B_1 and one vertex in X.
3. X vertices are adjacent only to B_1 and to U vertices.
   We denote by B_0 the nonempty subset of B_1 consisting of those vertices adjacent to vertices in X.

   Definition (annular edges, $\mathbf{A_0, A_1}$): Given $b_1 \in B_1$, the Valence and Adjacency axioms imply that there is a unique maximal path a_0 in $T - X$ which contains b_1 and has the property that $\partial a_0 \subset B_1$. Then actually $\partial a_0 \subset B_0$. We denote the set of such maximal paths by $A_0 = \{a_{0,i}\}_{i=1}^{n_0}$ and refer to the $a_{0,i}$ as *annular edges* (even though they are not necessarily edges of T, but rather unions of edges). Note that the Valence and Adjacency axioms imply that the vertices within any annular edge lie in B_1. We require
4. Every annular edge has an even number of B_1 vertices.

This implies that for each annular edge $a_{0,i} \in A_0$, as $a_{0,i}$ is traversed from one boundary vertex $b_0^- \in B_0$ to the other boundary vertex $b_0^+ \in B_0$, the edges can be consecutively labelled by an alternating sequence of symbols $A_1, C, A_1, C, ..., A_1, C, A_1$, beginning and ending with A_1. Possibly $a_{0,i}$ is itself already an edge of T, in which case the label is simply A_1. See Figure 2.2.

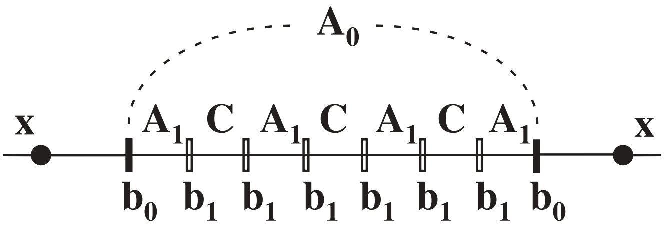

Fig. 2.2. Structure of an annular edge $a_0 \in A_0$.

We denote by A_1 and C respectively the collection of all subedges of annular edges labelled by A_1 and by C. Elements of A_1 will be denoted by $a_{1,j}$; those of C by c. Note that $\partial A_1 \supset \partial A_0$.

Stars. Given $x \in X$ we denote by $B(x)$ the set of all elements of B_1 adjacent to x; these are necessarily in B_0. Similarly, we denote by $U(x)$ the set of all elements of U adjacent to x. Below, let $[O]$ denote the smallest subtree of T containing O, a set of vertices, and for a subtree T' let $\partial T'$ be the set of vertices of valence one . Set

- $S_0(x) = [\{x\} \cup B(x)]$. Then

$$\partial S_0(x) = B(x)$$

and

$$\bigcup_{x \in X} \partial S_0(x) = \partial A_0.$$

- $S_1(x) = [\{x\} \cup B(x) \cup U(x)]$. Then $\partial S_1(x) = B(x) \cup U(x)$.
- $S_0 = \bigcup_{x \in X} S_0(x), S_1 = \bigcup_{x \in X} S_1(x)$.

Thus, the star $S_0(x)$ is a caricature of a component $\mathcal{S}_0(x)$, while the star $S_1(x)$ is a caricature of $\mathcal{S}_1(x)$.

A–3 **Dynamics.**

1. $f : X \to X$. This need not be surjective onto X.
2. $f : B_1 \sqcup U \to B_0$.
3. f is a homeomorphism on each edge of T
4. $f|a_{0,i}$ is a homeomorphism on each annular edge $a_{0,i}$.
5. $f : A_1 \to A_0$, i.e. the image of each edge $a_{1,j} \in A_1$ is an annular edge $a_{0,i} \in A_0$. This need not be surjective onto A_0. However, note that every annular edge contains at least one subedge in A_1. (This is the combinatorial analog of the second condition in the definition [1.13] of invariance of a multicurve.)
6. f maps an edge joining an x-vertex and a u-vertex homeomorphically onto an edge joining an x-vertex to a b_0-vertex.

In particular, any folding of T by f must be concentrated in the X vertices.

Definition (image subtrees): Recall that $u \in U$ is adjacent to exactly one x vertex. Since $f|[x, u]$ is a homeomorphism, there is a unique subtree $f_*(u) \subset T$ which is the closure of the connected component of $T - f(u)$ which does *not* contain $f(x)$. Said another way, $f_*(u)$ is the largest connected subtree containing $f(u)$ but not $f(x)$. See Figure 2.3.

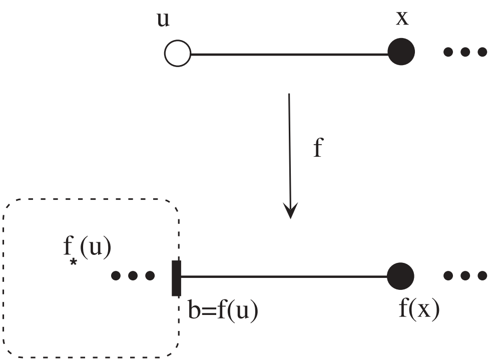

Fig. 2.3. The image subtree $f_*(u)$ is the largest subtree containing $f(u)$ but not $f(x)$, where x is the unique node adjacent to u.

Similarly, given $c \in C$, the map $f|c$ is a homeomorphism. We set $f_*(c)$ to be the closure of the connected component of $T - \{\partial f(c)\}$ which contains $f(c)$; see Figure 3.1.

Note that $f_*(u), f_*(c)$ may be quite large subtrees.

A–4 **Degree function.**

1. **Local homogeneity.**
   a) for all $y = f(x) \in X$, and all $b_0 \in B_0(y)$,

$$\sum_{v \in B(x), f(v)=b_0} \deg(v) = \deg(x).$$

   This is the combinatorial analog of the fact that if $F : \mathcal{S}_1(x) \to \mathcal{S}_0(y)$ is proper and has degree $\deg(x)$, then for each boundary component of $\mathcal{S}_0(y)$, the sum of the degrees of F on those preimages lying in $\mathcal{S}_1(x)$ must be equal to $\deg(x)$.

   b) The function deg is constant on any vertices in a given annular edge $a_{0,i}$.
   This is the combinatorial consequence of the fact (to be proved later, Lemma 5.6) that the restriction of F to the annuli in $\mathcal{A}_0$ is an unramified covering.

   Below, by $\deg(\partial c)$ we mean the value of the function deg on either boundary component of c; by Axiom A–4 1(b) above this is independent of which vertex is chosen.

2. **Global homogeneity.** For all $y \in f(X)$,

$$\sum_{f(x)=y} \deg(x) + \sum_{y \in f_*(u)} \deg(u) + \sum_{y \in f_*(c)} \deg(\partial c) = d,$$

   and for all $b_0 \in f(B_1) = B_0$,

$$\sum_{b \in B_1, f(b)=b_0} \deg(b) + \sum_{b_0 \in f_*(u)} \deg(u) + \sum_{b_0 \in f_*(c)} \deg(\partial c) = d.$$

   It may be that there is some redundancy in Axioms A–4 #1(a,b) and A–4 #2.

A–5 **Orientations.** We assume we have chosen arbitrarily a decomposition

$$B_0 = B_0^- \sqcup B_0^+$$

such that each annular subedge a_0 is bounded by a pair (b_0^-, b_0^+). Thus, we may think of annular subedges as oriented arcs $b_0^- \to b_0^+$ and coordinatize them with $I = [-1, +1] \simeq a_0$.

Definition (signed degree function δ): Recall that every edge $a_{1,j} \in A_1$ is a subedge of an oriented annular edge $a_{0,i}$. We equip $a_{1,j}$ with the induced orientation, and modify the degree function deg on the boundary of $a_{1,j}$ by setting

$$\delta(b) = \pm \deg(b)$$

with the sign chosen + if $f|a_{1,j}$ is orientation-preserving and minus otherwise. Note that δ has the same value on the two boundary vertices of each $a_{1,j} \in A_1$.

2.2 Map of spheres over a mapping tree

Let $\mathcal{S}$ be the disjoint union of $|X|$ copies of the two-sphere S^2, and let $\mathcal{S}(x)$ denote the xth component. An *embellished map of spheres lying over a mapping tree* $\mathbf{T} = (T, f, \deg)$ is an orientation-preserving branched covering

$$\mathcal{F}^{\mathcal{S}} : \mathcal{S} \to \mathcal{S}$$

together with a continuous map

$$\pi^{\mathcal{S}} : \mathcal{S} \to T$$

satisfying the following axioms.

A–6 **Embellished.** $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ is an embellished map of spheres $\mathcal{S}$ to itself for which $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$, where $\mathcal{Q}$ is closed, $\mathcal{Z}$ is finite, and $\mathcal{F}(\mathcal{Z}) \subset \mathcal{Z}$. (For the definition of embellished map of spheres, see Definition 1.12.) Note that $\mathcal{Z}$ consists of isolated points of $\mathcal{Y}$. Also, the conditions imply that the global postcritical set is well-defined. When combining, we will cut away a set of neighborhoods $\mathcal{B}$ of points in $\mathcal{Z}$ and glue in annuli in the resulting holes to obtain a new sphere S^2.

A–7 **Standard form.** There is a collection $\mathcal{B}$ of $|\mathcal{Z}|$ disjoint, closed, topological disks $\mathcal{B}_z$ such that for each z in $\mathcal{Z}$,

1. $z \in \text{int}(\mathcal{B}_z)$
2. $\mathcal{B}_z \subset \mathcal{S} - \mathcal{Q}$
3. Each component $(\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{B}_z)$ is either a component of $\mathcal{B}$, or is disjoint from $\mathcal{B}$
4. For each component $\mathcal{U}_u$ of $(\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{B}) - \mathcal{B}$, we have $|\mathcal{U}_u \cap Q| \leq 1$.
5. For each $z \in \mathcal{Z}$, there is a homeomorphism of pairs $(\mathcal{B}_z, z) \to (\overline{\Delta}, 0)$ such that $\mathcal{F}^{\mathcal{S}}|\partial\mathcal{B}_z$ in these local coordinates is given by $\zeta \to \zeta^k$ where $k \geq 1$ is the local degree of $\mathcal{F}^{\mathcal{S}}$ near z.
   Note that then we have also $k = \deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{B}_z)$. Equivalently, parameterizing $\partial\mathcal{B}_z$ by $\mathbb{R}/\mathbb{Z}$, the map on $\partial\mathcal{B}_z$ is given by $t \mapsto kt$ modulo one.
   We merely require the existence of this set of coordinates; at this point we do not make a choice of such coordinates. The effect of this is to rule out irrational rotations along cycles of boundary components in $\mathcal{B}$.

Definition (missing disks):
We set $\mathcal{U} = (\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{B}) - \mathcal{B}$. The components of $\mathcal{U}$ we will refer to as *missing disks.*

A–8 **Tameness.** We require that $\mathcal{Q}$ and $\mathcal{F}^{\mathcal{S}}$ are *tame* near any accumulation points of $\mathcal{Q}$. We leave this unspecified but require that the Tameness Restrictions of §1.9 hold. We emphasize that our results will be independent of any such choice.

A–9 **Covering.** There are some obvious compatibility requirements needed for the mapping tree to be a combinatorial caricature of the sphere maps $\mathcal{F}^{\mathcal{S}}$. For $x \in X$, let

$$\begin{aligned}
\mathcal{S}(x) &= \text{the copy of } S^2 \text{ above } x\\
\mathcal{B}(x) &= \mathcal{B} \cap \mathcal{S}(x)\\
\mathcal{U}(x) &= \mathcal{U} \cap \mathcal{S}(x)\\
\mathcal{S}_0(x) &= \overline{\mathcal{S}(x) - \mathcal{B}(x)}\\
\mathcal{S}_1(x) &= \overline{\mathcal{S}_0(x) - \mathcal{U}(x)} = \overline{\mathcal{S}(x) - \mathcal{B}(x) - \mathcal{U}(x)}.
\end{aligned}$$

1. For each $x \in X$, $\pi^{\mathcal{S}}$ is a map of triples

$$\pi^{\mathcal{S}} : (\mathcal{S}_0(x), \partial\mathcal{S}_1(x), \partial\mathcal{S}_0(x)) \to (S_0(x), \partial S_1(x), \partial S_0(x))$$

which is a bijection on connected components of each of the sets indicated. Note that $\partial S_1(x) = B(x) \cup U(x)$ and $\partial S_0(x) = B(x)$.
Bijections. Combining the above with the properties of $\mathcal{Z}, \mathcal{B}$, we have bijections between connected components:

$$\mathcal{Z} \longleftrightarrow \mathcal{B} \longleftrightarrow \partial\mathcal{B} = \partial\mathcal{S}_0 \overset{\pi^{\mathcal{S}}}{\longleftrightarrow} \partial S_0 = B_0 = B_0^- \sqcup B_0^+ = \partial A_0$$

and

$$\mathcal{U} \longleftrightarrow \partial\mathcal{U} \longleftrightarrow U$$

and we use these bijections to write

$$\partial A_0 = \partial A_0^- \sqcup \partial A_0^+$$

$$\partial\mathcal{B} = \partial\mathcal{B}^- \sqcup \partial\mathcal{B}^+$$

$$\mathcal{Z} = \mathcal{Z}^- \sqcup \mathcal{Z}^+.$$

We will use these frequently in the sequel. We denote connected components of $\mathcal{U}$ by $\mathcal{U}_u$ so that the correspondence $\mathcal{U} \supset \mathcal{U}_u \leftrightarrow u \in U$ is a bijection.

2. For each $x \in X$, the diagram

$$\begin{array}{ccc} \mathcal{S}_1(x) & \xrightarrow{\mathcal{F}^{\mathcal{S}}} & \mathcal{S}_0(f(x)) \\ \pi^{\mathcal{S}} \downarrow & & \downarrow \pi^{\mathcal{S}} \\ S_1(x) & \xrightarrow{f} & S_0(f(x)) \end{array}$$

commutes up to postcomposition of f by homeomorphisms of T which fix the vertices.

3. Given a connected embedded submanifold $E \subset \mathcal{S}$ which maps under $\mathcal{F}^{\mathcal{S}}$ onto an embedded submanifold of $\mathcal{S}$, we denote by $\deg(\mathcal{F}^{\mathcal{S}}, E)$ the positive topological degree of the restriction $\mathcal{F}^{S}SS|E$, i.e. the cardinality of a generic fiber.
We require that the unsigned (positive) degrees satisfy

$$\begin{aligned} \deg(\mathcal{F}^{\mathcal{S}},\ \mathcal{S}(x)) &= \deg(x) \\ \deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{B}_z) &= \deg(b_0) \\ \deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{U}_u) &= \deg(u) \end{aligned}$$

where $b_0 \leftrightarrow z$ under the bijection defined in the above axiom A-9(1).

4. a) **Nonperipheral.**
If x has valence 1 in $S_0(x)$, then $|\mathcal{Q} \cap \mathcal{S}_0(x)| \geq 2$.
b) **Nonparallel.** If x has valence 2 in $S_0(x)$, then $\mathcal{Q} \cap \mathcal{S}_0(x) \neq \emptyset$.
These two conditions will guarantee that, when combining, core curves of the annuli in $\mathcal{A}_0$ will form a multicurve.

2.3 Map of annuli over a mapping tree

Below, the *standard annulus* is

$$A = \mathbb{T} \times I = (\mathbb{R}/\mathbb{Z}) \times [-1, 1].$$

Points in $\mathbb{T}$ are denoted t and are taken modulo one. Points in I are denoted s. We equip A with the *standard orientation* which is the product of the orientations on $\mathbb{T}$ and on I. With this convention, the induced orientation on the boundary component $\mathbb{T} \times \{-1\}$ is opposite that of $\mathbb{T}$, while the orientation induced on the boundary component $\mathbb{T} \times \{+1\}$ agrees with that of $\mathbb{T}$. Note that with this convention, there is a canonical representative $t \mapsto (t, 0)$ for a generator of the fundamental group of A which we call the *core curve* of A. This allows us to measure the *signed degree* of any map between standard annuli. Moreover, the standard annulus has a preferred *linear direction* from the negative boundary component to the positive one.

Let $\mathbf{T}$ be a mapping tree with A_0, A_1 as above. A *map of annuli over* $\mathbf{T}$ is an orientation-preserving covering

$$\mathcal{F}^{\mathcal{A}} : \mathcal{A}_1 \to \mathcal{A}_0$$

where

- $\mathcal{A}_0 = A \times A_0$,
- $\mathcal{A}_1 \subset \mathcal{A}_0$,
- $\partial\mathcal{A}_1 \supset \partial\mathcal{A}_0$,
- $\mathcal{A}_1$ is homeomorphic to $A \times A_1$ via an orientation-preserving homeomorphism h such that $h|\partial\mathcal{A}_0 = \mathrm{id}_{\partial\mathcal{A}_0}$,

together with a continuous map

$$\pi^{\mathcal{A}} : \mathcal{A}_0 \to A_0 \subset T$$

satisyfing the following axioms:

A–10 **Subannuli.** We require that the inclusion map $\mathcal{A}_1 \hookrightarrow \mathcal{A}_0$ has signed degree one on each component. This is consistent with the manner in which we oriented edges of A_1 using the orientations induced from ambient edges in A_0.

A–11 **Covering.** Let $\mathcal{A}_{a_{1,j}}$ denote the component of $\mathcal{A}_1$ corresponding to $A \times \{a_{1,j}\}$ and write $\mathcal{A}_{a_{0,i}} = A \times \{a_{0,i}\}$. We require that the annuli in $\mathcal{A}_0$ lie over oriented edges in the obvious way, i.e. that for all $a_{0,i} \in A_0$ and $a_{1,j} \in A_1$ we have $\pi^{\mathcal{A}}|\mathcal{A}_{a_{0,i}}$ and $\pi^{\mathcal{A}}|\mathcal{A}_{a_{1,j}}$ are direction-preserving and proper maps onto $a_{0,i}$ and $a_{1,j}$ respectively. Since $a_{0,i} \simeq [-1,1]$ we might as well require additionally that

$$\pi^{\mathcal{A}}((\mathbb{T}, s) \times a_{0,i}) = s \in [-1,1] \simeq a_{0,i}$$

just to fix the picture in our minds.
Definition (missing annuli): We set

$$\mathcal{C} = \overline{\mathcal{A}_0 - \mathcal{A}_1}.$$

Each component of $\mathcal{C}$ is an annulus compactly contained in $\mathcal{A}_0$ and we call components of $\mathcal{C}$ the *missing annuli.*
Bijections. By construction we have bijections on the level of connected components

$$\begin{aligned}
&\mathcal{A}_0 \overset{\pi^{\mathcal{A}}}{\longleftrightarrow} A_0 \\
&\partial\mathcal{A}_0 \overset{\pi^{\mathcal{A}}}{\longleftrightarrow} \partial A_0 = B_0 \\
&\mathcal{A}_1 \overset{\pi^{\mathcal{A}}}{\longleftrightarrow} A_1 \\
&\partial\mathcal{A}_1 \overset{\pi^{\mathcal{A}}}{\longleftrightarrow} \partial A_1 = B \\
&\mathcal{C} \overset{\pi^{\mathcal{A}}}{\longleftrightarrow} C
\end{aligned}$$

Furthermore, we require that for all $a_1 \in A_1$, the diagram

$$\begin{array}{ccc} \mathcal{A}_{a_{1,j}} & \xrightarrow{\mathcal{F}^{\mathcal{A}}} & \mathcal{A}_{f(a_{1,j})} \\ \pi^{\mathcal{A}} \downarrow & & \downarrow \pi^{\mathcal{A}} \\ a_{1,j} & \xrightarrow{f} & f(a_{1,j}) \end{array}$$

commutes up to postcomposition with homeomorphisms of T fixing the vertices, and that the covering $\mathcal{F}^{\mathcal{A}}$ on this component has signed degree $\delta(b)$ where b is either boundary component of a_1.

A–12 **Standard form.** We require that in the coordinates given for $\mathcal{A}_0, \mathcal{A}_1$, if $\mathcal{F}_j^{\mathcal{A}} = \mathcal{F}^{\mathcal{A}}|\mathcal{A}_{a_{1,j}}$, and if δ is the signed degree of this map, then

$$\mathcal{F}_j^{\mathcal{A}}(s, -1) = (\delta(j)s, \ \ \mathrm{sgn}(\delta(j)) \cdot (-1))$$

$$\mathcal{F}_j^{\mathcal{A}}(s, +1) = (\delta(j)s, \ \ \mathrm{sgn}(\delta(j)) \cdot (+1))$$

where sgn is the function which is $+1$ on positive numbers and -1 on negative numbers. Note that while the absolute value of the signed degree is constant on subannuli of a given annulus, the signs may differ. This is because different subannuli within a given annulus may map to different annuli in $\mathcal{A}_0$ and our choices for the orientations on the edges of A_0 was arbitrary.

3

Combinations

In this section, we suppose that sphere maps and annulus maps lying over a common mapping tree are given. We introduce the notion of *gluing data*, which comes in two flavors. *Topological gluing data* describes how to glue the annulus and sphere maps together to obtain a map with domain a subset of an abstract S^2 and range equal to all of this same S^2. This map is as yet undefined on two regions: *missing disks* denoted $\mathcal{U}$ and *missing annuli* denoted $\mathcal{C}$. The missing disks may contain points of $\mathcal{Q}$, possibly critical, whose images are not yet determined and are specified by a choice of *critical gluing data.* The topological and critical gluing data must satisfy certain additional axioms in order for the resulting map $F : S^2 \to S^2$ to be a well-defined embellished branched covering.

The outcome of this process is summarized in §3.4, Theorem 3.2.

3.1 Topological gluing data

Recall that we have bijections

$$\partial\mathcal{B} \xrightarrow{\pi^{\mathcal{S}}} B_0 \xleftarrow{\pi^{\mathcal{A}}} \partial\mathcal{A}_0.$$

Furthermore, $B_0 = B^- \sqcup B^+$ and we have a bijective correspondence $B^- \leftrightarrow B^+$.

A *topological gluing* is a *choice* of a continuous map

$$\rho : \partial\mathcal{B} \to \partial\mathcal{A}_0$$

such that

- ρ is a conjugacy, i.e. $\rho \circ \mathcal{F}^{\mathcal{S}} = \mathcal{F}^{\mathcal{A}} \circ \rho$;
- ρ is orientation-reversing (with respect to the orientations induced on boundary components from orientations of the surfaces);

K.M. Pilgrim: LNM 1827, pp. 49–57, 2003.

- the map $\pi^{\mathcal{A}} \circ \rho \circ (\pi^{\mathcal{S}})^{-1}$ induces the involution $b_0^{\pm} \mapsto b_0^{\mp}$ on connected components.

  The *new sphere* S^2 is the quotient space

$$(\mathcal{S} - \mathcal{B}) \sqcup_\rho \mathcal{A}_0 \approx S^2.$$

The projection maps $\pi^{\mathcal{S}}, \pi^{\mathcal{A}}$ then define a continuous projection $\pi : S^2 \to T$.

3.2 Critical gluing data

At this point it is possible, after making some choices of extensions over $\mathcal{U}$ and $\mathcal{C}$, to define a new continuous map $F : S^2 \to S^2$ on the quotient space $\mathcal{S}_0 \sqcup_\rho \mathcal{A}_0 \approx S^2$. However, if there are $q \in \mathcal{Q}$ with $\mathcal{F}^{\mathcal{S}}(q) \in \mathcal{Z}$ then we lose control of the orbit of q under our new map: by the definition of $\mathcal{U}$, $\mathcal{F}^{\mathcal{S}}(q) \in \mathcal{Z}$ is equivalent to $q \in \mathcal{U} \cap \mathcal{Q}$, and we have not yet defined the new map over $\mathcal{U}$. Therefore, we specify the images of such points q first.

We do this by specifying the desired *mapping scheme* of the map to be constructed. Loosely, a mapping scheme is a set-theoretic object recording the dynamics and local degrees of a map $F : (S^2, Y) \to (S^2, Y)$ on the set Y.

Definition 3.1 (Mapping scheme). *A* mapping scheme *is a triple* (Y, τ, ω), *where* Y *is a set,* $\tau : Y \to Y$ *is the* dynamics function, *and* $\omega : Y \to \mathbb{N}$ *is the* degree function. *It is of* degree d spherical type *if the* Riemann-Hurwitz *condition*

$$\sum_{y \in Y} (\omega(y) - 1) = 2d - 2$$

and the local degree *condition*

$$\text{for all } y \in \tau(Y), \qquad \sum_{f(x)=y} \omega(x) \leq d$$

are satisfied. It is called complete *if the sharper condition*

$$\text{for all } y \in \tau(Y), \qquad \sum_{f(x)=y} \omega(x) = d$$

is satisfied.

Definition (critical gluing): A *critical gluing* is a choice of function

$$\kappa : \mathcal{Q} \cap \mathcal{U} \to \mathcal{Q}.$$

A critical gluing defines a *new mapping scheme* (Q, τ, ω), where $Q = \mathcal{Q}$, given by

$$\tau(q) = \begin{cases} \mathcal{F}^{\mathcal{S}}(q) & \text{if } \mathcal{F}^{\mathcal{S}}(q) \notin \mathcal{Z}, \\ \kappa(q) & \text{if } \mathcal{F}^{\mathcal{S}}(q) \in \mathcal{Z} \end{cases}$$

and

$$\omega(q) = \deg(\mathcal{F}^{\mathcal{S}}, q).$$

The map κ, however, is defined purely set-theoretically, and some restrictions are clearly necessary if we hope for the existence of an embellished branched covering $F : (S^2, Q) \to (S^2, Q)$ obtained by extending our maps $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}$. Thus we require some *topological compatibility.*

Definition (images of missing disks and annuli). First, some notation. Let $\mathcal{U}_u$ be any component of $\mathcal{U}$. Let

$$\mathcal{F}_*(\mathcal{U}_u) = \pi^{-1} f_*(u).$$

This is the topological disk in S^2 bounded by the component of $\partial\mathcal{B}$ lying over $b_0 = f(u)$ and which does not contain $\pi^{-1}(x)$, where $u \in U(x)$. We refer to $\mathcal{F}_*(\mathcal{U}_u)$ as the *image of the missing disk* $\mathcal{U}_u$.

Similarly, suppose $\mathcal{C}_c$ is any component of $\mathcal{C}$. Then $\mathcal{C}_c = \pi^{-1}(c) \subset \mathcal{A}_1$. Let

$$\mathcal{F}_*(\mathcal{C}_c) = \pi^{-1}(f_*(c)).$$

This is the annulus in S^2 bounded by the curves $\pi^{-1}(f(b_1))$ and $\pi^{-1}(f(b_1'))$, where $\partial c = [b_1, b_1']$. See Figure 3.1. We refer to $\mathcal{F}_*(\mathcal{C}_c)$ as the *image of the missing annulus* $\mathcal{C}_c$.

Note that $\mathcal{F}_*(\mathcal{U}_u)$ and $\mathcal{F}_*(\mathcal{C})$, as subsets of S^2, depend only on $\mathcal{S}_0$ and T, not on the maps involved or the choice of topological gluing.

Topological compatibility.

A–13 Recall that each component $\mathcal{U}_u$ contains at most one point of $\mathcal{Q}$. If this point is q, we require

$$\tau(q) \in \mathcal{F}_*(\mathcal{U}_u) \cap \mathcal{Q}.$$

Note that this is a perfectly reasonable condition if we are to map the disk $\mathcal{U}_u$ in a reasonable way onto $\mathcal{F}_*(\mathcal{U}_u)$. Axiom A–7 #4, which asserts that $|\mathcal{U}_u \cap \mathcal{Q}| \leq 1$, is designed to guarantee that there is an essentially unique way of extending the map over $\mathcal{U}_u$, and to avoid any complications in lifting equivalences; see the proof of the Uniqueness of Combinations Theorem in §4.

A–14 The mapping scheme (Q, τ, ω) is a complete spherical type mapping scheme of degree d, where d is the degree of the mapping tree.

Despite the set-theoretic appearance of this second axiom, it will have strong topological consequences, since it rules out the existence of extraneous preimages of points of $\tau(Q)$ under the map F (yet to be defined).

We are now ready to complete the definition of our new map.

3.3 Construction of combination

Gluing and extending. The new sphere is $S^2 = \mathcal{S}_0 \sqcup_\rho \mathcal{A}_0$. Identify $\mathcal{U}$ and $\mathcal{C}$ with subsets of our new sphere S^2. Since ρ is a conjugacy, the function given by $\mathcal{F}^{\mathcal{S}}$ on $\mathcal{S}_1$ and by $\mathcal{F}^{\mathcal{A}}$ on $\mathcal{A}_1$ is a continuous surjective map $\mathcal{S}_1 \cup_\rho \mathcal{A}_1 \to \mathcal{S}_0 \sqcup_\rho \mathcal{A}_0 = S^2$. However, there are some missing pieces in the domain of our map, namely the missing disks $\mathcal{U}$ and the missing annuli $\mathcal{C}$.

Extending over $\mathcal{U}$. To extend over $\mathcal{U}$, suppose $\mathcal{U}_u$ is a component of $\mathcal{U}$. Recall that $\mathcal{F}_*(\mathcal{U}) = \pi^{-1}(f_*(u))$ denotes the image of the missing disk $\mathcal{U}_u$. There are two cases to consider:

Case $\mathcal{Q} \cap \mathcal{U}_u = \emptyset$. In this case there are no critical points of $\mathcal{F}^{\mathcal{S}}$ in $\mathcal{U}_u$, since $\mathcal{Y} \cap \mathcal{U}_u = \emptyset$ and $\mathrm{Crit}(\mathcal{F}^{\mathcal{S}}) \subset \mathcal{Y}$. Since $\mathcal{U}_u$ is a component of a preimage of a disk in $\mathcal{S}$ under $\mathcal{F}^{\mathcal{S}}$ we have that $\deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{U}_u)$ has degree one. We extend over $\mathcal{U}_u$ by sending $(\mathcal{U}_u, \partial\mathcal{U}_u) \to (\mathcal{F}_*(\mathcal{U}_u), \partial\mathcal{F}_*(\mathcal{U}_u))$ by an arbitrary homeomorphism.

Case $\mathcal{Q} \cap \mathcal{U}_u \neq \emptyset$. By Axiom A–7 #4 this occurs only when $\mathcal{Q} \cap \mathcal{U}_u = \{q\}$, and by Topological Compatibility Axiom A–13 in this case we have $\tau(q) \in \mathcal{F}_*(\mathcal{U}_u) \cap \mathcal{Q}$. Using again the fact that $\mathcal{U}_u$ is a preimage of a disk in $\mathcal{S}$ under $\mathcal{F}^{\mathcal{S}}$, and the fact that there is at most one critical point of $\mathcal{F}^{\mathcal{S}}$ in $\mathcal{U}_u$, we have that $\deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{U}_u) = \deg(\mathcal{F}^{\mathcal{S}}, q) = \omega(q)$. We extend over $\mathcal{U}_u$ by sending $(\mathcal{U}_u, q) \to (\mathcal{F}_*(\mathcal{U}_u), \tau(q))$ via a covering ramified of degree $\deg(u) = \omega(q)$ over $\tau(q)$. This is possible by Covering Axiom A–9 #3, which guarantees that $\deg(\mathcal{F}^{\mathcal{S}}, \partial\mathcal{U}_u) = \deg(u)$.

The extension over $\mathcal{U}$ we denote by $\mathcal{F}^{\mathcal{U}}$ and refer to these maps as the *missing disk maps* .

Extending over $\mathcal{C}$. To extend over $\mathcal{C}$, suppose $\mathcal{C}_c$ is any component of $\mathcal{C}$. Recall that $\mathcal{F}_*(\mathcal{C}_c) = \pi^{-1}(f_*(c))$ denotes the image of the missing annulus $\mathcal{C}_c$. This is an annulus in S^2 bounded by $\pi^{-1}(f(b_1))$ and $\pi^{-1}(f(b_1'))$ which are boundary components of a component of $\mathcal{S}_0$. Moreover the degrees by which these components map are the same and are equal to $\deg(c)$. We extend by an unbranched covering $\mathcal{C}_c \to \mathcal{F}_*(\mathcal{C}_c)$ of this same degree.

The extension over $\mathcal{C}$ we denote by $\mathcal{F}^{\mathcal{C}}$ and refer to these maps as the *missing annulus maps* .

Definition of combination. Finally, we define $F : S^2 \to S^2$ by setting

$$F = \begin{cases} \mathcal{F}^{\mathcal{S}} & \text{on } \mathcal{S}_1 \\ \mathcal{F}^{\mathcal{A}} & \text{on } \mathcal{A}_1 \\ \mathcal{F}^{\mathcal{U}} & \text{on } \mathcal{U} \\ \mathcal{F}^{\mathcal{C}} & \text{on } \mathcal{C}. \end{cases}$$

We have now produced a continuous orientation-preserving map $F : S^2 \to S^2$ which we call an *amalgam.*

3.4 Summary: statement of Combination Theorem

Proposition 3.3, proved in the next subsection, together with the construction described, implies the following result.

Theorem 3.2 (Combination theorem). *Given data*

1. *a mapping tree* $\mathbf{T} = (T, f, \deg)$,
2. *a map* $\mathcal{F}^{\mathcal{S}}$ *of spheres over* $\mathbf{T}$
3. *a map* $\mathcal{F}^{\mathcal{A}}$ *of annuli over* $\mathbf{T}$
4. *a topological gluing* ρ
5. *a critical gluing* κ
6. *missing disk maps* $\mathcal{F}^{\mathcal{U}}$
7. *missing annulus maps* $\mathcal{F}^{\mathcal{C}}$

satisfying Axioms A–1 through A–14, there is a uniquely determined embellished branched covering

$$F : (S^2, Q) \to (S^2, Q).$$

The dependence of F on the data is a subtle issue and will be investigated in Section 4.

3.5 Properties of combinations

In this subsection, we assume that the map F is an amalgam produced using a particular choice of the data above. Below, recall that Q is defined as the image of $\mathcal{Q}$ under the inclusion map $\mathcal{S}_0 \hookrightarrow S^2$.

Proposition 3.3 (Amalgams are embellished.).

1. *Suppose that Axioms A1–A13 are satisfied. Then*
   a) F *is a branched covering of degree* d.
   b) $F(Q) \subset Q$.
   c) $\mathrm{Crit}(F) \subset Q$.
   Hence, (Q, τ, ω) *is a degree* d *mapping scheme of spherical type.*

2. *If in addition Axiom A14 is satisfied, then the mapping scheme (Q, τ, ω) is complete, and hence $F : (S^2, Q) \to (S^2, Q)$ is an embellished branched covering.*

Proof:

1a) The map F is clearly a branched covering. That the degree is d follows easily from the Global Homogeneity Axioms A–4 #2 and the Covering Axioms A–9 and the definition of F.

1b) For all $q \in Q$ we have $F(q) = \tau(q)$ by construction. To see this, recall that $Q \subset \mathcal{S}_0 = \mathcal{S}_1 \cup \mathcal{U}$. If $q \in Q \cap \mathcal{S}_1$ then $F(q) = \mathcal{F}^{\mathcal{S}}(q) = \tau(q)$ by the definition of τ, while if $q \in Q \cap \mathcal{U}$ then $F(q) = \tau(q)$ by the definition of $F = \mathcal{F}^{\mathcal{U}}$ on $\mathcal{U}$. Since $\tau : Q \to Q$, we have that $F(Q) \subset Q$ and so Q is forward-invariant.

1c) The critical points of F are of two types: those in $\mathcal{S}_1$, which are in Q since they are critical points of $\mathcal{F}^{\mathcal{S}}$, and those in $\mathcal{U}$. By construction of $\mathcal{F}^{\mathcal{U}}$, if o is a critical point of $\mathcal{F}^{\mathcal{U}}$, then the component $\mathcal{U}_u$ containing o contains a point q of Q, and by construction of $\mathcal{F}^{\mathcal{U}}$ this point must be o itself. Hence $\mathrm{Crit}(F) \subset Q$ and in fact $\mathrm{Crit}(F) = \mathrm{Crit}(Q)$, where the latter set is defined as those $q \in Q$ for which $\omega(q) > 1$.

2. We have $F|Q = \tau$ by the definition of F and τ. Also, for all $q \in Q$, $\deg(F, q) = \omega(q)$. Topological Compatibility Axiom A14 implies that each $p \in \tau(Q)$ has d preimages under τ in Q, counted with multiplicity. Hence each $p \in F(Q)$ has d preimages under F in Q, counted with multiplicity. Since the degree of F is d and $F|Q = \tau|Q$, we have $F^{-1}(F(Q)) \subset Q$. The inclusion in the other direction is trivial. □

As a consequence of Topological Compatibility Axiom 14 and the last paragraph of the previous proof, we have

Proposition 3.4 (Image of Q constrained). *For all $c \in C$, we have $F(Q) \cap F(\mathcal{C}_c) = \emptyset$.*

Proof: Otherwise, there are strictly more than d preimages of points in the intersection. □

Proposition 3.5. *If $\mathcal{A}_{a_{0,i}} \subset \mathcal{F}_*(\mathcal{C}_c)$, then $\mathcal{A}_{a_{0,i}}$ is inessential in $\mathcal{F}_*(\mathcal{C}_c)$.*

Proof: The annulus $\mathcal{F}_*(\mathcal{C}_c)$ lies over a subtree of T as shown in Figure 3.1. □To set up the next statement, let

$$A_0 = \{a_{0,1}, ..., a_{0,i}, ..., a_{0,n_0}\}$$

and

$$A_1 = \{a_{1,1}, ..., a_{1,j}, ..., a_{1,n_1}\}.$$

Let Γ denote the multicurve in $S^2 - Q$ whose elements are the n_0 core curves of the annuli $\mathcal{A}_{0,i}$ in $\mathcal{A}_0$.

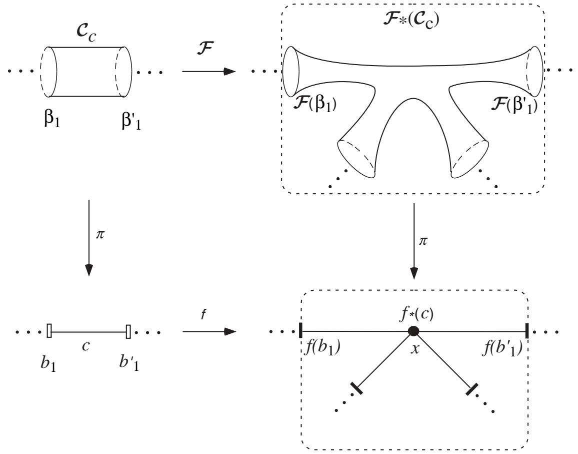

Fig. 3.1. The annulus $\mathcal{F}_*(\mathcal{C}_c)$ is bounded by two curves in $\partial\mathcal{S}_0(x)$ for a unique x, and lies over a subtree of T as indicated.

Proposition 3.6 (Core curves form multicurve). *Γ is an F-invariant multicurve, and the Thurston linear transformation $F_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ is given by the matrix $N_{\mathbf{T}}$, where*

$$(N_{\mathbf{T}})_{i,j} = \sum_{\substack{a_{1,k} \subset a_{0,i} \\ f(a_{1,k}) = a_{0,j}}} \frac{1}{|\delta(a_{1,k})|},$$

where $\delta(a_{1,k})$ is the signed degree of either boundary component of $a_{1,k}$.

In particular, F_Γ depends only on T.

Proof: Covering Axioms A–9 #4(a,b) guarantees that elements of Γ are essential, nonperipheral, and pairwise nonhomotopic. To check invariance, fix $\gamma \in \Gamma$. The preimages $\widetilde{\gamma}$ of γ under F consist of curves in

1. $\mathcal{A}_1$,
2. possibly $\mathcal{U}$, and
3. possibly $\mathcal{C}$.

The curves in case (1) are homotopic rel Q to an element of Γ. The curves in case (2) are all inessential or peripheral, by Axiom A–7 #4. The curves in case (3) are all inessential since any such preimage is the preimage under a missing annulus map $\mathcal{F}^{\mathcal{C}}$, and by Proposition 3.5, γ is inessential in any annulus $\mathcal{F}_*(\mathcal{C})$. To check the second invariance condition, let $a_{0,i} \in A_0$. By the definition of A_1 and A_0, the annular edge $a_{0,i}$ contains an element $a_{1,j} \in A_1$ as a subedge. Hence $f(a_{0,i}) \supset f(a_{0,j})$, an element of A_0. Lifting to the sphere under $\pi^{\mathcal{A}}$ shows that there is a preimage of an element of Γ homotopic to $a_{0,i}$. Finally, the Covering Axioms for the annulus maps, and the observation that the only preimages contributing in the definition of F_Γ are those in Case 1, show that the transformation F_Γ has the indicated form. □

The next results concern the persistence of invariant multicurves under amalgamation. To set up the statement, let $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$, and denote by $\mathrm{scc}(\mathcal{S}, \mathcal{Y})$ and $\mathrm{scc}(S^2, Q)$ the sets of homotopy classes of essential, simple, possibly peripheral closed curves in $\mathcal{S} - \mathcal{Y}$ and $S^2 - Q$, respectively. Let $[\Gamma] \subset \mathrm{scc}(S^2, Q)$ denote the classes of the elements of the multicurve Γ as in the previous Proposition 3.6. A simple closed curve in $\mathcal{S} - \mathcal{Y}$ is homotopic rel $\mathcal{Y}$ to a curve in $\mathcal{S}_0 - \mathcal{Q}$, and by Axiom 7, part 4(a), the image of an essential curve in $\mathcal{S}_0 - \mathcal{Q}$ under the inclusion $\iota : \mathcal{S}_0 \hookrightarrow S^2$ is essential in $S^2 - Q$. Hence ι induces a function

$$\iota_* : \mathrm{scc}(\mathcal{S}, \mathcal{Y}) \to \mathrm{scc}(S^2, Q)$$

with the following easily verified properties:

Proposition 3.7 (Curves under inclusion).

1. *A class belongs to the image of ι_* if and only if up to homotopy in $S^2 - Q$ it is disjoint from Γ. (Parallel curves are disjoint.)*
2. *The image under ι_* of a homotopy class $[\alpha]$ is represented by a curve which is peripheral in $S^2 - Q$ if and only if α is peripheral in both $\mathcal{S} - \mathcal{Y}$ and $\mathcal{S} - \mathcal{Q}$, i.e. α is homotopic in $\mathcal{S} - \mathcal{Y}$ to a curve surrounding a single element of $\mathcal{Q}$.*
3. *The image under ι_* of a homotopy class $[\alpha]$ lies in $[\Gamma]$ if and only if α is peripheral in $\mathcal{S} - \mathcal{Y}$ and $\mathcal{S} - \mathcal{Z}$, i.e. α is homotopic in $\mathcal{S} - \mathcal{Y}$ to a curve surrounding a single element of $\mathcal{Z}$. (Equivalently, α is homotopic in $\mathcal{S} - \mathcal{Y}$ to a boundary component of $\mathcal{B}$.)*
4. *The images under ι_* of two classes $[\alpha^+], [\alpha^-]$ coincide if and only if (i) their common image lies in Γ, and (ii) α^+, α^- are respectively homotopic to two boundary components of $\mathcal{B}$ given by b_0^-, b_0^+. Moreover, every class in $[\Gamma]$ has exactly two preimages under ι_* arising in this fashion.*
5. *If $\Gamma_1^{\mathcal{S}}$ is a multicurve in $\mathcal{S} - \mathcal{Y}$, then $\iota_* | [\Gamma_1^{\mathcal{S}}]$ is a bijection, and the image $[\Gamma_1] = \iota_*([\Gamma_1^{\mathcal{S}}])$ is a set of classes represented by a multicurve Γ_1 in $S^2 - Q$.*

Proposition 3.8 (Invariant multicurves persist). *Let $\Gamma_1^{\mathcal{S}}$ be a multicurve in $\mathcal{S} - \mathcal{Y}$ and let $\Gamma_1 \subset S^2 - Q$ represent the image $\iota_*([\Gamma_1^{\mathcal{S}}])$.*

1. *If Γ_1 is F-invariant, then $\Gamma_1^{\mathcal{S}}$ is $\mathcal{F}^{\mathcal{S}}$-invariant.*
2. *Suppose $\Gamma_1^{\mathcal{S}}$ is invariant under $\mathcal{F}^{\mathcal{S}}$.*
   a) *If no element of Γ_1 is essential in the image $\mathcal{F}_*(\mathcal{C}_c)$ of some missing annulus $\mathcal{C}$, then Γ_1 is invariant under F. In this case, the bijection $\iota_* : [\Gamma_1^{\mathcal{S}}] \to [\Gamma_1]$ defines a conjugacy from the Thurston linear map determined by $\mathcal{F}^{\mathcal{S}}$ and $\Gamma_1^{\mathcal{S}}$ to the Thurston linear map determined by F and Γ_1.*

b) *Otherwise, there are preimages of elements of Γ_1 under F which are homotopic in $S^2 - Q$ to elements of Γ. In this case, $\Gamma \sqcup \Gamma_1$ is F-invariant, and the Thurston linear transformation $F_{\Gamma \sqcup \Gamma_1}$ has matrix of the form*

$$\begin{pmatrix} N_{\mathbf{T}} & * \\ 0 & N_1 \end{pmatrix}$$

where $N_{\mathbf{T}}$ is the matrix for the Thurston linear map determined by F and Γ, and N_1 is the matrix for the Thurston linear map determined by $\mathcal{F}^{\mathcal{S}}$ and $\Gamma_1^{\mathcal{S}}$.

Proof: First, by Proposition 3.7, (3) above, and the fact that multicurves consist of nonperipheral curves, no element of Γ_1 is homotopic to an element of Γ. Next, the hypothesis implies that we may assume that $\Gamma_1 \subset \mathcal{S}_0$. Finally, recall that by construction of the amalgam F we have $F = \mathcal{F}^{\mathcal{S}}$ on $\mathcal{S}_1$.

With these observations, conclusion (1) follows immediately, using the fact that a preimage under $\mathcal{F}^{\mathcal{S}}$ of a simple closed curve in $\mathcal{S}_0 \subset \mathcal{S}$ is either in $\mathcal{S}_1$ (and hence corresponds to a preimage under F) or in $\mathcal{U}$ (and hence is peripheral in $\mathcal{S} - \mathcal{Y}$ by Covering Axiom A–7#4).

Let us now prove (2). Choose $\gamma_1^{\mathcal{S}} \in \Gamma_1^{\mathcal{S}}$ and let $\gamma_1 = \iota(\gamma_1^{\mathcal{S}})$. Let δ be a component of $F^{-1}(\gamma_1)$. By the construction of the amalgam F, there are four possibilities for where δ resides:
(i) in $\mathcal{U}$, in which case δ is peripheral;
(ii) in $\mathcal{A}_1$, but this is impossible since then $\gamma_1 \subset \mathcal{A}_0$ and we are assuming $\gamma_1 \subset \mathcal{S}_0$ which is disjoint from $\mathcal{A}_0$;
(iii) in $\mathcal{C}$;
(iv) in $\mathcal{S}_1$.
Now suppose that the hypothesis in 2(a) holds, i.e. that γ_1 is inessential in the image of every missing annulus. If $\delta \subset \mathcal{C}_c$ for some c, then δ must be inessential in $\mathcal{C}_c$, and so since $\mathcal{C}_c \cap Q = \emptyset$ we conclude that δ is inessential in $S^2 - Q$. Hence, if δ is essential in $S^2 - Q$, it lies in $\mathcal{S}_1$. Since the inclusion ι is a conjugacy between $\mathcal{F}^{\mathcal{S}}$ and F on $\mathcal{S}_1$, the conclusions in 2(a) follow immediately.

Now suppose that the hypothesis in 2(b) holds for this γ_1 and δ. Then δ is an essential curve in $\mathcal{C}_c$, and hence δ is homotopic in $S^2 - Q$ to a core curve of $\mathcal{A}_0$, i.e. to an element of Γ. Thus $\Gamma \sqcup \Gamma_1$ satisfies condition (1) in the definition of invariant multicurve (Definition 1.13). By Proposition 3.6, $N_{\mathbf{T}}$ has a nonzero entry in every row. Since Γ_1 is an invariant multicurve, it satisfies condition (2) in the definition, and this implies that the matrix N_1 has a nonzero entry in every row. Thus, $\Gamma \sqcup \Gamma_1$ also satisfies condition (2). Hence, $\Gamma \sqcup \Gamma_1$ is an invariant multicurve. □

4

Uniqueness of combinations

Here, we investigate the dependence of the combinatorial class of an amalgam F on the data used in its definition, and introduce several notions of equivalence to capture this dependence.

4.1 Structure data and amalgamating data.

First, we pin down the global topological structure of the input data, and the set-theoretic dynamics on the set $\mathcal{Y}$.

Definition 4.1 (Structure data). *The* structure data *for amalgamation consists of:*

1. *a mapping tree* $\mathbf{T}$ *and a set* $\mathcal{S}$ *of spheres, indexed by* X*;*
2. *the sets* $\mathcal{Q}, \mathcal{Z}$ *as subsets of* $\mathcal{S}$*;*
3. *the restriction of a map of spheres* $\mathcal{F}^{\mathcal{S}}$ *over* $\mathbf{T}$ *to the set* $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$*, so that the set-theoretic data of the dynamics on* $\mathcal{Y}$ *is fixed;*
4. *the local degrees of* $\mathcal{F}^{\mathcal{S}}$ *near points in* $\mathcal{Y}$*;*
5. *a critical gluing* $\kappa : \mathcal{Q} \cap \mathcal{U} \to \mathcal{Q}$*.*

While most of these conditions are reasonable, the second warrants some explanation. If we are free to vary these sets while keeping, essentially, the same dynamics (and local degrees), then we must introduce a further notion of equivalence. However, we can always arrange so that these sets are the same by a suitable conjugation. So this is not really a serious restriction.

Next, we formally introduce the parameters on which the construction of an amalgam depends.

Definition 4.2 (Amalgamating data). *Fix a structure data for amalgamation.* Amalgamating data with respect to this structure data *is a 5-tuple*

$$\mathbf{D} = \{\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}, \rho, \mathcal{F}^{\mathcal{U}}, \mathcal{F}^{\mathcal{C}}\}$$

where the terms and the dependencies are indicated as follows:

K.M. Pilgrim: LNM 1827, pp. 59–68, 2003.

1. *$\mathcal{F}^{\mathcal{S}}$ is a map of spheres over* $\mathbf{T}$*, satisfying conditions (1)-(5) in Definition 4.1;*
2. *$\mathcal{F}^{\mathcal{A}}$ is a map of annuli over* $\mathbf{T}$*;*
3. *ρ is a topological gluing (depends on $\mathcal{F}^{\mathcal{S}}$ and $\mathcal{F}^{\mathcal{A}}$)*
4. *$\mathcal{F}^{\mathcal{U}}$ are missing disk maps (depends on $\mathcal{F}^{\mathcal{S}}, \rho$)*
5. *$\mathcal{F}^{\mathcal{C}}$ are missing annulus maps (depends on $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}$, and ρ)*

Given $\mathbf{D}$*, we denote by* $F(\mathbf{D})$ *the amalgam produced with the data* $\mathbf{D}$.

4.2 Combinatorial equivalence of sphere and annulus maps

Next, we formulate notions of combinatorial equivalence of sphere and annulus maps.

Definition 4.3 (Combinatorial equivalence of sphere maps over T). *Let $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}'}$ be two families of sphere maps lying over a common mapping tree T such that*

1. $\mathcal{S}_0 = \mathcal{S}_0', \mathcal{S}_1 = \mathcal{S}_1'$
2. $\mathcal{F}^{\mathcal{S}}|\partial\mathcal{S}_1 = \mathcal{F}^{\mathcal{S}'}|\partial\mathcal{S}_1'$

We say $\mathcal{F}^{\mathcal{S}}$ and $\mathcal{F}^{\mathcal{S}'}$ are combinatorially equivalent over $\mathbf{T}$ *if there exist tame orientation-preserving homeomorphisms $\mathcal{H}_0, \mathcal{H}_1 : \mathcal{S} \to \mathcal{S}'$ such that the following hold:*

1. $\mathcal{H}_0 \circ \mathcal{F} = \mathcal{F}' \circ \mathcal{H}_1$
2. $\mathcal{H}_0|_{\partial\mathcal{S}_1} = \mathcal{H}_1|_{\partial\mathcal{S}_1} = \mathrm{id}_{\partial\mathcal{S}_1}$
3. $\mathcal{H}_0|_{\mathcal{S}_1} \simeq \mathcal{H}_1$ *through tame homeomorphisms which are the identity on* $\partial\mathcal{S}_1 \sqcup \mathcal{Q}$

This notion of equivalence is distinct from the definition of equivalence of an embellished map of spheres per se, as given in Definition 1.12. The difference is that here we require some control over the boundary values. With this definition, however, it is easy to check that if e.g. $\mathcal{Q}$ is finite and no additional regularity beyond continuity is required in the definition of tameness, then the set of such classes naturally a discrete topological space. This need not be the case if e.g. $\mathcal{Q}$ is infinite, where to recover discreteness some tameness assumptions might be necessary.

Definition 4.4 (Combinatorial equivalence of annulus maps). *Suppose $\mathcal{F}^{\mathcal{A}}, \mathcal{F}^{\mathcal{A}'}$ are two families of annulus maps lying over a common mapping tree* $\mathbf{T}$ *such that*

1. $\mathcal{A}_0 = \mathcal{A}_0', \mathcal{A}_1 = \mathcal{A}_1'$,
2. $\mathcal{F}^{\mathcal{A}}|\partial\mathcal{A}_1 = \mathcal{F}^{\mathcal{A}'}|\partial\mathcal{A}_1'$.

We say $\mathcal{F}^{\mathcal{A}}$ is combinatorially equivalent *to $\mathcal{F}^{\mathcal{A}'}$ if there exist orientation-preserving homeomorphisms $\mathcal{H}_0^{\mathcal{A}} : \mathcal{A}_0 \to \mathcal{A}'_0$ and $\mathcal{H}_1^{\mathcal{A}} : \mathcal{A}_1 \to \mathcal{A}'_1$ with the following properties:*

1. $\mathcal{H}_0^{\mathcal{A}} \circ \mathcal{F}^{\mathcal{A}} = \mathcal{F}^{\mathcal{A}'} \circ \mathcal{H}_1^{\mathcal{A}}$
2. $\mathcal{H}_0^{\mathcal{A}} | \partial \mathcal{A}_1 = \mathcal{H}_1^{\mathcal{A}} | \partial \mathcal{A}_1 = \mathrm{id}_{\partial \mathcal{A}_1}$
3. *the homeomorphism $\overline{\mathcal{H}_1}^{\mathcal{A}} : \mathcal{A}_0 \to \mathcal{A}'_0$ obtained by extending $\mathcal{H}_1^{\mathcal{A}}$ by the* identity *over $\mathcal{C} = \mathcal{A}_0 - \mathcal{A}_1$ is isotopic to $\mathcal{H}_0^{\mathcal{A}}$ through homeomorphisms which are the identity on $\partial \mathcal{A}_0$.*

Note that the extension over $\mathcal{C}$ in (3) is required to be the identity.

4.3 Statement of Uniqueness of Combinations Theorem

Assume a fixed structure data for amalgamation is given.

One would like a statement which asserts that given a set of amalgamating data $\mathbf{D}$, the combinatorial class of an amalgam $F = F(\mathbf{D})$ depends only on the classes of sphere maps $\mathcal{F}^{\mathcal{S}}$ and of annulus maps $\mathcal{F}^{\mathcal{A}}$. This is not possible: choices must be made for the extension $\mathcal{F}^{\mathcal{C}}$ of F over the missing annuli $\mathcal{C}$. That is, for fixed $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}$ and varying $\mathcal{F}^{\mathcal{C}}$, it is possible to produce amalgams F in distinct classes. It turns out that the choice of missing disk maps $\mathcal{F}^{\mathcal{U}}$ is irrelevant (Proposition 4.7).

Rather than formulate a lengthy notion of equivalence of amalgamating data, we will show that once a basepoint $\mathbf{D}$ has been fixed, the amalgamating data $\mathbf{D}'$ can be altered to produce new amalgamating data $\mathbf{D}''$ such that (i) $F(\mathbf{D}')$ is equivalent to $F(\mathbf{D}'')$, and (ii) $\mathbf{D}''$ is in "simple form with respect to $\mathbf{D}$" (Definition 4.11). Hence, in order to enumerate the set of classes of amalgams $F(\mathbf{D}')$ which can be produced by varying the amalgamating data $\mathbf{D}'$, it suffices to assume that this data is in simple form with respect to $\mathbf{D}$.

We then have

Theorem 4.5 (Uniqueness of combinations). *Fix a basepoint $\mathbf{D}$ of amalgamating data. Suppose amalgamating data $\mathbf{D}'$ is in simple form with respect to $\mathbf{D}$. If the sphere maps $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}'}$ are combinatorially equivalent, and if the annulus maps $\mathcal{F}^{\mathcal{A}}, \mathcal{F}^{\mathcal{A}'}$ are combinatorially equivalent, then the amalgams $F(\mathbf{D}), F(\mathbf{D}')$ are combinatorially equivalent.*

Under the same hypotheses, we have

Corollary 4.6 (Bound on classes with fixed sphere class). *For a fixed combinatorial class of sphere maps, the number of combinatorial classes of amalgams is bounded by the number of combinatorial classes of annulus maps.*

In §§7.1, 7.2 respectively we state and prove a theorem giving the number of combinatorial classes of annulus maps, making the bound in the above Corollary explicit.

4.4 Proof of Uniqueness of Combinations Theorem

4.4.1 Missing disk maps irrelevant

Proposition 4.7 ($[F(\mathbf{D}')]$ is independent of $\mathcal{F}^{\mathcal{U}'}$). *The combinatorial class $[F(\mathbf{D}')]$ is independent of the choice of the missing disk maps $\mathcal{F}^{\mathcal{U}'}$.*

Proof: Assume F', F'' are amalgams produced from amalgamating data $\mathbf{D}', \mathbf{D}''$ which differ only in the choice of missing disk maps $\mathcal{F}^{\mathcal{U}'}, \mathcal{F}^{\mathcal{U}''}$. We claim that there is a combinatorial equivalence $\mathcal{H}_0, \mathcal{H}_1$ between F' and F'' where $\mathcal{H}_0 = \mathrm{id}_{S^2}$.

Note that this makes sense. The assumption $\mathcal{F}^{\mathcal{S}'} = \mathcal{F}^{\mathcal{S}''}$ implies in particular that the subsets $\mathcal{S}_0', \mathcal{S}_0''$ coincide. Henceforth, let us simply write $\mathcal{S}_0$ for these sets. By assumption $\rho' = \rho''$, so the domains of definition of F', F'', which are both spheres, may be regarded as the same sphere S^2. Then $\mathcal{A}_0' = \mathcal{A}_0''$ and we denote this common subset of S^2 by simply $\mathcal{A}_0$. Also, $\mathcal{U}' = \mathcal{U}''$ and we denote this set as simply $\mathcal{U}$ as well. Note that it now makes sense to speak of id_{S^2}. Let $\mathcal{H}_0 = \mathrm{id}_{S^2}$.

To find $\mathcal{H}_1$, set $\mathcal{H}_1 = \mathrm{id}$ on $\mathcal{A}_0 \sqcup \mathcal{S}_1$. To complete the definition of $\mathcal{H}_1$, we need only extend it over $\mathcal{U}$. Above each $u \in U$, $\mathcal{F}^{\mathcal{U}'}$ and $\mathcal{F}^{\mathcal{U}''}$ are branched coverings of disks, with the same boundary values, ramified of the same degree over the same single point. Hence for each $u \in U$ there is no obstruction to lifting the restriction of $\mathcal{H}_0$ to the set $\mathcal{F}_*'(\mathcal{U}_u) = \mathcal{F}_*''(\mathcal{U}_u)$ under F' and F'' to obtain a homeomorphism which is the identity on $\partial\mathcal{U}_u \cup \{u\}$. Thus now $\mathcal{H}_1$ is defined on all of S^2 and $\mathcal{H}_0 \circ F' = F'' \circ \mathcal{H}_1$ by construction.

It remains to show $\mathcal{H}_0 \simeq \mathcal{H}_1$ rel Q. This follows immediately from the Alexander Trick (see below). □

The Alexander Trick ([Bir], Lemma 4.4.1) *If $g : D^n \to D^n$ is a homeomorphism from the unit n-ball to itself which fixes the $(n-1)$-sphere $S^{n-1} = \partial D^n$ pointwise, then g is isotopic to the identity under an isotopy which fixes S^{n-1} pointwise. If $g(0) = 0$ then the isotopy may be chosen to fix 0.*

Hence in the following, we suppress mention of the missing disk maps $\mathcal{F}^{\mathcal{U}}$ in the amalgamating data.

4.4.2 Reduction to fixed boundary values

In what follows, we assume that we have chosen a basepoint $\mathbf{D} = \{\mathcal{F}^{\mathcal{S}}, \rho, \mathcal{F}^{\mathcal{A}}, \mathcal{F}^{\mathcal{C}}\}$ and we set $F = F(\mathbf{D})$.

Proposition 4.8 (Reduction to fixed boundary values). *Let $F' = F(\mathbf{D}')$. Then F' is combinatorially equivalent to $F'' = F(\mathbf{D}'')$, where in the data $\mathbf{D}''$ we have*

- $\mathcal{S}_0^{''} = \mathcal{S}_0$,
- $\mathcal{S}_1'' = \mathcal{S}_1$ *(and so $\mathcal{U}'' = \mathcal{U}$),*
- $\rho'' = \rho$

- $\mathcal{A}_1'' = \mathcal{A}_1$ *(and so* $\mathcal{C}'' = \mathcal{C}$*),*
- $\mathcal{F}^{\mathcal{U}''}|\partial\mathcal{U}'' = \mathcal{F}^{\mathcal{U}}|\partial\mathcal{U}$*, and*
- $\mathcal{F}^{\mathcal{A}''}|\partial\mathcal{A}_1'' = \mathcal{F}^{\mathcal{A}}|\partial\mathcal{A}_1$*.*

Proof: The homeomorphism $\phi_0 = (\rho')^{-1} \circ \mathrm{id}_{\partial\mathcal{A}_0} \circ \rho$ conjugates $\mathcal{F}^{\mathcal{S}}|\partial\mathcal{S}_0$ to $\mathcal{F}^{\mathcal{S}'}|\partial\mathcal{S}_0'$. Lift this homeomorphism under F and F' to a conjugacy $\phi_1 : \partial\mathcal{S}_1 \sqcup \partial(\mathcal{A}_1 - \mathcal{A}_0) \to \partial\mathcal{S}_1' \sqcup \partial(\mathcal{A}_1' - \partial\mathcal{A}_0')$. Extend ϕ_1 arbitrarily to a tame homeomorphism Φ fixing $\mathcal{Q} \sqcup \mathcal{Z}$ and sending the pairs $(\mathcal{S}_0, \mathcal{S}_1) \to (\mathcal{S}_0', \mathcal{S}_1')$, $(\mathcal{A}_0, \mathcal{A}_1) \to (\mathcal{A}_0', \mathcal{A}_1')$. Then $\Phi : S^2 \to S^2$ is continuous. Now set $F'' = \Phi^{-1} \circ F' \circ \Phi$, i.e. "pull back" F' under Φ. Pulling back the data $\mathbf{D}'$ to $\mathbf{D}''$ we find that the resulting map $F'' = F(\mathbf{D}'')$ then has the indicated properties. □

Thus the set of all combinatorial classes of amalgams obtained with fixed structure data coincides with the set of all combinatorial classes of amalgams produced using the subset of amalgamating data where the spaces $\mathcal{S}_0, \mathcal{S}_1, \mathcal{A}_0, \mathcal{A}_1, \mathcal{U}, \mathcal{C}$ are fixed, the topological gluing is fixed, and the boundary values of the maps are fixed and equal to ρ.

Assumption. Hence, once the basepoint $\mathbf{D}$ is chosen, we may assume that whenever $\mathbf{D}' = \{\mathcal{F}^{\mathcal{S}'}, \rho', \mathcal{F}^{\mathcal{A}'}, \mathcal{F}^{\mathcal{C}'}\}$ is amalgamating data defining F', we have in particular that the boundary values $\rho' = \rho$. Summarizing, we refer to this as *having fixed the structure data and the boundary values* . The data required to form an amalgam then consists only of the sphere, annulus, and missing disk and annulus maps; by Proposition 4.7 we suppress mention of the missing disk maps, since they do not affect the class of the amalgam. In short, we now write

$$\mathbf{D} = \{\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{A}}, \mathcal{F}^{\mathcal{C}}\}$$

for the basepoint data defining F and

$$\mathbf{D}' = \{\mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{A}'}, \mathcal{F}^{\mathcal{C}'}\}$$

for that defining F'.

4.4.3 Reduction to simple form

Recall from the previous subsection that we assume that the structure data and the boundary values are fixed.

Proposition 4.9. *Fix a basepoint* $F = F(\mathbf{D})$*, where*

$$\mathbf{D} = \{\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{A}}\}.$$

Let $F' = F(\mathbf{D}')$*, where*

$$\mathbf{D}' = \{\mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{C}'}, \mathcal{F}^{\mathcal{A}'}\}$$

and assume $\mathcal{F}^{\mathcal{S}'}$ *is combinatorially equivalent to* $\mathcal{F}^{\mathcal{S}}$*.*

Then F' is combinatorially equivalent to an amalgam $F'' = F(\mathbf{D}'')$ which is produced using a map $\mathcal{F}^{\mathcal{S}''}$ of spheres combinatorially equivalent over $\mathbf{T}$ to $\mathcal{F}^{\mathcal{S}}$ via a pair $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ of tame homeomorphisms with $\mathcal{H}_0^{\mathcal{S}} = \mathrm{id}$.

The content of this proposition is that we may assume $\mathcal{H}_0^{\mathcal{S}} = \mathrm{id}$ without altering the class of amalgam.

Proof: (of Proposition 4.9) Let $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ be a pair of tame homeomorphisms yielding an equivalence. By assumption, the boundary values are the same, hence we may extend $\mathcal{H}_0^{\mathcal{S}}$ by the identity on $\mathcal{A}_0$ to obtain a homeomorphism $\mathcal{H}_0 : S^2 \to S^2$. Let $F'' = (\mathcal{H}_0)^{-1} \circ F' \circ \mathcal{H}_0$. Then F'' is tamely topologically conjugate to F'. Moreover, the map of spheres $\mathcal{F}^{\mathcal{S}''} = (\mathcal{H}_0)^{-1} \circ F' \circ \mathcal{H}_0$ is combinatorially equivalent to $\mathcal{F}^{\mathcal{S}}$ via the pair $\mathcal{H}_0^{\mathcal{S}} \circ (\mathcal{H}_0^{\mathcal{S}})^{-1} = \mathrm{id}_{\mathcal{S}_0}$ and $\mathcal{H}_1^{\mathcal{S}} \circ (\mathcal{H}_0^{\mathcal{S}})^{-1}$, as required. □

Assumption. We may therefore assume that in any set of amalgamating data $\mathbf{D}'$ for which the map of spheres $\mathcal{F}^{\mathcal{S}'}$ is combinatorially equivalent over $\mathbf{T}$ via a pair $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ to the map of spheres $\mathcal{F}^{\mathcal{S}}$ in our basepoint $\mathbf{D}$, we have $\mathcal{H}_0 = \mathrm{id}_{\mathcal{S}_0}$.

At this point it is tempting to try to glue together equivalences between the sphere and annulus maps. However, the extensions $\mathcal{F}^{\mathcal{C}}$ over the missing annuli $\mathcal{C}$ are a potential source of obstructions to the lift $\mathcal{H}_1$ of an equivalence $\mathcal{H}_0 = \mathcal{H}_0^{\mathcal{S}} \cup \mathcal{H}_0^{\mathcal{A}}$ being isotopic to $\mathcal{H}_0$. The Unwinding Trick provides one more degree of freedom in altering the input data so as to achieve the desired "simple form", for which this obstruction vanishes.

Below, note that once the boundary values and structure data are fixed, the missing annuli $\mathcal{C}_c$ and their images $\mathcal{F}_*(\mathcal{C}_c)$ are well-defined as subsets of S^2.

Proposition 4.10 (The unwinding trick). *Fix a basepoint $F = F(\mathbf{D})$, where*

$$\mathbf{D} = \{\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{A}}\}.$$

Let $F' = F(\mathbf{D}')$, where

$$\mathbf{D}' = \{\mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{C}'}, \mathcal{F}^{\mathcal{A}'}\}$$

and for which $\mathcal{F}^{\mathcal{S}'}$ is combinatorially equivalent to $\mathcal{F}^{\mathcal{S}}$ via a pair $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ with $\mathcal{H}_0 = \mathrm{id}_{\mathcal{S}_0}$.

Then F' is combinatorially equivalent to $F'' = F(\mathbf{D}'')$, where

$$\mathbf{D}'' = \{\mathcal{F}^{\mathcal{S}''} = \mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{C}''}, \mathcal{F}^{\mathcal{A}''}\},$$

and the data $\mathbf{D}''$ has the following property:

Any homeomorphism $\mathcal{H}_0 : S^2 \to S^2$ which is the identity on $\mathcal{S}_0$, when restricted to the image $\mathcal{F}_*(\mathcal{C}_c)$ of any missing annulus $\mathcal{C}_c$, lifts under $\mathcal{F}^{\mathcal{C}}$ and $\mathcal{F}^{\mathcal{C}''}$ to a homeomorphism $\mathcal{H}_1^{\mathcal{C}} : \mathcal{C}_c \to \mathcal{C}_c$ which is isotopic to the identity rel $\partial\mathcal{C}_c$.

Furthermore, $F'' = F' \circ h$, where $h : S^2 \to S^2$ is a homeomorphism supported on the interiors of the annuli in $\mathcal{A}_0$, and the homeomorphism h can be chosen so as to depend only on $\mathcal{F}^{\mathcal{C}'}$ and not on $\mathcal{F}^{\mathcal{A}'}$.

Proof: Recall (Figure 3.1) that the annulus $\mathcal{F}_*(\mathcal{C}_c)$ is the union of

- the unique component $\mathcal{S}_0(x)$ of $\mathcal{S}_0$ whose boundary contains $\partial\mathcal{F}_*(\mathcal{C}_c)$, and
- finitely many disks in S^2.

Hence altering $\mathcal{H}_0$ on these disks will not change the isotopy class (rel boundary of $\mathcal{F}_*(\mathcal{C}_c)$) of the map $\mathcal{H}_0|\mathcal{F}_*(\mathcal{C}_c)$. Therefore it is enough to prove the result for *some* particular such homeomorphism $\mathcal{H}_0$ which is the identity on $\mathcal{S}_0$; for convenience we take $\mathcal{H}_0 = \mathrm{id}_{S^2}$.

Fix $c \in C$. The maps $\mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{C}'}$ when restricted to the annuli $\mathcal{C}_c = \mathcal{C}'_c$ are covering maps of the same degree. Each annulus $\mathcal{C}_c$ is an essential subannulus in the interior of some directed annulus $\mathcal{A}_{0,i(c)} \in \mathcal{A}_0$. Thus, $\mathcal{C}_c$ has a preferred positive boundary component $\partial^+\mathcal{C}_c$ which coincides with the negative boundary component of some annulus $\mathcal{A}_{1,j(c)} \in \mathcal{A}_1$ as well as a preferred negative boundary component $\partial^-\mathcal{C}_c$ (see Figure 4.1).

Let $\mathcal{H}_1^{\mathcal{C}_c} : \mathcal{C}_c \to \mathcal{C}'_c = \mathcal{C}_c$ be the particular lift of the identity map $\mathcal{H}_0$ on $\mathcal{F}_*(\mathcal{C}_c) = \mathcal{F}'_*(\mathcal{C}'_c)$ under $\mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{C}'}$ which is the identity on the *negative* boundary component of $\mathcal{C}_c$. There are two potential complications: (1) $\mathcal{H}_1^{\mathcal{C}_c}$ need not be the identity on the other positive boundary component of $\mathcal{C}_c$, and (2) even if it is, it need not be isotopic to the identity rel $\partial\mathcal{C}_c$.

We resolve these difficulties simultaneously as follows. First, note that the combinatorial class of F' is unchanged if F' is precomposed with a homeomorphism h isotopic to the identity rel Q. We will find such an h such that the identity map on the annuli $\mathcal{F}_*(\mathcal{C}_c),\ c \in C$ lifts under $F' \circ h$ to a homeomorphism $\mathcal{H}_1^{\mathcal{C}_c}$ which is isotopic to the identity rel boundary. The homeomorphism h will be supported on a neighborhood of $\cup_{c\in C}\partial^+\mathcal{C}_c$.

Fix $c \in C$. Choose arbitrarily points $x^\pm \in \partial^\pm\mathcal{C}_c$ and a curve $\alpha : [0,1] \to \mathcal{C}_c$ from x^- to x^+. Let $\widetilde{\alpha}$ be the lift of $(\mathcal{H}_0 \circ \mathcal{F}^{\mathcal{C}}|\mathcal{C}_c)(\alpha)$ under $\mathcal{F}^{\mathcal{C}'}$ based at x^- (recall that $\mathcal{H}_0 = \mathrm{id}_{S^2}$). Then the terminal endpoint $\tilde{x}^+$ is a point on $\partial^+(\mathcal{C}_c)$ lying over $(\mathcal{F}^{\mathcal{C}}|\mathcal{C}_c)(x^+) = (\mathcal{F}^{\mathcal{C}'}|\mathcal{C}_c)(x^+)$, since we have already arranged that the boundary values of $\mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{C}'}$ are the same.

Then there exists an "unwinding" homeomorphism $h_c : \mathcal{C}_c \cup \mathcal{A}_{1,j(c)} \to \mathcal{C}_c \cup \mathcal{A}_{1,j(c)}$ with the following properties (see Figure 4.1):

1. $h_c = \mathrm{id}$ on a neighborhood of the boundary of the annulus $\mathcal{C}_c \cup \mathcal{A}_{1,j(c)}$.
2. h_c is isotopic to the identity rel $\partial(\mathcal{C}_c \cup \mathcal{A}_{1,j(c)})$
3. $\mathcal{F}^{\mathcal{C}'} \circ h|\partial\mathcal{C}_c = \mathcal{F}^{\mathcal{C}'}$
4. $h_c(\alpha)$ is homotopic rel endpoints via a homotopy in $\mathcal{C}_c$ to the curve $\widetilde{\alpha}$.

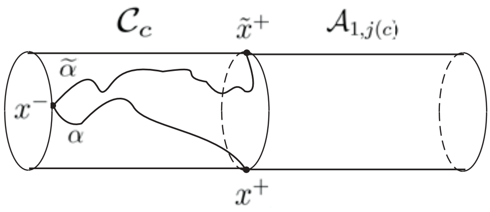

Fig. 4.1. The homeomorphism h_c sends α to a curve homotopic to $\widetilde{\alpha}$.

The homeomorphism h_c can be obtained by extending the unique deck transformation of $\mathcal{F}^{\mathcal{C}}|\partial^+(\mathcal{C}_c)$ sending x^+ to $\widetilde{x}^+$ to a neighborhood of $\partial^+(\mathcal{C}_c)$ so as to arrange that $h_c(\alpha)$ is homotopic to $\widetilde{\alpha}$.

Consider now the composition $(\mathcal{F}^{\mathcal{C}'}|\mathcal{C}_c) \circ h_c$. The lift of $(\mathcal{H}_0 \circ \mathcal{F}^{\mathcal{C}}|\mathcal{C}_c)(\alpha)$ based at x^- under this map is then $h_c^{-1} \circ (\mathcal{F}^{\mathcal{C}'}|\mathcal{C}_c)^{-1}(\alpha) = h_c^{-1}(\widetilde{\alpha})$, and by property (4) above, we have that this curve is homotopic in $\mathcal{C}_c$ to α itself. Hence, the identity map $\mathcal{H}_0$ lifts under $\mathcal{F}^{\mathcal{C}}$ and $\mathcal{F}^{\mathcal{C}'} \circ h_c$ to a map denoted $\mathcal{H}_1|\mathcal{C}_c$ which on $\mathcal{C}_c$ sends the curve α to a curve which is homotopic to α fixing endpoints. From this it follows that $\mathcal{H}_1|\mathcal{C}_c$ is isotopic to the identity through maps fixing $\partial\mathcal{C}_c$.

We then let $h : S^2 \to S^2$ be the homeomorphism obtained by applying h_c on each of the disjoint annuli $\mathcal{C}_c \cup \mathcal{A}_{1,j(c)}$. Then $F'' = F' \circ h$ is combinatorially equivalent to F' and the arguments in the previous paragraph show that desired properties hold.

□

Definition 4.11 (Simple form). *A set* $\mathbf{D}''$ *of amalgamating data is said to be in* simple form *with respect to a basepoint* $\mathbf{D}$ *if*

1. $\mathbf{D}, \mathbf{D}''$ *are amalgamating data with respect to the same structure data,*
2. *the boundary values* ρ, ρ'' *are the same,*
3. *the sphere maps* $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}''}$ *are equivalent over* $\mathbf{T}$ *via homeomorphisms* $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ *with* $\mathcal{H}_0^{\mathcal{S}} = \mathrm{id}$*, and*
4. *any homeomorphism* $\mathcal{H}_0 : S^2 \to S^2$ *which is the identity on* $\mathcal{S}_0$*, when restricted to the image* $\mathcal{F}_*(\mathcal{C}_c)$ *of any missing annulus* $\mathcal{C}_c$*, lifts under* $\mathcal{F}^{\mathcal{C}}$ *and* $\mathcal{F}^{\mathcal{C}''}$ *to a homeomorphism* $\mathcal{H}_1^{\mathcal{C}} : \mathcal{C}_c \to \mathcal{C}_c$ *which is isotopic to the identity rel* $\partial\mathcal{C}_c$*.*

Summarizing: given fixed structure data and a basepoint $\mathbf{D}$ of amalgamating data, the set of classes of amalgams which can be produced using any other amalgamating data with respect to this structure data coincides with the a priori smaller set of classes of amalgams which can be produced using amalgamating data which is in simple form with respect to **D**.

Moreover, the last conclusion in the Unwinding Trick shows that simple form is preserved under varying just the annulus maps:

Corollary 4.12. *Suppose* $\mathbf{D}' = \{\mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{C}'}, \mathcal{F}^{\mathcal{A}'}\}$ *is in simple form with respect to* $\mathbf{D}$*. Then* $\{\mathcal{F}^{\mathcal{S}'}, \mathcal{F}^{\mathcal{C}'}, \mathcal{F}^{\mathcal{A}''}\}$ *is also in simple form with respect to* $\mathbf{D}$*, where* $\mathcal{F}^{\mathcal{A}''}$ *is any other choice of annulus maps.*

4.4.4 Conclusion of proof of Uniqueness Theorem

Let $F = F(\mathbf{D})$ be the basepoint and $F' = F(\mathbf{D}')$ be in simple form with respect to F. Recall that the simple form implies that the pair $\mathcal{H}_0^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{S}}$ yielding a combinatorial equivalence between $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}'}$ may be taken with $\mathcal{H}_0^{\mathcal{S}} = \mathrm{id}_{\mathcal{S}_0}$. Let $\mathcal{H}_0^{\mathcal{A}}, \mathcal{H}_1^{\mathcal{A}}$ be a pair of homeomorphisms giving an equivalence between the annulus maps. We must show that F and F' are combinatorially equivalent.

Let $\mathcal{H}_0 : S^2 \to S^2$ be given by

$$\mathcal{H}_0 = \begin{cases} \mathcal{H}_0^{\mathcal{S}} = \mathrm{id}_{\mathcal{S}_0} & \text{on } \mathcal{S}_0 \\ \mathcal{H}_0^{\mathcal{A}} & \text{on } \mathcal{A}_0 \end{cases}$$

We now lift $\mathcal{H}_0$ under F and F' to a homeomorphism $\mathcal{H}_1 : S^2 \to S^2$ in the following way.

On the missing disks $\mathcal{U}, \mathcal{U}'$ there is no obstruction to lifting $\mathcal{H}_0$ to a map which is the identity on the boundary since, for each $u \in U$, the maps $\mathcal{F}^{\mathcal{U}}, \mathcal{F}^{\mathcal{U}'}$ are either homeomorphisms, or are both ramified coverings of disks with one critical point and with the same degree. We denote the collection of such lifts by $\mathcal{H}_1^{\mathcal{U}}$.

On the missing annuli $\mathcal{C}$, $\mathcal{H}_0$ lifts to a homeomorphism $\mathcal{H}_1^{\mathcal{C}}$ which is isotopic to the identity rel boundary, by the definition of simple form [4.11] and the Unwinding Trick [Prop. 4.10].

On the subsets $\mathcal{S}_1, \mathcal{A}_1$ we use $\mathcal{H}_1^{\mathcal{S}}, \mathcal{H}_1^{\mathcal{A}}$ as in the definitions [4.3, 4.4] of combinatorial equivalence. By definition, these are the identity on their boundaries.

Finally, we set

$$\mathcal{H}_1 = \begin{cases} \mathcal{H}_1^{\mathcal{S}} & \text{on } \mathcal{S}_1 \\ \mathcal{H}_1^{\mathcal{A}} & \text{on } \mathcal{A}_1 \\ \mathcal{H}_1^{\mathcal{C}} & \text{on } \mathcal{C} \\ \mathcal{H}_1^{\mathcal{U}} & \text{on } \mathcal{U}. \end{cases}$$

Thus, $\mathcal{H}_1$ is a continuous lift of $\mathcal{H}_0$. We now verify that $\mathcal{H}_0$ is isotopic to $\mathcal{H}_1$ through maps which fix Q. In fact, we shall show that $\mathcal{H}_0, \mathcal{H}_1$ are isotopic rel $\partial\mathcal{S}_1 \cup Q$.

On $\mathcal{U}$, the maps $\mathcal{H}_0, \mathcal{H}_1$ agree on $\partial\mathcal{U}$ and, on each component $\mathcal{U}_u$, and they fix any (necessarily unique, by Axiom A-7, 4) point $q \in \mathcal{U}_u$, if such a point q exists. By the Alexander Trick, there is an isotopy rel $\partial\mathcal{U}_u \cup \{q\}$ joining $\mathcal{H}_0|\mathcal{U}_u$ to $\mathcal{H}_1|\mathcal{U}_u$.

On $\mathcal{S}_1$, the maps $\mathcal{H}_0, \mathcal{H}_1$ are isotopic rel $\partial\mathcal{S}_1 \cup \mathcal{Q}$ by the definition [4.3] of combinatorial equivalence of sphere maps.

Fix now a component $\mathcal{A}_{0,i}$ of $\mathcal{A}_0$. If $\mathcal{A}_{0,i}$ contains no components of $\mathcal{C}$, then $\mathcal{H}_0$ and $\mathcal{H}_1$ are isotopic rel $\partial\mathcal{A}_{0,i}$ by the definition of combinatorial equivalence of annulus maps [4.4]. Otherwise, the map $\mathcal{H}_1$ restricted to each component $\mathcal{C}_c$ of $\mathcal{C} \cap \mathcal{A}_{0,i}$ is isotopic to the identity rel $\partial\mathcal{C}_c$, by the construction of $\mathcal{H}_1^{\mathcal{C}}$ and the Unwinding Trick. By the definition of combinatorial equivalence of annulus maps, the extension $\overline{\mathcal{H}}_1^{\mathcal{A}}$ of $\mathcal{H}_1^{\mathcal{A}}$ over each component of $\mathcal{C} \cap \mathcal{A}_{0,i}$ by the identity map yields a map which is isotopic to $\mathcal{H}_0^{\mathcal{A}}$ rel boundary $\mathcal{A}_{0,i}$. It follows that the map $\mathcal{H}_1|A_{0,i}$ is isotopic to $\mathcal{H}_0^{\mathcal{A}}|\mathcal{A}_{0,i}$. This holds over all the disjoint annuli $\mathcal{A}_{0,i}$ in $\mathcal{A}_0$, hence $\mathcal{H}_0|\mathcal{A}_0$ is isotopic to $\mathcal{H}_1|\mathcal{A}_0$ rel boundary $\mathcal{A}_0$.

Together, the isotopies over $\mathcal{U}$, $\mathcal{S}_1$, and $\mathcal{A}_0$ piece together to yield an isotopy between $\mathcal{H}_0$ and $\mathcal{H}_1$ through maps fixing $\partial\mathcal{S}_1 \cup Q$. $\square$

5

Decomposition

5.1 Statement of Decomposition Theorem

In this section, we give a converse, decomposition procedure for amalgamation.

We say that two pairs, each consisting of an embellished map of spheres and an invariant multicurve, are *combinatorially equivalent* if there is a combinatorial equivalence between them sending one multicurve to the other, up to isotopy.

The main result of this section is

Theorem 5.1 (Decomposition). *Let $G : (S^2, Q) \to (S^2, Q)$ be an embellished map of a sphere to itself and Γ a nonempty invariant multicurve. Then there is a homeomorphism $h : (S^2, Q) \to (S^2, Q)$ isotopic to the identity rel Q such that $F = G \circ h$ is an amalgam with annuli $\mathcal{A}_0$ whose core curves are homotopic to Γ. The data which expresses F as an amalgam has the following properties:*

1. **Boundary values.** *The boundary values are determined by the pair $(F, \mathcal{A}_0)$.*
2. **Structure data.**
   a) **Mapping tree.** *The mapping tree $\mathbf{T}$ depends only on the equivalence class of pair (G, Γ). The set of spheres $\mathcal{S} = S^2 \times X$, where X is the set of X-vertices of $\mathbf{T}$.*
   b) **Set** $\mathcal{Q}$. *The map $F : (S^2, Q) \to (S^2, Q)$ and the isotopy class of Γ in $S^2 - Q$ determine uniquely the set $\mathcal{Q}$ which is canonically identified with Q.*
   c) **Set** $\mathcal{Z}$. *The subset $\mathcal{Z}$ is in bijective correspondence with the set of boundary components of $\mathcal{S}_0$. The locations of the points $\mathcal{Z}$, with the following exceptions, depend only on F and the isotopy class of Γ.*

      The exceptions arise as those z corresponding to boundary components b for which there exists no component $\mathcal{U}_u$ of $\mathcal{U}$ which intersects $\mathcal{Q}$ and whose boundary maps to b. Such components b are necessarily either

K.M. Pilgrim: LNM 1827, pp. 69–77, 2003.

totally invariant under F, or have all preimages other than possibly itself mapping by degree one.

We write

$$\mathcal{Z} = \mathcal{Z}^{exc} \sqcup \mathcal{Z}^{det}$$

for the exceptional and determined elements of $\mathcal{Z}$.

d) **Restrictions and local degrees on $\mathcal{Q}$.** *The mapping scheme $(\mathcal{Q} \sqcup \mathcal{Z}, \mathcal{F}, \deg_{\mathcal{F}})$ is uniquely determined up to isomorphism by F and the isotopy class of Γ in the following sense: a choice of exceptional points $\mathcal{Z}^{exc}$, together with F and $\mathcal{Q}$, determines uniquely a mapping scheme with underlying set $\mathcal{Q} \sqcup \mathcal{Z}$. A different choice of exceptional points $\mathcal{Z}' = \mathcal{Z}^{exc'} \sqcup \mathcal{Z}^{det}$ determines a canonically defined isomorphism $\mathcal{Q} \sqcup \mathcal{Z} \to \mathcal{Q} \sqcup \mathcal{Z}'$ of the corresponding mapping schemes which is the identity off the exceptional set.*
However, the set $\mathcal{Q} \cap \mathcal{F}^{-1}(\mathcal{Z})$ is canonically determined by F and the isotopy class of Γ.

e) **Critical gluing.** *The critical gluing $\kappa : \mathcal{Q} \cap \mathcal{F}^{-1}(\mathcal{Z}) \to \mathcal{Q}$ is canonically determined by F and the isotopy class of Γ.*

3. **Annulus maps. Missing annulus maps. Missing disk maps.** *The annulus maps $\mathcal{F}^{\mathcal{A}}$, missing annulus maps $\mathcal{F}^{\mathcal{C}}$, and missing disk maps $\mathcal{F}^{\mathcal{U}}$ depend only on F and $\mathcal{A}_0$.*

4. **Sphere maps.** *The (non-proper) restriction of the sphere maps $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ to the subset $\mathcal{S}_0$,*

$$\mathcal{F}^{\mathcal{S}}|\mathcal{S}_0 : \mathcal{S}_1 \sqcup \mathcal{U} \to \mathcal{S}$$

depends only on F and $\mathcal{A}_0$.

Remark. To see how elements of $\mathcal{Z}^{exc}$ can arise, consider a decomposition procedure inverse to mating. Here, $\Gamma = \{\gamma\}$ consists of the "equator" and is totally invariant. Cut along γ to obtain two disks $D_\pm$ viewed as disjoint subsets of two different spheres $S^2_\pm$, each of which is identified with the domain S^2 of G. There is still no natural way to extend the map $G|_{D_+}$ over the disk $S^2 - D_+$–since the boundary ∂D_+ maps by itself with degree (typically) ≥ 2, any extension over $S^2 - D_\pm$ must introduce branching, and the locations of the branch points and values must be specified in advance.

The embellished map of spheres $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ is called *a decomposition of G along Γ*. When we wish to emphasize the dependence of the map $\mathcal{F}$ on F, we say that $\mathcal{F}$ is *produced from F*. Technically, this is not quite right: the construction of $\mathcal{F}$ from F will involve some choices; the Decomposition Theorem [5.1] records what those choices are. In Section 6 we analyze the dependence of the class of $\mathcal{F}$ on these choices.

The map F will have a special form, called *standard form with respect to Γ*, which is defined in the next section.

5.2 Standard form with respect to a multicurve

Definition 5.2 (Standard form). *An embellished branched covering $F : (S^2, Q) \to (S^2, Q)$ is said to be in* standard form *with respect to a multicurve Γ if there exist a set $A_0 = \{\mathcal{A}_{a_{0,1}}, ..., \mathcal{A}_{a_{0,n_0}}\}$ of $n_0 = |\Gamma|$ disjoint, closed annuli such that*

1. *each $\gamma_i \in \Gamma$ is homotopic to a core curve of some annulus in A_0;*
2. *if $\mathcal{A}_0 = \cup_{i=1}^{n_0} \mathcal{A}_{a_{0,i}}$, and if $\mathcal{A}_1$ denotes the components of preimages of $\mathcal{A}_0$ which are both essential and nonperipheral in $S^2 - Q$, then $\mathcal{A}_1 \subset \mathcal{A}_0$, and $\partial \mathcal{A}_1 \supset \partial \mathcal{A}_0$;*
3. *In suitable local coordinates, the map $F|\partial \mathcal{A}_1 : \partial \mathcal{A}_1 \to \partial \mathcal{A}_0$ is given by $t \mapsto \delta^+ t$ modulo 1, where δ^+ is the (positive) degree of F on the given boundary component.*

The standard form F is obtained by precomposing G with a homeomorphism isotopic to the identity rel Q such that suitably thickened elements of Γ_0 form the required set of annuli $\mathcal{A}_0$.

Proposition 5.3 (Standard form). *Let $G : (S^2, Q) \to (S^2, Q)$ be an embellished branched covering and Γ a G-invariant multicurve. Then there is a homeomorphism $h : (S^2, Q) \to (S^2, Q)$ isotopic to the identity rel Q such that $F = G \circ h$ is in standard form with respect to Γ.*

Proof: Thicken the components of Γ to disjoint, closed annuli $\mathcal{A}_0$. By invariance of Γ, each essential, nonperipheral preimage of $\partial \mathcal{A}_0$ is homotopic rel Q to an element of $\partial \mathcal{A}_0$, and each component of $\partial \mathcal{A}_0$ is homotopic to at least one such preimage. For fixed i, consider the collection of all preimages of $\partial \mathcal{A}_0$ homotopic to γ_i in $S^2 - Q$. There are exactly two outermost such preimages, say b_0^-, b_0^+, where outermost means that neither b_0^- nor b_0^+ separates a pair of curves in this collection. Together, b_0^- and b_0^+ bound a closed annulus $\mathcal{A}_{a_{0,i}}^{\#}$ which is isotopic in $S^2 - Q$ to the annulus $\mathcal{A}_{a_{0,i}}$. Let $h_1 : (S^2, Q) \to (S^2, Q)$ be a homeomorphism of S^2 isotopic to the identity rel Q and sending each annulus $\mathcal{A}_{a_{0,i}}$ to the annulus $\mathcal{A}_{a_{0,i}}^{\#}$; such an h_1 exists by the Tameness Restriction # 3 (§1.9). Then $G \circ h_1$ satisfies the first two conditions in the definition [5.2] of standard form, and the third is achieved by precomposing $G \circ h_1$ with another suitable tame homeomorphism h_2 isotopic to the identity rel Q to obtain $F = G \circ h_1 \circ h_2 = G \circ h$. $\square$

5.3 Maps in standard forms are amalgams

Next, we claim that F is an amalgam.

Proposition 5.4 (Maps in standard form are amalgams). *If F is in standard form with respect to a multicurve Γ, then F is an amalgam whose core curves are isotopic to Γ.*

Proof: We proceed by brutally verifying the axioms one by one. For an overview of this process, see the discussion at the beginning of §2.

Definition of $\mathcal{S}_0, \mathcal{S}_1, \mathcal{A}_0, \mathcal{A}_1, \mathcal{U}, \mathcal{C}$. We set $\mathcal{A}_0 = \sqcup_i \mathcal{A}_{a_{0,i}}$ and $\mathcal{A}_1 = \sqcup_j \mathcal{A}_{a_{1,j}}$ as before. Let $\mathcal{C} = \overline{\mathcal{A}_0 - \mathcal{A}_1}$ and index the connected components $\mathcal{C}_c$ by elements $c \in C$. Let $\mathcal{S}_0 = \overline{S^2 - \mathcal{A}_0}$, and let X denote the set of connected components of $\mathcal{S}_0$. Set $\mathcal{S}_0(x)$ to be the xth connected component of $\mathcal{S}_0$. Finally, let $\mathcal{S}_1$ denote the union of those connected components of $F^{-1}(\mathcal{S}_0)$ satisfying the following two properties: it is contained in some component $\mathcal{S}_0(x)$, and its boundary contains $\partial\mathcal{S}_0(x)$.

Lemma 5.5. *The inclusion map $\mathcal{S}_1 \hookrightarrow \mathcal{S}_0$ induces a bijection between the components of $\mathcal{S}_1$ and those of $\mathcal{S}_0$.*

Proof: Injectivity is clear, and surjectivity follows from F-invariance and the extremal property (2) of the annuli in the definition of standard form [5.2]. □This lemma allows us to write $\mathcal{S}_1(x)$ for the unique component of $\mathcal{S}_1$ which is contained in $\mathcal{S}_0(x)$, and whose boundary contains $\partial\mathcal{S}_0(x)$.

Finally, we set $\mathcal{U} = \overline{\mathcal{S}_0 - \mathcal{S}_1}$, and index components $\mathcal{U}_u$ of $\mathcal{U}$ by the set U.

Definition of tree T. Let B_1 denote the set of boundary components of $\mathcal{A}_1$, and B_0 those of $\mathcal{A}_0$. Observe that $B_0 \subset B_1$ and that $B_1 - B_0$ consists precisely of the set of boundary components of $\mathcal{C}$.

We set T to be the tree whose set of vertices is $X \sqcup U \sqcup B_1$ and whose edges are defined as follows:

- First, given $u \in U$, we have $\partial\mathcal{U}_u \subset \mathcal{S}_1(x)$ for a unique $x \in X$. We join x and u by an edge.
- Next, given $c \in C$, we have $\partial\mathcal{C}_c = b_1^- \sqcup b_1^+$ for components b_1^-, b_1^+ of $\partial\mathcal{A}_1$. We join b_1^- and b_1^+ by an edge and label this edge with c.
- Then, given $x \in X$, $\partial\mathcal{S}_0(x) \subset \partial\mathcal{A}_0 = B_0$. We join each $b_0 \in \partial\mathcal{S}_0(x)$ to x by an edge.
- Finally, we join two elements of B_1 by an edge if the corresponding boundary components of $\mathcal{A}_1$ together bound an annulus $\mathcal{A}_{a_{1,j}}$ in $\mathcal{A}_1$. We label this edge with $a_{1,j}$.

With this definition, it is straightforward to verify that the Valence and Adjacency Axioms (Axioms 1,2) are satisfied.

Having now in hand the tree T, we set A_0, A_1 as in the definition of a mapping tree given in §2.1. We will need

Lemma 5.6. *For each annulus $A_{a_{0,i}} \in \mathcal{A}_0$, $F|\mathcal{A}_{a_{0,i}}$ is an unbranched covering, and the image $F(\mathcal{A}_{a_{0,i}})$ is disjoint from $F(Q)$. Hence, for $c \in \mathcal{C}$, $F|\mathcal{C}_c$ is a proper unbranched covering, and the image $F(\mathcal{C}_c)$ is disjoint from $F(Q)$.*

Note that this is not obvious–although $\partial\mathcal{A}_{a_{0,i}}$ is made up of preimages, $\mathcal{A}_{a_{0,i}}$ need not a priori be a preimage of an annulus.

Proof: Let b_0^-, b_0^+ denote the boundary components of $\mathcal{A}_{a_{0,i}}$. Since $\mathcal{A}_{a_{0,i}} \cap Q = \emptyset$, and Q is an embellishment, $\mathcal{A}_{a_{0,i}}$ contains no critical points. It follows that $F(b_0^-) \neq F(b_0^+)$. So $F(b_0^-) \cup F(b_0^+)$ bounds an annulus, call it B. Equip b_0^-, b_0^+ with the orientation induced from $\mathcal{A}_{a_{0,i}}$, and push these orientations down to orientations on $F(b_0^-), F(b_0^+)$. The extremal condition (2) on $\partial A_{0,i}$ in the definition of standard form [5.2] implies that these orientations agree with the orientations induced by B. Now suppose B contains a point $p \in F(Q)$. Choose arbitrarily $o \in \partial B$ and a path $\alpha : [0,1] \to B$ from o to p whose interior avoids $F(Q)$. Lifting α based at $\tilde{o} \in \partial \mathcal{A}_{a_{0,i}}$, we obtain a path $\widetilde{\alpha}$ in $A_{0,i}$ terminating at a point of $F^{-1}(F(Q)) = Q$. This is impossible since $\mathcal{A}_{a_{0,i}} \cap Q = \emptyset$ by construction. So $B \cap F(Q) = \emptyset$. It follows that $F|\mathcal{A}_{a_{0,i}}$ is proper. Since a proper local homeomorphism is a covering map, the first statement of the Lemma is proved. The second follows immediately since boundary components of $\mathcal{C}_c$ are preimages of curves. $\square$

A similar argument, using the fact that boundary components of $\mathcal{U}$ are inessential or peripheral rel Q and the fact that $Q = F^{-1}(F(Q))$, establishes

Lemma 5.7. *For each $\mathcal{U}_u \in \mathcal{U}$, $F|\mathcal{U}_u$ is a proper branched covering which is ramified at a maximum of one point in the interior of $\mathcal{U}_u$. Furthermore, the image $F(\mathcal{U}_u)$ contains at most one point of $F(Q)$.*

Definition of map $f : T \to T$. We first define T on the vertices of T.

Using Lemma 5.5, for $x \in X$ we set $f(x) = y$ if $F(\mathcal{S}_1(x)) = \mathcal{S}_0(y)$. Consider now $b_1 \in B_1$. By definition b_1 is a boundary component of $\mathcal{A}_1$ and hence $F(b_1)$ is a boundary component b_0 of $\mathcal{A}_0$; we set $f(b_1) = b_0$. Next consider a component $\mathcal{U}_u$ of $\mathcal{U}$. It is a disk bounded by a single boundary component β of some component $\mathcal{S}_1(x)$. Thus $F(\beta)$ is a boundary component of $\mathcal{S}_0(f(x))$, i.e. is some $b_0 \in B_0$. We then set $f(u) = b_0$.

It remains to extend f over edges of T. It is straightforward to verify that the images of any pair of distinct, adjacent vertices is a pair of distinct vertices. Since any two vertices in a tree are joined by a unique path, there is an essentially unique way to extend f over the edges of T such that f is continuous and the restriction of f to edges is a homeomorphism; this verifies Axioms A–3, #s 1,2,3. The remaining axioms A–3 #s 4,5 follow from Lemma 5.6.

Degree function. We set the degree function deg in the obvious way, namely

- $\deg(x) = \deg(F|\mathcal{S}_1(x) \to \mathcal{S}_0(f(x)))$
- $\deg(b_1) = \deg(F|b_1 : b_1 \to F(b_1))$
- $\deg(u) = \deg(F|\beta : \beta \to F(\beta))$ where $\beta = \partial \mathcal{U}_u$.

The first Local Homogeneity Axiom (Axiom A-4 #1(a)) then follows from the fact that $F|\mathcal{S}_1(x) : \mathcal{S}_1(x) \to \mathcal{S}_0(f(x))$ is a ramified covering map, and the second (Axiom A-4 #1(b)) from Lemma 5.6. The Global Homogeneity Axioms (Axioms A–4 # 2) follow immediately from the definition of the degree function, Lemmas 5.6 and 5.7, and the fact that F is a branched covering of degree d.

Orientations. We equip the annular edges $a_{0,i}$ of A_0 with arbitrarily chosen orientations, and use this to define the signed degree function δ.

Sphere maps. For each $x \in X$, let $\mathcal{S}(x) = S^2$ be the domain (and range) of F. *We emphasize that $\mathcal{S}(x)$ is not constructed as an abstract manifold by gluing disks onto components of* $\mathcal{S}_0(x)$. Let $\mathcal{B}(x) = \overline{\mathcal{S}(x) - \mathcal{S}_0(x)}$ and set $\mathcal{B} = \cup_x \mathcal{B}(x)$; this is a collection of disks.

Equip each disk in $\mathcal{B}$ with a preferred center point z in its interior, so that $\mathcal{B}_z$ refers to the unique disk with center z. The points in $\mathcal{Z}$ are not quite chosen arbitrarily: if $b = \partial\mathcal{B}_z = \partial F(\mathcal{U}_u)$ then by Lemma 5.7 above, the disk $F(\mathcal{U}_u) = \mathcal{B}_z$ contains at most one point of $F(Q)$. If such a point exists, we choose it to be the center, and denote the resulting collection of determined points by $\mathcal{Z}^{det}$. If no such point exists, we choose arbitrarily a center point z of $\mathcal{F}(\mathcal{U}_u)$, and denote the resulting collection of exceptional points whose locations are not canonically determined by $\mathcal{Z}^{exc}$.

Set $\mathcal{Z} = \mathcal{Z}^{exc} \sqcup \mathcal{Z}^{det}$, $\mathcal{Q} = Q$, $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$, and $\mathcal{S} = \sqcup_x \mathcal{S}(x)$. Then $\mathcal{S} = S^2 \times X$ is a union of spheres and $\mathcal{S}_0, \mathcal{S}_1$ may be canonically identified with subsets of $\mathcal{S}$.

At this point we have a proper map $\mathcal{F}^{\mathcal{S}}|\mathcal{S}_1 : \mathcal{S}_1 \to \mathcal{S}_0$ given by the restriction of F to $\mathcal{S}_1$. We need to extend this map to $\mathcal{S}$. In fact, we have a bit more: for each $x \in X$, the map F restricts to a (non-proper) map

$$\mathcal{F}^{\mathcal{S}}|\mathcal{S}_0(x) = F|\mathcal{S}_0(x) : \mathcal{S}_0(x) \to \mathcal{S}(f(x)),$$

i.e. in addition $\mathcal{F}^{\mathcal{S}}$ is already defined on $\mathcal{U}$. It remains therefore to extend $\mathcal{F}^{\mathcal{S}}$ over $\mathcal{B}$.

Given $z \in \mathcal{Z}$, $\partial\mathcal{B}_z \subset \partial\mathcal{S}_0 \subset \partial\mathcal{S}_1$ and so $F(\partial\mathcal{B}_z) \subset \partial\mathcal{S}_0 = \partial\mathcal{B}$ is a boundary component of another disk in $\mathcal{B}$ whose center we denote (abusing notation) as $\mathcal{F}^{\mathcal{S}}(z)$. We then "radially" extend the map $\mathcal{F}^{\mathcal{S}} : \partial\mathcal{B}_z \to \partial\mathcal{B}_{\mathcal{F}^{\mathcal{S}}(z)}$ to a map $(\mathcal{B}_z, z) \to (\mathcal{B}_{\mathcal{F}^{\mathcal{S}}(z)}, \mathcal{F}^{\mathcal{S}}(z))$ which is a covering ramified only at z and of local degree equal to that of the degree on the boundary.

Since the modification of G to F in Proposition 5.3 can be taken to be away from Q, and $\mathcal{F}^{\mathcal{S}} = F$ near any accumulation points of $\mathcal{Q} = Q$, the Tameness Axiom A–8 is satisfied. The Standard Form Axioms A–7 are also satisfied; verification of A–7#4 follows from F-invariance of Γ_0. The projection map $\pi^{\mathcal{S}}$ can be then defined as follows: project the disks $\mathcal{B}$ to B_0, the disks $\mathcal{U}$ to U, and extend continuously to obtain a projection of $\mathcal{S}$ onto the mapping tree T. The Covering Axioms A–9 #s 1–4 are then satisfied.

At this point we need to verify the Embellished Axiom A–6.

Lemma 5.8. *The map*

$$\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$$

is an embellished branched covering.

Proof: The set $\mathcal{Y}$ is clearly closed, forward-invariant, and contains the critical points of $\mathcal{F}^{\mathcal{S}}$. Thus it suffices to prove that $(\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{Y}) = \mathcal{Y}$. The containment $\supset$ is trivial. To prove the other direction, let $y \in (\mathcal{F}^{\mathcal{S}})(\mathcal{Y})$.
Case 1: $y \in \mathcal{Q}$**.** Since $\mathcal{Z}$ is forward-invariant under $\mathcal{F}^{\mathcal{S}}$, and since $\mathcal{F}^{\mathcal{S}}(\mathcal{Q} \cap \mathcal{U}) \subset \mathcal{Z}$, we must have that $y \in \mathcal{F}^{\mathcal{S}}(\mathcal{Q} - \mathcal{Q} \cap \mathcal{U})$. Hence $y \in \mathcal{F}^{\mathcal{S}}(\mathcal{S}_1 \cap \mathcal{Q})$. By construction $\mathcal{Q} = Q$ and $\mathcal{F}^{\mathcal{S}}|\mathcal{S}_1 = F|\mathcal{S}_1$, hence $y \in F(\mathcal{S}_1 \cap Q)$. Thus $y \in F(Q)$. Also, $y \in \mathcal{S}_0$. By construction $(\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{S}_0) = \mathcal{S}_1$, and so $(\mathcal{F}^{\mathcal{S}})^{-1}(\mathcal{S}_0) = \mathcal{S}_1$. Hence $(\mathcal{F}^{\mathcal{S}})^{-1}(y) \subset F^{-1}(F(Q)) = Q = \mathcal{Q}$, since $F : (S^2, Q) \to (S^2, Q)$ is embellished. Therefore $(\mathcal{F}^{\mathcal{S}})^{-1}(y) \subset \mathcal{Q} \subset \mathcal{Y}$.
Case 2: $y \in \mathcal{Z}$**.** Then any preimage p of y under $\mathcal{F}^{\mathcal{S}}$ which is not a priori in $\mathcal{Y}$ lies in a disc $\mathcal{U}_u$ for which $\mathcal{U}_u \cap \mathcal{Q} = \emptyset$.

Now, if $y = \mathcal{F}^{\mathcal{S}}(q)$, some $q \in \mathcal{Q}$, then $F(\mathcal{U}_u)$ must contain a point of $F(Q)$. Since $F : (S^2, Q) \to (S^2, Q)$ is embellished, $F^{-1}(F(Q)) = Q$, and so $\mathcal{U}_u$ must contain a point of Q. Hence we may assume that $y \notin \mathcal{F}^{\mathcal{S}}(\mathcal{Q})$.

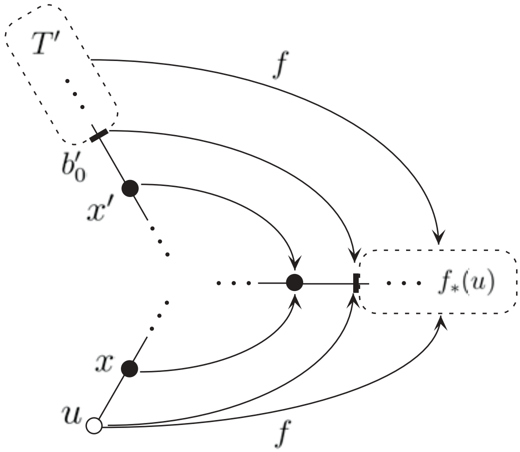

Fig. 5.1. If $F(\mathcal{U}_u)$ contains no points of $F(Q)$, then $f : T' \to f_*(u)$ is a homeomorphism.

In this case, there exists $z \in \mathcal{Z}$ such that $\mathcal{F}^{\mathcal{S}}(z) = y$. In terms of the original map F, this corresponds to the following situation: there exist $u \in U$, $x, x' \in X$, and $b_0' \in B_0$ such that $f(x) = f(x')$, $f(u) = f(b_0')$, $u \in S_1(x)$, and $b_0' \in S_0(x')$; see Figure 5.1. We claim that the image subtree $f_*(u)$ contains the image under f of at least one x vertex corresponding to a component $\mathcal{S}_0(x)$ for which $\mathcal{S}_0(x) \cap Q \neq \emptyset$. To see this, suppose otherwise. Then the disk $F(\mathcal{U}_u)$ contains in particular no critical values of F, hence all components of its preimages are disks mapping homeomorphically under F. In particular, the disk D in S^2 bounded by the component of $\partial\mathcal{S}_0(x')$ lying above b_0' maps homeomorphically to $F(\mathcal{U}_u)$ and lies over a corresponding maximal subtree T' bounded by b_0' mapping homeomorphically to $f_*(u)$ under f. By Covering

Axiom A–9 #4, the disk D must contain elements of Q. Hence $F(D) = F(\mathcal{U}_u)$ must contain points of $F(Q)$, and we are in the case treated by the previous paragraph. $\square$

Annulus maps. Set

$$\mathcal{F}^{\mathcal{A}} = \mathcal{F}|\mathcal{A}_1 : \mathcal{A}_1 \to \mathcal{A}_0.$$

Choose arbitrarily a direction-preserving identification of $\mathcal{A}_0$ with $A \times A_0$ so that the map on $\mathcal{F} : \partial\mathcal{A}_0 \to \partial\mathcal{A}_0$ is given as in the Standard form axiom A–12 for annulus maps. For each annular edge $a_{0,i} \in A_0$, choose a direction-preserving identification of $a_{0,i}$ with $[-1, 1]$. To define a projection, let $\pi : \mathcal{A}_0 \to \sqcup_{A_0} a_0$ be the map which forgets the radial coordinate.

The Subannuli Axiom A–10 follows since each annulus in $\mathcal{A}_1$ is essential in $\mathcal{A}_0$. The Covering Axiom A–11 follows easily as well.

Topological gluing data. Since, as a subset of the original sphere S^2 we have $\partial\mathcal{S}_0 = \partial\mathcal{A}_0$, we have a preferred topological gluing $\rho : \partial\mathcal{B} \to \partial\mathcal{A}_0$ giving a conjugacy between boundary values.

Critical gluing data. Suppose $q \in \mathcal{Q}$ maps to $z \in \mathcal{Z}$ under $\mathcal{F}^{\mathcal{S}}$. Since $\mathcal{Q} = Q$ and $F(Q) \subset Q$, we obtain a correspondence $q \mapsto \mathcal{F}^{\mathcal{S}}(q) = z \leftrightarrow F(q) \in \mathcal{Q}$. In this way, we the critical gluing data

$$\kappa : \mathcal{Q} \cap \mathcal{F}^{-1}(\mathcal{Z}) \to \mathcal{Q}$$

defined by $\kappa(q) = F(q)$ for $q \in \mathcal{Q} \cap \mathcal{F}^{-1}(\mathcal{Z})$ is uniquely determined by $\mathcal{F}$. The resulting mapping scheme is then canonically isomorphic to that of F on Q, hence the Topological Compatibility Axiom A–14 is satisfied. Finally, Topological Compatibility Axiom A–13 follows from invariance of Γ_0 and the definition of κ given above.

Gluing and extending. Gluing is done in the obvious way using the map ρ. Given F, for the missing annulus maps one may take $\mathcal{F}^{\mathcal{C}} = F|\mathcal{C}$. The missing disk maps $\mathcal{F}^{\mathcal{U}}$ were already defined by restricting F to $\mathcal{U}$.

This completes the verification of the axioms, and the proof of Proposition 5.4. $\square$

Thus, an embellished branched covering $F : (S^2, Q) \to (S^2, Q)$ in standard form with respect to an invariant multicurve $\Gamma_0 \subset S^2 - Q$ is indeed an amalgam.

5.4 Proof of Decomposition Theorem

Most of the statements follow directly from the construction in the previous subsection. Statement 2(a) warrants a brief discussion. First, given two different standard forms F, F' for the same pair (G, Γ), it is easy to show that the resulting mapping trees $\mathbf{T}, \mathbf{T}'$ are canonically isomorphic. The only arbitrary choices are in the dynamics function, which is only defined up to postcomposition with vertex-fixing homeomorphisms anyway. Similarly, given a combinatorial equivalence of pairs $(G, \Gamma) \to (G', \Gamma')$ and standard forms F, F', there is an induced isomorphism between the corresponding mapping trees $\mathbf{T}, \mathbf{T}'$. $\square$

6

Uniqueness of decompositions

6.1 Statement of Uniqueness of Decompositions Theorem

Recall that two pairs, each consisting of an embellished map of spheres and an invariant multicurve, are *combinatorially equivalent* if there is a combinatorial equivalence between them sending one multicurve to the other, up to isotopy.

Theorem 6.1 (Uniqueness of Decomposition). *Let $G : (S^2, Q) \to (S^2, Q)$ be an embellished branched cover and Γ an invariant multicurve. Let $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ be a decomposition of G along Γ. Then the combinatorial class of $\mathcal{F}$ is determined by the equivalence class of pairs (G, Γ).*

Given $G : (S^2, Q) \to (S^2, Q)$ and an invariant multicurve Γ_0, the construction of a standard form F depends on a choice of annuli $\mathcal{A}_0$ and on the choice of homeomorphism h for which $F = G \circ h$ and $h \simeq \text{id}$ relative to Q. Furthermore, the construction of the map of spheres $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ depends on these choices, and on others. Hence the conclusion of the theorem is indeed non-obvious.

The proof will be similar in spirit to the proof of the Uniqueness of Combinations Theorem [6.1]: by applying conjugacies we will reduce to the case when the structure data and boundary values are fixed. The conclusion will follow by an appeal to the Tameness Restriction # 2, §1.9.

6.2 Proof of Uniqueness of Decomposition Theorem

Let $G : (S^2, Q) \to (S^2, Q)$ and $G' : (S^2, Q') \to (S^2, Q')$ be embellished maps. Let Γ, Γ' be multicurves invariant under G, G', respectively. Suppose the pair (G, Γ) is combinatorially equivalent to (G', Γ'). Let $F = G \circ h$ and $F' = G' \circ h'$ be standard forms of G, G' with respect to Γ, Γ' having sets of annuli $\mathcal{A}_0, \mathcal{A}_0'$, respectively. Let $\mathcal{F}^{\mathcal{S}} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ be the embellished map of spheres

K.M. Pilgrim: LNM 1827, pp. 79–81, 2003.

produced by decomposing F along $\mathcal{A}_0$, using any set of choices, and similarly define $\mathcal{F}^{\mathcal{S}'} : (\mathcal{S}', \mathcal{Y}') \to (\mathcal{S}', \mathcal{Y}')$.

Suppose $h_0, h_1 : (S^2, Q) \to (S^2, Q')$ yield a combinatorial equivalence between (G, Γ) and (G', Γ'). Set $H_0 = h_0$ and $H_1 = (h')^{-1} \circ h_1 \circ h$. Then $H_0, H_1 : (S^2, Q) \to (S^2, Q')$ form a combinatorial equivalence between F and F'.

Let $\mathcal{S}_0 = S^2 - \mathcal{A}_0$ and $\mathcal{S}_0' = S^2 - \mathcal{A}_0'$. Let $\mathcal{S}_1, \mathcal{S}_1'$ be the subsets of $\mathcal{S}_0, \mathcal{S}_0'$ as defined in §5.3, regarded for now as subsets of S^2. Then $H_0(\mathcal{S}_1)$ and $\mathcal{S}_1'$ are ambiently homeomorphic via a map isotopic to the identity rel Q', hence by Tameness Restriction #3 there exists a tame homeomorphism $H : (S^2, Q') \to (S^2, Q')$ isotopic to the identity rel Q' sending $H_0(\mathcal{S}_1)$ onto $\mathcal{S}_1'$.

Set $\psi_0 = H \circ H_0 : (S^2, \mathcal{S}_1, Q) \to (S^2, \mathcal{S}_1', Q')$ and lift ψ_0 under F, F' to obtain a map $\psi_1 : (S^2, Q) \to (S^2, Q')$ such that the pair (ψ_0, ψ_1) is a combinatorial equivalence between F and F'; this is possible since $H \circ H_0$ is isotopic to H_0 rel Q. Then

$$F' = \psi_0 \circ F \circ \psi_1^{-1}$$

and so

$$F'' = \psi_0^{-1} \circ F' \circ \psi_0 = F \circ \psi_1^{-1} \circ \psi_0 = G \circ (h \circ \psi_1^{-1} \circ \psi_0)$$

is a map topologically conjugate to F'. Furthermore, F'' is a standard form of G with respect to Γ. To see this, note first that by construction F'' is obtained by precomposing G with the map $h \circ \psi_1^{-1} \circ \psi_0$ which is isotopic to the identity rel Q. Next, ψ_0 carries $\mathcal{S}_1$ to $\mathcal{S}_1'$ by construction, hence in particular $\psi_0 : \mathcal{A}_0 \to \mathcal{A}_0'$. Since F' has associated annuli $\mathcal{A}_0'$, the conjugate F'' therefore has associated annuli $\psi_0^{-1}(\mathcal{A}_0') = \mathcal{A}_0$.

Recall that $\mathcal{F}^{\mathcal{S}'}$ denotes the map of spheres produced from F' as in the Decomposition Theorem [5.1], using any set of choices. The Decomposition Theorem is natural in the following sense: since we identify $\mathcal{S}'$ with copies of the domain S^2 of F', as opposed to an abstract manifold, a choice of map of spheres $\mathcal{F}^{\mathcal{S}'}$ produced using F' and a homeomorphism ψ_0 conjugating F' to F'' induces a topological conjugacy

$$\psi_0^{\mathcal{S}} : \mathcal{S}' \to \mathcal{S}''$$

between the map of spheres $\mathcal{F}^{\mathcal{S}'}$ and another map of spheres denoted $\mathcal{F}^{\mathcal{S}''}$ which is produced from F'' using a set of choices which is uniquely determined by F' and ψ_0.

Putting the previous two paragraphs together, we have established that any choice of map of spheres $\mathcal{F}^{\mathcal{S}'}$ produced using F' is tamely topologically conjugate, hence combinatorially equivalent, to a map of spheres $\mathcal{F}^{\mathcal{S}''}$ produced using a map F'' which is topologically conjugate to F', and for which the surface $\mathcal{S}_1''$ coincides with the surface $\mathcal{S}_1$ of F.

It therefore suffices to prove

Proposition 6.2. *Suppose F, F' are standard forms of a single map G with respect to a multicurve Γ, with the property that the surfaces $\mathcal{S}_1, \mathcal{S}'_1$ coincide. Let $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}'}$ be the associated sphere maps as in the Decomposition Theorem [5.1], produced using any set of choices. Then $\mathcal{F}^{\mathcal{S}}$ and $\mathcal{F}^{\mathcal{S}'}$ are combinatorially equivalent.*

Let F, F' and $\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{S}'}$ be as in the above Proposition 6.2. Note that now $\mathcal{S} = \mathcal{S}'$ and $\mathcal{S}_1 = \mathcal{S}'_1$.

Proposition 6.3. *There is a one-parameter family of maps*

$$\mathcal{F}^t : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y}^t), \quad t \in [0,1]$$

such that

1. $\mathcal{F}^t | \mathcal{Q} = \mathrm{id}_{\mathcal{Q}}$*, all* t
2. $\mathcal{F}^0 = \mathcal{F}^{\mathcal{S}}$
3. $\mathcal{F}^1 | \mathcal{S}_1 = \mathcal{F}' | \mathcal{S}_1$.
4. *the subsets* $\mathcal{Y}^t$ *vary (tamely) isotopically in* S^2

Proof: Since $F = G \circ h$ and $F' = G \circ h'$ where h, h' are (tamely) isotopic to the identity rel Q, we have $F = F' \circ H$ where $H = (h')^{-1} \circ h$ is isotopic to the identity rel Q. Let H_t be such an isotopy, with $H_0 = \mathrm{id}_{S^2}$ and $H_1 = H$. Note that while $H_t(\mathcal{S}_1)$ may vary, $H_0(\mathcal{S}_1) = H_1(\mathcal{S}_1) = \mathcal{S}_1$. As before the maps H_t induce homeomorphisms $\mathcal{H}_t^{\mathcal{S}} : \mathcal{S} \to \mathcal{S}$. We set

$$\mathcal{F}^t = \mathcal{F}^{\mathcal{S}} \circ \mathcal{H}_t^{-1} : \mathcal{S} \to \mathcal{S}.$$

By construction, (1)–(4) hold. $\square$

Proof of Prop. 6.2 and conclusion of proof. Since $\mathcal{F}^{\mathcal{S}} = \mathcal{F}^0$ and $\mathcal{F}^1$ are joined by a path in which the marked sets $\mathcal{Y}_t$ vary isotopically, $\mathcal{F}^{\mathcal{S}}$ and $\mathcal{F}^1$ are combinatorially equivalent by Tameness Restriction #2. To prove Proposition 6.2 it therefore remains only to prove that $\mathcal{F}^{\mathcal{S}'}$ and $\mathcal{F}^1$ are combinatorially equivalent.

The two maps agree on $\mathcal{S}_1$, whose complement is a union of disks with preferred, possibly different centers as the only possible ramification points; the possibility of varying centers occurs when $\mathcal{Z}^{exc}$ is nonempty. It follows easily that $\mathcal{F}^{\mathcal{S}'}$ and $\mathcal{F}^1$ are combinatorially equivalent, by precomposing with a homeomorphism isotopic to the identity rel $\partial \mathcal{S}_1$ so as to arrange that the centers are the same, and then applying the Alexander Trick; cf. the proof of Proposition 4.7.

Thus, $\mathcal{F}$ and $\mathcal{F}'$ are combinatorially equivalent, hence Proposition 6.2 and the Uniqueness of Decompositions Theorem are proven. $\square$

7

Counting classes of annulus maps

7.1 Statement of Number of Classes of Annulus Maps Theorem

Fix a mapping tree $\mathbf{T}$, and let $\mathcal{A}_0, \mathcal{A}_1$ be as in the definition of a map of annuli over $\mathbf{T}$ (§2.3). Here, we bound the number of combinatorial classes of annulus maps.

To set up the statement, recall that A_0, A_1 are annular edges and edges of T, respectively, which lie below components of $\mathcal{A}_0$ and $\mathcal{A}_1$ under the projection $\pi^{\mathcal{A}}$. By abuse of notation, we set

$$A_0 = \{1, 2, ..., i, ..., n_0\}$$

$$A_1 = \{1, 2, ..., j, ..., n_1\}.$$

The dynamics function $f : T \to T$ determines a map (also denoted f)

$$f : A_1 \to A_0$$

given by

$$f(j) = i \quad \text{if } f(a_{1,j}) = a_{0,i}.$$

In addition we have a surjection

$$\epsilon : A_1 \to A_0$$

given by

$$\epsilon(j) = i \quad \text{if} \quad a_{1,j} \subset a_{0,i}$$

i.e. if the annular edges satisfy $a_{1,j} \subset a_{0,i}$ (recall the definition given in §2.1). Recall that in addition, we have a signed degree function

K.M. Pilgrim: LNM 1827, pp. 83–88, 2003.

$$\delta : A_1 \to \mathbb{Z} - \{0\}.$$

Define matrices $L, M, (\epsilon)$ by

$$L_{ii} = \begin{cases} \operatorname{lcm}\{|\delta(j)| \mid f(j) = i\} & \text{if } f^{-1}(i) \neq \emptyset \\ 1 & \text{otherwise} \end{cases}$$

$$M_{ji} = \begin{cases} \frac{L_{ii}}{|\delta(j)|} & \text{if } f(j) = i \\ 0 & \text{otherwise} \end{cases}$$

and

$$(\epsilon)_{ij} = \begin{cases} 1 & \text{if } \epsilon(j) = i \\ 0 & \text{otherwise} \end{cases}$$

Theorem 7.1 (Number of classes of annulus maps). *There is a bijection between the set of combinatorial classes of annulus maps over T and the elements of the quotient group*

$$\mathbb{Z}^{A_0}/(L - (\epsilon)M)\mathbb{Z}^{A_0}.$$

In particular, this set is finite if and only if the matrix

$$N_{\mathbf{T}} = (\epsilon)ML^{-1}$$

does not have one as an eigenvalue.

The matrix $N_{\mathbf{T}}$ is the matrix of the Thurston linear map (F_Γ) of any amalgam produced using $\mathbf{T}$; see Proposition 3.6. The equality $N_{\mathbf{T}} = (\epsilon)ML^{-1}$ is part of the conclusion of the theorem.

7.2 Proof of Number of Classes of Annulus Maps Theorem

7.2.1 Homeomorphism of annuli. Index.

Homeomorphism of annuli. Index. We denote by $\mathcal{G}$ the space of all homeomorphisms $g : A \to A$ such that $g|\partial A = \mathrm{id}$, endowed with the uniform topology, and by $\mathcal{G}^0$ the subspace of maps isotopic to the identity through maps which are the identity on the boundary. There is a split exact sequence

$$1 \to \mathcal{G}^0 \to \mathcal{G} \underset{\sigma}{\overset{\pi}{\rightleftarrows}} \mathbb{Z} \to 0$$

where the map π may be described as follows. Identify the ends of the annulus to form a torus. The induced map on the fundamental group of this torus is a matrix of the form

$$(g_*) = \begin{pmatrix} 1 & \pi(g) \\ 0 & 1 \end{pmatrix}$$

where the integer $\pi(g)$ represents the number of times the image of the longitude wraps around the torus in the meridional direction. A section σ is given by

$$\sigma(n) = (s, t) \mapsto (s + nt, t)$$

i.e. by the nth power of a standard Dehn twist. Thus, indices characterize isotopy classes (relative to the boundary) of self-homeomorphisms of annuli which are the identity on the boundary. The index is additive in the sense that if an annulus A_0 is the union of two subannuli A_1, A_1' along a common boundary component, and if g_1, g_1' are two homeomorphisms from A_1, A_1' to A_1, A_1' which are the identity on the boundary, then the index of the homeomorphism of A_0 to itself obtained by gluing together g_1, g_1' is the sum of the indices of g_1, g_1'.

7.2.2 Characterization of combinatorial equivalence by group action.

The set of all annulus maps lying over T forms a topological space $\mathcal{E}$, using the uniform topology. There is a preferred basepoint $\mathcal{F}_0^{\mathcal{A}}$ sending the jth annulus of $\mathcal{A}_1$ to its image via

$$(s, t) \mapsto (\delta(j)s, \operatorname{sgn}(\delta(j))t).$$

Let $\mathcal{G}^{A_0}$, $\mathcal{G}^{A_1}$ denote the Cartesian products of $|A_0|, |A_1|$ copies of $\mathcal{G}$, respectively, and let $\pi^{A_0} : \mathcal{G}^{A_0} \to \mathbb{Z}^{A_0}$ and $\pi^{A_1} : \mathcal{G}^{A_1} \to \mathbb{Z}^{A_1}$ be the homomorphisms obtained by applying π in each factor. Let

$$\begin{aligned} \mathcal{K} &= \ker(\pi^{A_0} \times \pi^{A_1} : \mathcal{G}^{A_0} \times \mathcal{G}^{A_1} \to \mathbb{Z}^{A_0} \oplus \mathbb{Z}^{A_1}) \\ &= (\mathcal{G}^0)^{A_0} \times (\mathcal{G}^0)^{A_1} \end{aligned}$$

The group $\mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$ acts transitively and faithfully on $\mathcal{E}$ by

$$(g_0, g_1).\mathcal{F}^{\mathcal{A}} = g_0 \circ \mathcal{F}^{\mathcal{A}} \circ g_1^{-1}.$$

Let $\mathcal{H}$ be the stabilizer of the basepoint $\mathcal{F}_0^{\mathcal{A}}$ under this action. Since the action is transitive, the action of $\mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$ on $\mathcal{E}$ is isomorphic to the action of $\mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$ on the set of cosets $(\mathcal{G}^{A_0} \times \mathcal{G}^{A_1})/\mathcal{H}$ by left multiplication. Furthermore, additivity of indices and the definition of combinatorial equivalence (Definition 4.4) show that $\mathcal{F}^{\mathcal{A}}, \mathcal{F}^{\mathcal{A}'} \in \mathcal{E}$ are combinatorially equivalent if and only if $(g_0, g_1).\mathcal{F}^{\mathcal{A}} = \mathcal{F}^{\mathcal{A}'}$ for some $(g_0, g_1) \in \mathcal{D}$, where

$$\mathcal{D} = \{(g_0, g_1) | (\epsilon)\pi^{A_1}(g_1) = \pi^{A_0}(g_0)\}$$

and where (ϵ) is the matrix given previously. Clearly $\mathcal{D} \supset \mathcal{K}$. Note that since $\mathcal{D}$ is the preimage of a subgroup of an abelian group under a homomorphism, $\mathcal{D}$ is a normal subgroup of $\mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$.

7.2.3 Reduction to abelian groups

Orbit equivalence.

Lemma 7.2. *Suppose G acts transitively on a set E with stabilizer H. Let $\pi : G \to \pi(G)$ be an epimorphism with kernel K, and suppose D is a subgroup of G containing K.*

Then the action of the subgroup D on $E = G/H$ descends to an action of $\pi(D)$ on $\pi(D)/\pi(H)$, and the natural map $E = G/H \to \pi(G)/\pi(H)$ is a bijection between the set of orbits of D on G/H and that of $\pi(D)$ on $\pi(G)/\pi(H)$.

Proof: By transitivity, the action of G on E is isomorphic to the action of G on G/H by left multiplication. Since π is an epimorphism, the map $G/H \to \pi(G)/\pi(H)$ is a surjection. Together with π, this defines an epimorphism between the actions of G on G/H and $\pi(G)$ on $\pi(G)/\pi(H)$. By restricting to D we obtain similarly an epimorphism between the actions of D on G/H and $\pi(D)$ on $\pi(G)/\pi(H)$. Thus the map $G/H \to \pi(G)/\pi(H)$ sends D-orbits surjectively to $\pi(D)$-orbits and is a surjection on the set of orbits.

We claim in fact that this map between sets of orbits is injective. For suppose $\pi(D)\pi(g)\pi(H) = \pi(D)\pi(g')\pi(H)$. Then $DgH = KDg'H$. But by hypothesis $K \subset D$, so $DgH = Dg'H$. □

Applying Lemma 7.2 with $G = \mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$, $H = \mathcal{H}$, $K = \mathcal{K}$, $\pi = \pi^{A_0} \times \pi^{A_1}$, and $E = \mathcal{E}$, we see that the set of combinatorial equivalence classes of elements of $\mathcal{E}$ is in bijective correspondence with the orbits of $\pi(\mathcal{D}) < \mathbb{Z}^{A_0} \oplus \mathbb{Z}^{A_1}$ acting by (left) addition on the set of cosets $(\mathbb{Z}^{A_0} \oplus \mathbb{Z}^{A_1})/\pi(\mathcal{H})$.

7.2.4 Computations and conclusion of proof

Lemma 7.3. *Let $\delta \in \mathbb{Z} - \{0\}$ and suppose $F^A : A \to A$ is given by $F^A(s,t) = (\delta s, \mathrm{sgn}(\delta)t)$. If $g_0, g_1 \in \mathcal{G}$ and $g_0 \circ F^A = F^A \circ g_1$, then $\mathrm{sgn}(\delta)\pi(g_0) = \delta\pi(g_1)$.*

Proof: The maps $g_0 \circ F^A$ and $F^A \circ g_1$ induce maps of the torus obtained by identifying the boundary components. With respect to the standard basis of the torus, functoriality of induced maps on fundamental groups yields

$$(g_0 \circ F^A)_* = \begin{pmatrix} \delta & \pi(g_0) \\ 0 & \mathrm{sgn}(\delta) \end{pmatrix}$$

and

$$(F^A \circ g_1)_* = \begin{pmatrix} \delta & \delta\pi(g_1) \\ 0 & \mathrm{sgn}(\delta) \end{pmatrix}.$$

□

Lemma 7.4. *We have*

$$(\pi^{A_0} \times \pi^{A_1})(\mathcal{H}) = \text{col}\begin{pmatrix} L \\ M \end{pmatrix}$$

where the right-hand side is the subgroup of $\mathbb{Z}^{A_0} \oplus \mathbb{Z}^{A_1}$ *generated by the columns of the indicated matrix.*

Proof: Suppose $(g_0, g_1) \in \mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$. Set $y = \pi^{A_0}(g_0) = (y_1, ..., y_i, ..., y_{n_0})$ and $x = \pi^{A_1}(g_1) = (x_1, ...x_j, ..., x_{n_1})$. Then Lemma 7.3 implies that $(y, x) \in (\pi^{A_0} \times \pi^{A_1})(\mathcal{H})$ if and only if

$$\text{sgn}(\delta(j)) \cdot y_{f(j)} = \delta(j) \cdot x_j \text{ for all } j = 1, ..., n_1$$

which is equivalent to

$$y_{f(j)} = |\delta(j)| \cdot x_j \text{ for all } j = 1, ..., n_1.$$

The solution set of these equations (over the integers) has a basis equal to the columns of the block matrix

$$\begin{array}{cc} \begin{pmatrix} L \\ M \end{pmatrix} & \begin{array}{c} A_0 \\ A_1 \end{array} \end{array}$$

Here, the rows correspond to the indicated subsets. The matrix L has $|A_0|$ rows and columns, is diagonal, and has all diagonal entries nonzero. M has $|A_1|$ rows and $|A_0|$ columns,

$$L_{ii} = \begin{cases} \text{lcm}\{|\delta(j)| \mid f(j) = i\} & \text{if } f^{-1}(i) \neq \emptyset \\ 1 & \text{otherwise} \end{cases}$$

and

$$M_{ji} = \begin{cases} \frac{L_{ii}}{|\delta(j)|} & \text{if } i = f(j) \\ 0 & \text{otherwise} \end{cases}$$

□

Lemma 7.5. *We have*

$$(\pi^{A_0} \times \pi^{A_1})(\mathcal{D}) = \text{col}\begin{pmatrix} (\epsilon) \\ I \end{pmatrix}.$$

Proof: This follows directly from the definition of $\mathcal{D}$. □

Conclusion of proof: Since the ambient group is abelian, the orbits of the action of $(\pi^{A_0} \times \pi^{A_1})(\mathcal{D})$ on the quotient $(\mathbb{Z}^{A_0} \oplus \mathbb{Z}^{A_1})/(\pi^{A_0} \times \pi^{A_1})(\mathcal{H})$ is thus the quotient of $\mathbb{Z}^{A_0} \oplus Z^{A_1}$ by the group generated by the columns of the two matrices

$$\begin{pmatrix} L \\ M \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} (\epsilon) \\ I \end{pmatrix}.$$

i.e. by the columns of the matrix

$$\begin{pmatrix} L & (\epsilon) \\ M & I \end{pmatrix}.$$

Elementary row and column operations performed over the integers will not change the isomorphism type of the quotient. Thus, the quotient is isomorphic to

$$\mathbb{Z}^{A_0}/(L - (\epsilon)M)\mathbb{Z}^{A_0}$$

and this proves the first part of Theorem 7.1.

To prove the last statement, notice that the quotient is infinite if and only if $L - (\epsilon)M$ is singular as a real matrix. Since L is invertible over the reals, this occurs if and only if $I - (\epsilon)ML^{-1}$ is singular, i.e. if and only if $(\epsilon)ML^{-1}$ has one as an eigenvalue. Simply applying the definitions of $(\epsilon), M, L$ and multiplying shows that

$$((\epsilon)ML^{-1})_{ij} \quad = \quad \sum_{\substack{\epsilon(k)=i \\ f(k)=j}} \frac{1}{|\delta(k)|} \quad = \quad (N_{\mathbf{T}})_{ij}.$$

8

Applications to mapping class groups

Our methods also shed light on certain combinatorial automorphisms of branched coverings. In this section, we use throughout the notation of §7.

8.1 The Twist Theorem

Definition 8.1 (Combinatorial automorphisms. Mapping class group.). *Let $G : (S^2, Q) \to (S^2, Q)$ be an embellished branched covering and $\Gamma_0 \subset S^2 - Q$ a G-invariant multicurve. The* group of combinatorial automorphisms *of G, denoted* $\mathrm{Mod}(G)$*, is the set of pairs (h_0, h_1) of homeomorphisms for which (h_0, h_1) yield a combinatorial equivalence from G to itself. The* mapping class group *of G is the quotient of the combinatorial automorphism group by the equivalence relation $(h_0, h_1) \sim (h'_0, h'_1)$ if h_0 and h_1 are respectively isotopic to h'_0 and h'_1 through tame homeomorphisms fixing Q.*

Note that when Q is finite, the map sending (h_0, h_1) to h_0 descends to an embedding of $\mathrm{Mod}(G)$ into the classical mapping class group $\mathrm{Mod}(S^2, Q)$ consisting of orientation-preserving homeomorphisms of S^2 sending Q to itself, modulo isotopy.

Remark: According to M. Rees, the mapping class group of G is isomorphic to the fundamental group of B, where B is a space of maps which are combinatorially equivalent to G (in a sense similar to that given here). In particular, G is combinatorially equivalent to a rational map if and only if B is contractible. See [Ree5], §6.3.

Theorem 8.2 (Twist Theorem). *Let $G : (S^2, Q) \to (S^2, Q)$ be an embellished branched covering and Γ an invariant multicurve. Let $\mathbf{T}$ be the mapping tree determined by decomposing G along Γ and let A_0, A_1, L, M, ϵ be as in Section §7.1.*

K.M. Pilgrim: LNM 1827, pp. 89–94, 2003.

Then the mapping class group of G contains a subgroup $\mathrm{Tw}(G, \Gamma)$ *isomorphic to the intersection in* $\mathbb{Z}^{A_0} \times \mathbb{Z}^{A_1}$ *of the subgroups generated by the columns of the matrices*

$$\begin{pmatrix} L \\ M \end{pmatrix} \quad \textit{and} \quad \begin{pmatrix} (\epsilon) \\ I \end{pmatrix}.$$

The group $\mathrm{Tw}(G, \Gamma)$ *is nontrivial if and only if the Thurston linear map* $G_\Gamma = N_{\mathbf{T}} = (\epsilon) M L^{-1}$ *has one as an eigenvalue. In this case, it is an infinite free abelian group generated by powers of Dehn twists supported in the annuli* $\mathcal{A}_0$ *and* $\mathcal{A}_1$*.*

Note that the conclusion makes sense, since by the Uniqueness of Decompositions Theorem [6.1] the mapping tree $\mathbf{T}$ depends only on the equivalence class of pair (G, Γ) and not on any choices. Also, the above subgroup is actually isomorphic to an analogous mapping class group for annulus maps; see §8.2.1.

8.2 Proof of Twist Theorem

8.2.1 Combinatorial automorphisms of annulus maps

Definition 8.3 (Combinatorial automorphisms of annulus maps). *Let* $\mathcal{F}^{\mathcal{A}} \in \mathcal{E}$ *be a map of annuli over* $\mathbf{T}$*. A* combinatorial automorphism *of* $\mathcal{F}^{\mathcal{A}}$ *is a pair* $(g_0, g_1) \in \mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$ *such that* $(g_0, g_1).\mathcal{F}^{\mathcal{A}} = \mathcal{F}^{\mathcal{A}}$ *and* $(g_0, g_1) \in \mathcal{D}$*.*

The set of combinatorial automorphisms of $\mathcal{F}^{\mathcal{A}}$ therefore form a group under composition. Since $\mathcal{G}^{A_0} \times \mathcal{G}^{A_1}$ acts transitively on $\mathcal{E}$ and $\mathcal{D}$ is normal, the combinatorial automorphism groups of any two elements $\mathcal{F}^{\mathcal{A}}$ are conjugate. We define the *mapping class group* of $\mathcal{F}^{\mathcal{A}}$ to be the image under $\pi^{A_0} \times \pi^{A_1}$ of the group of combinatorial automorphisms of $\mathcal{F}^{\mathcal{A}}$. Thus, the isomorphism class of mapping class group depends only on $\mathcal{E}$, so that we may speak of the *mapping class group* $\mathrm{MCG}(\mathcal{E})$ of $\mathcal{E}$.

Lemma 8.4. *Up to isomorphism,*

$$\begin{aligned} MCG(\mathcal{E}) &= (\pi^{A_0} \times \pi^{A_1})(\mathcal{H} \cap \mathcal{D}) \\ &= (\pi^{A_0} \times \pi^{A_1})(\mathcal{H}) \cap (\pi^{A_0} \times \pi^{A_1})(\mathcal{D}) \\ &= \mathrm{col} \begin{pmatrix} L \\ M \end{pmatrix} \cap \mathrm{col} \begin{pmatrix} (\epsilon) \\ I \end{pmatrix}. \end{aligned}$$

Proof: The first equality follows directly from the definition of the mapping class group, using the basepoint $\mathcal{F}_0^{\mathcal{A}}$ whose stabilizer is $\mathcal{H}$. The second inequality is a consequence of the general fact that if $\phi : G \to G'$ is a homomorphism with kernel K, and if H and $D \supset K$ are subgroups of G, then $\phi(H \cap D) = \phi(H) \cap \phi(D)$. The third follows from the second and the computations of the images of $\mathcal{H}$ and $\mathcal{D}$. $\square$

Lemma 8.5. *The following conditions are equivalent:*

1. *$\mathcal{E}$ has finitely many equivalence classes*
2. *$N_{\mathbf{T}}$ does not have one as an eigenvalue*
3. *the mapping class group of $\mathcal{E}$ is trivial.*

Proof: Equivalence of the first two conditions follows from Theorem 7.1. Suppose the second condition holds. The proof of Theorem 7.1 shows that this holds if and only if the matrix

$$\begin{pmatrix} L & (\epsilon) \\ M & I \end{pmatrix}$$

has full rank equal to $|A_0| + |A_1|$. The matrices

$$\begin{pmatrix} L \\ M \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} (\epsilon) \\ I \end{pmatrix}$$

have rank $|A_0|$ and $|A_1|$ respectively. Thus, (2) is equivalent to the condition that the subgroups generated by the columns of the two matrices above have trivial intersection, i.e. $(\pi^{A_0} \times \pi^{A_1})(\mathcal{H}) \cap (\pi^{A_0} \times \pi^{A_1})(\mathcal{D})$ is trivial. By the previous Lemma, this is equivalent to triviality of the mapping class group. $\square$

8.2.2 Conclusion of proof of Twist Theorem

A combinatorial equivalence (h_0', h_1') between maps G, G' induces an isomorphism between their combinatorial automorphism groups by sending $(h_0, h_1) \mapsto (h_0' \circ h_0 \circ (h_0')^{-1}, h_1' \circ h_1 \circ (h_1')^{-1})$ which descends to an isomorphism between mapping class groups. Therefore, it is enough to show that some representative of G has the indicated properties.

By the Decomposition Theorem [5.1], G is combinatorially equivalent to an amalgam $F = F(\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{C}}, \mathcal{F}^{\mathcal{A}})$. By the Unwinding Trick (Proposition 4.10) and Corollary 4.12, there exist $\mathcal{F}^{\mathcal{C}''}$ such that $F'' = F(\mathcal{F}^{\mathcal{S}}, \mathcal{F}^{\mathcal{C}''}, \mathcal{F}^{\mathcal{A}})$ is in simple form with respect to F. Given any pair of homeomorphisms $\mathcal{H}_0^{\mathcal{A}} : \mathcal{A}_0 \to \mathcal{A}_0$ and $\mathcal{H}_1^{\mathcal{A}} : \mathcal{A}_1 \to \mathcal{A}_1$ which yield combinatorial automorphisms of $\mathcal{F}^{\mathcal{A}}$, the proof of the Uniqueness of Combinations Theorem [4.5] shows that the extension $\mathcal{H}_0$ of $\mathcal{H}_0^{\mathcal{A}}$ by the identity over $\mathcal{S}_0$ lifts under F and F'' to a homeomorphism $\mathcal{H}_1$ such that $(\mathcal{H}_0, \mathcal{H}_1)$ is a pair giving a combinatorial equivalence from F to F''. Thus we have a function from the group of combinatorial automorphisms of $\mathcal{F}^{\mathcal{A}}$ to the set of combinatorial equivalences from F to F''. Choosing arbitrarily a basepoint equivalence, we may identify the set of equivalences from F to F'' with the combinatorial automorphism group of F so that this function is a homomorphism. It descends to a homomorphism from the mapping class group of annulus maps $\mathcal{E}$ to the mapping class group of F. By construction, it is generated by powers of Dehn twists along the annuli in $\mathcal{A}_0$. Since the annuli in $\mathcal{A}_0$ are nonperipheral, this homomorphism is injective. Thus, the mapping class group of $\mathcal{E}$ embeds into that of F, and the theorem follows from the Lemmas in the previous subsection. $\square$

8.3 When Thurston obstructions intersect

The Twist Theorem [8.2] sheds light on what happens when Thurston obstructions intersect. *For simplicity, we restrict our attention to a map of a single sphere* $G : (S^2, Q) \to (S^2, Q)$ *for which* Q *is finite.*

8.3.1 Statement of Intersecting Obstructions Theorem

We shall need some basic facts from the theory of nonnegative square matrices; see ([Gan], Volume 2, ch. XIII). Such a matrix A always has a largest nonnegative (or *leading*) eigenvalue $\lambda(A)$ equal to its spectral radius.

Irreducible matrices. A nonnegative n-by-n matrix A is called *irreducible* if no permutation of the indices places it in block lower-triangular form. In this case, given i, j there exists q with $0 \le q \le n$ such that $(A^q)_{i,j} \neq 0$. By Frobenius' Theorem, A has a positive eigenvalue $\lambda(A)$ equal to its spectral radius, and up to scale, there is exactly one positive eigenvector corresponding to $\lambda(A)$.

Reducible matrices. By applying a permutation of the indices, A may be assumed to be in the form

$$\begin{pmatrix} A_{11} & 0 & ... & 0 \\ A_{21} & A_{22} & ... & 0 \\ ... & ... & ... & 0 \\ A_{s1} & A_{s2} & ... & A_{ss} \end{pmatrix}$$

where the blocks A_{jj} are irreducible. Moreover,

$$\lambda(A) = \sup_j \lambda(A_{jj}).$$

We shall need to consider linear transformations for multicurves Γ which are not necessarily invariant. For an embellished branched map $G : (S^2, Q) \to (S^2, Q)$ Define $G_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ by the same formula. We say Γ is *irreducible* if the matrix (G_Γ) is irreducible. Let Γ be a multicurve and $A = (G_\Gamma)$. Let Γ_j denote the set of indices (curves in Γ) corresponding to the jth block in the decomposition given in the previous subsection, so that $A_{jj} = (G_{\Gamma_j})$. We call the Γ_j's the *irreducible components* of Γ. As sets, they are well-defined up to permutations of the indices j.

Irreducible obstructions. In the same setting, we say that an irreducible multicurve Γ for G is an *irreducible Thurston obstruction* if its leading eigenvalue is at least one. Therefore, by the above results, any Thurston obstruction contains an irreducible obstruction with the same eigenvalue.

Intersection numbers. Given simple closed curves $\alpha, \beta \in S^2 - Q$ their *intersection number* is

$$\alpha \cdot \beta = \inf\{\#\alpha' \cap \beta' | \alpha' \simeq \alpha, \beta' \simeq \beta\}$$

where $\simeq$ denotes free homotopy in $S^2 - Q$. The intersection number of multicurves $\Gamma_1 \cdot \Gamma_2$ is defined by extending (bi)linearly. Note that in particular $\Gamma_1 \cdot \Gamma_2 = 0$ if and only if up to homotopy either $\Gamma_1 = \Gamma_2$ or $\Gamma_1 \cap \Gamma_2 = \emptyset$.

Theorem 8.6 (Intersecting Irreducible Obstructions). *Suppose $G : (S^2, Q) \to (S^2, Q)$ is an embellished branched covering with Q finite, and Γ_1, Γ_2 are two irreducible obstructions. If $\Gamma_1 \cdot \Gamma_2 \neq 0$ then the leading eigenvalues of Γ_1 and Γ_2 are both equal to one.*

For explicit examples, see [DH3], Example 3, pp. 295-6.

Proof: While it seems likely that a direct combinatorial proof can be given, we prefer to make use of the connection between G_Γ and the growth of moduli of annuli when one tries to realize G analytically. We shall freely exploit the facts from conformal geometry in [DH3].

We consider the induced map σ_G on Teichmüller space considered by Douady and Hubbard [DH3]. Fix a basepoint τ_0 with underlying surface $X_0 = \mathbb{P}^1 - Q_0$ and let $\tau_n = \sigma_G^{\circ n}(\tau_0)$ have underlying surface $X_n = \mathbb{P}^1 - Q_n$.

The hypothesis and the Grötzsch Inequality imply that for each γ in Γ_1, Γ_2, there is a lower bound, independent of n, on the modulus of the largest embedded annulus in X_n with core curve homotopic to γ. This implies that the length of the unique geodesic on X_n homotopic to γ is bounded above independent of n.

If the leading eigenvalue of, say, Γ_1 is strictly greater than one, then the irreducibility assumption implies that the moduli of annuli about each $\gamma \in \Gamma_1$ on X_n tend to infinity. This implies that the length of each geodesic on X_n homotopic into Γ_1 tends to zero as $n \to \infty$.

However, a geodesic on a Riemann surface which intersects a short geodesic on a Riemann surface must be long, by the well-known Collar Lemma. That is, eventually, there are long geodesics homotopic to elements of Γ_2 on X_n, a contradiction.

□

8.3.2 Maps with intersecting obstructions have large mapping class groups

The Twist Theorem (Thm. 8.2) and the Intersecting Irreducible Obstructions Theorem (Thm. 8.6) have as an immediate consequence the following theorem, which sheds light on what happens when obstructions intersect.

Theorem 8.7 (Intersecting obstructions yield large mapping class groups). *Let $G : (S^2, Q) \to (S^2, Q)$ is an embellished branched covering such that Q is finite. If Γ_1, Γ_2 are two irreducible obstructions with nonzero intersection number, then the mapping class group of G contains two distinct free abelian subgroups, each generated by a collection of homeomorphisms which have the form of powers of Dehn twists about elements of Γ_1, Γ_2, respectively.*

It is tempting to try to deduce, from the existence of a pair of intersecting obstructions, the existence of an (up to isotopy) G-invariant subsurface $S \subset S^2 - Q$ such that $G|S$ is a homeomorphism. This is the case of Douady and Hubbard's Example 3 mentioned above: the union of the curves $\beta, \alpha, \delta_3, \delta_4$ constitute such a surface S.

However, the Lattès example, given in the next section, shows that this need not always be the case.

Question: Under what conditions does the presence of two intersecting obstructions imply the existence of such an invariant subsurface S on which G acts as a homeomorphism?

It seems likely that a sufficient condition is that the obstructions intersect in a suitably complicated fashion.

9

Examples

In this section, we give several examples. To make them concrete, and to connect them with the geometric objects they are supposed to emulate, we choose our examples from the theory of complex dynamical systems, though this is not really essential and indeed many of the constructions could just as easily be defined in the more general setting of branched coverings.

9.1 Background from complex dynamics

To explain our examples, we need to generalize the theory of holomorphic dynamical systems from maps defined on a single copy of the Riemann sphere, identified with $\mathbb{P}^1$, to maps on multiple copies.

Definition 9.1 (Rational map of spheres). *A map of spheres $\mathcal{R} : \mathcal{S} \to \mathcal{S}$ is* rational *if $\mathcal{S}$ is a disjoint union of copies of $\mathbb{P}^1$ and the map $\mathcal{R}$ is holomorphic.*

Periodic cycles and their multipliers, or eigenvalues, are defined just as for the case when $\mathcal{S}$ is a single sphere. In particular, an *attractor* (*repeller*) of $\mathcal{R}$ is a periodic cycle with multiplier λ with $|\lambda| < 1$ (> 1). A family of holomorphic functions defined on a domain $\Omega \subset \mathbb{P}^1$ into $\mathcal{S}$ is *normal* if every sequence contains a subsequence converging locally uniformly, where the spherical metric is used on each component to measure the distance between image values. The *Fatou set* of $\mathcal{R}$ is the subset of points $w \in \mathcal{S}$ for which the iterates of $\mathcal{R}$ are normal near w. The *Julia set* of $\mathcal{R}$ is the complement of the Fatou set. Since $\mathcal{R}$ need not be surjective, the Julia set need not coincide with the closure of the repelling periodic points. The *immediate basin* of an attractor is the union of the connected components of the Fatou set containing the points in the attracting cycle. The *grand orbit* of $w \in \mathcal{S}$ is the set $\cup_{n,m} \mathcal{R}^{-m}\mathcal{R}^n(z)$.

Proposition 9.2 (Rational standard form). *Let $\mathcal{R} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ be an embellished rational map of spheres satisfying Embellished Axiom A–6.*

K.M. Pilgrim: LNM 1827, pp. 95–103, 2003.

Suppose $\mathcal{Z}$ is contained in the grand orbit of a union of attractors. Then there is a canonical embellished map of spheres $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ combinatorially equivalent to $\mathcal{R}$ and which satisfies the Standard Form Axioms A–7.

Proof: The images of $\mathcal{Z}$ under $\mathcal{R}$ stabilize under iteration to a union of attracting cycles. Let $\{z_0, z_1, ..., z_{p-1}\}$ be one such attractor and $\Omega_0, \Omega_1, ..., \Omega_{p-1}$ its immediate basin. Since $\mathcal{Z}$ is finite and $\mathcal{Q}$ is closed, $\mathcal{Q}$ does not accumulate on $\mathcal{Z}$. Hence there are no critical points of $\mathcal{R}$ which are not in $\mathcal{Z}$ and which accumulate on but do not lie in $\mathcal{Z}$. The "Critical points in the basin" Theorem ([Mil4], Thm. 9.3) applied to $\mathcal{R}^{\circ p}$ at the point z_i implies that for each i, there exists a biholomorphic map $\phi : (\Omega_i, z_i) \to (\Delta, 0)$ such that $(\phi_i(z))^{\omega(z_i)} = \phi_{i+1 \bmod p}(\mathcal{R}(z_i))$ for all $z \in \Omega_i$, where $\omega(z_i)$ is the local degree of $\mathcal{R}$ near z_i. In particular, the basins Ω_i are open disks disjoint from $\mathcal{Q}$.

Furthermore, the maps ϕ_i are unique up to postcomposition with multiplication by certain roots of unity. From this follows the existence of canonical radial coordinates r on each basin Ω_i such that $\mathcal{R}$ preserves the set of radial lines. Let $\mathcal{B}_{z_i}$ be the closed disk in Ω_i given by $r \leq 1/2$ and let $\mathcal{B}'_{z_i} = (\mathcal{R}|\Omega_i)^{-1}(\mathcal{B}_{z_{i+1}})$ (the constant $1/2$ is arbitrary but fixed once and for all). Then $\mathcal{B}'_{z_i} \supset \mathcal{B}_{z_i}$.

Let $\mathcal{H} : \mathcal{S} \to \mathcal{S}$ be a homeomorphism which is the identity off outside of a small neighborhood of $\cup_i \mathcal{B}'_{z_i}$, which fixes the angular coordinates in Ω_i, and which sends $\mathcal{B}_{z_i}$ onto $\mathcal{B}'_{z_i}$. Then $\mathcal{H}$ is isotopic to the identity rel $\mathcal{Y}$. It is straightforword though tedious to write an explicit formula for $\mathcal{H}$ in these radial coordinates which depends only on the set of local degrees $\omega(z_i)$, and so $\mathcal{H}$ is canonical. Then $\mathcal{F} = \mathcal{R} \circ \mathcal{H}$ is combinatorially equivalent to $\mathcal{R}$, and the disks $\mathcal{B}_{z_i}$ are permuted under $\mathcal{F}$. If $z \in \mathcal{Z}$ satisfies $\mathcal{F}^r(z) = z_i$ we set $\mathcal{B}_z$ to be the unique component of $\mathcal{F}^{-r}(\mathcal{B}_{z_i})$ containing z. In this way we obtain a collection $\mathcal{B}$ of disks, and Standard Form Axioms A–7 #s 1,2,3, and 5 are then satisfied.

To verify Axiom A–7 # 4, we first note that each Fatou component in the grand orbit of a point in $\mathcal{Z}$ must be a disk. For otherwise, there is a critical point with an infinite forward orbit accumulating at an attractor in $\mathcal{Z}$, and this is impossible. It follows that the radial coordinates defined above may be pulled back to the Fatou components containing points in the grand orbit of $\mathcal{Z}$ such that each such component Ω has a unique preferred center given by its intersection with the grand orbit of $\mathcal{Z}$. Hence each such component Ω contains at most one point of $\mathcal{Q}$. It follows easily from the construction of $\mathcal{B}$ that Axiom A–7 #4 holds. □

9.2 Matings

Matings. The simplest combinations are *matings*; see e.g. [Tan2]. Loosely, mating is gluing together polynomials of the same degree along their actions on the circle at infinity of the complex plane. Below, we interpret this construction in our language.

Let $\widehat{\mathbb{C}}_\pm$ denote two disjoint copies of the Riemann sphere. Let $R_\pm : \widehat{\mathbb{C}}_\pm \to \widehat{\mathbb{C}}_\pm$ be two monic postcritically finite polynomial maps of the same degree d. Set $\mathcal{S} = \widehat{\mathbb{C}}_- \sqcup \widehat{\mathbb{C}}_+$ and let $\mathcal{R} : \mathcal{S} \to \mathcal{S}$ be defined by applying R_- on $\widehat{\mathbb{C}}_-$ and R_+ on $\widehat{\mathbb{C}}_+$. Let $\mathcal{Q} = \mathcal{Q}_- \sqcup \mathcal{Q}_+$ be the union of the full preimages of the postcritical sets of R_-, R_+ in the finite plane, and let $\mathcal{Z} = \{\infty_-, \infty_+\}$ and $\mathcal{Y} = \mathcal{Q} \sqcup \mathcal{Z}$. Then $\mathcal{R}$ satisfies the hypotheses of Proposition 9.2 (see Axiom A-7).

Let $\mathcal{F}$ be the canonical standard form of $\mathcal{R}$. Note that $\mathcal{B}$ and its complement are both completely invariant under $\mathcal{R}$, i.e. coincide with their grand orbits. Hence $\mathcal{S}_0 = \mathcal{S}_1 = \overline{\mathcal{S} - \mathcal{B}}$. Regard $\mathcal{F}$ as a map of spheres over the mapping tree $\mathbf{T}$ defined as follows. There are two vertices x_-, x_+, each of which has degree d and a single boundary vertex b_-, b_+. The underlying tree T has a single annular edge joining b_- and b_+ which we orient so that it points to b_+. The dynamics function $f : T \to T$ is the identity map. Then Axioms A–1 through A–9 are satisfied. We set $\mathcal{F}^{\mathcal{S}} = \mathcal{F}|\mathcal{S}_0$.

Let $\mathcal{A}_0$ be a single copy of the standard annulus. Let $\mathcal{F}^{\mathcal{A}} : \mathcal{A}_0 \to \mathcal{A}_0$ be any degree d self-map of $\mathcal{A}_0$ satisfying Axioms A–10 through A–12.

We now glue these two maps together; see Figure 9.1. Choose arbitrarily an orientation-reversing conjugacy $\rho : \partial\mathcal{S}_0 \to \partial\mathcal{A}_0$ as the topological gluing. Since $\mathcal{B}$ is completely invariant, there are no missing disks, and the Topological Compatibility Axiom A–13 is trivially satisfied. Axiom A–14 is also satisfied. To see this, note that $\sum_{q\in\mathcal{Q}_-} \omega(q) = d - 1 = \sum_{q\in\mathcal{Q}_+} \omega(q)$, so that $\sum_{q\in Q=\mathcal{Q}} \omega(q) = 2(d-1)$ and the Riemann-Hurwitz condition is satisfied. Also, by the definition of $\mathcal{Q}_\pm$, $\mathcal{Q}_\pm = \mathcal{R}_\pm^{-1}(\mathcal{R}_\pm(\mathcal{Q}_\pm))$, hence each $q \in \mathcal{Q}_\pm$ has exactly d preimages, counted with multiplicity. Hence the amalgam F is well-defined. We call such an amalgam a *mating* of R_- with R_+.

By construction this amalgam F depends only the topological gluing ρ and the choice of the annulus maps $\mathcal{F}^{\mathcal{A}}$. By the Uniqueness of Combinations Theorem, the set of classes of matings then coincides with the set of maps produced using a fixed topological gluing. How many such classes are there? By Corollary 4.6 it is enough to count the number of classes of annulus maps over $\mathbf{T}$, given in Theorem 7.1. We have $L = (d), M = (1), (\epsilon) = 1$, and therefore there are

$$|\mathbb{Z}/(L - (\epsilon)M)\mathbb{Z}| = |\mathbb{Z}/(d-1)\mathbb{Z}| = d - 1$$

such classes.

This bound coincides with the count produced from the usual definition of mating, where the polynomials are glued together using any one of the $d-1$ choices of orientation-reversing conjugacies between the actions of the two polynomials on the circles at infinity. Any mating produced in this fashion has an invariant multicurve consisting of a single curve mapping to itself by full degree d. By the Decomposition Theorem [5.1], such a map is combinatorially equivalent to an amalgam. Hence, the two definitions of mating yield precisely the same set of resulting classes.

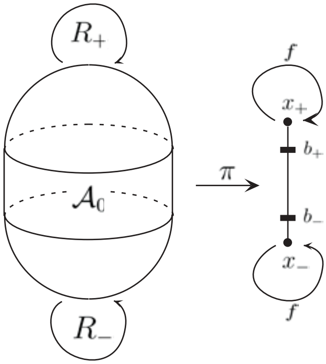

Fig. 9.1. Mating.

9.3 Generalized matings

In this section, we generalize the mating construction.

Suppose $R_-, R_+ : \mathbb{P}^1 \to \mathbb{P}^1$ are degree $d_-, d_+ \geq 2$ endomorphisms with fixed points z_-^0, z_+^0 having the same local degree $k \geq 2$ and which are isolated points of the postcritical sets P_-, P_+. We wish to combine R_-, R_+ by "gluing along boundaries of neighborhoods of z_-^0, z_+^0", much as we did for matings. In this more general setting, however, the points z_-^0, z_+^0 may have preimages other than themselves.

To realize this in our setting, begin with the map

$$R_- \sqcup R_+ : \mathbb{P}^1_- \sqcup \mathbb{P}^1_+ \to \mathbb{P}^1_- \sqcup \mathbb{P}^1$$

acting on two disjoint copies $\mathbb{P}^1_-, \mathbb{P}^1_+$ of $\mathbb{P}^1$ onto itself. We will extend this map, by adding new, disjoint spheres which are preperiodic, to a rational embellished map of spheres

$$\mathcal{R} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$$

having the same image $\mathbb{P}^1_- \sqcup \mathbb{P}^1_+$ as follows. We set

$$\mathcal{S} = \mathbb{P}^1 \times X, \quad X = X_- \sqcup X_+$$

where $X_\pm = R_\pm^{-1}(z_\pm^0)$. The set $\mathcal{Y}$ will be

$$\mathcal{Y} = \mathcal{R}^{-1}(P_- \sqcup P_+)$$

i.e. the full preimage, under $\mathcal{R}$, of the union of the postcritical sets of R_-, R_+; see Figure 9.2.

To define the extension on the copy of $\mathbb{P}^1$ corresponding to, say, $z_-^i \in R_-^{-1}(z_-^0)$, choose arbitrarily two distinct points $\overline{z}_-^i$ and q_-^i in this copy of $\mathbb{P}^1$. Map this copy holomorphically, using a degree $d_-^i = \deg(R_-, z_-^i)$ map, such that

- $\overline{z}^i_- \mapsto z^0_+$ by local degree d^i_-;
- $q^i_- \mapsto$ an element of $P_+ - \{z^0_+\}$ by local degree d^i_-.

Repeat the construction for points in $\mathcal{R}_+^{-1}(z^0_+)$. Then $\sum_i d^i_- = d_- - k$ and $\sum_i d^i_+ = d_+ - k$.

Note that with this definition of $\mathcal{R}$,

- the degree of $\mathcal{R}$ is $d = k + d_- - k + d_+ - k = d_- + d_+ - k$, and
- $\mathcal{Y}$ is the full preimage of the postcritical set $P_- \sqcup P_+$ of $\mathcal{R}$.

Thus $\mathcal{R} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ is an embellished map of spheres which is rational. Furthermore, note that $\mathcal{R}$ is canonical: once the set-theoretic data of the images of the $z^i_\pm$ and $q^i_\pm$ are specified, any two such $\mathcal{R}$ are holomorphically conjugate.

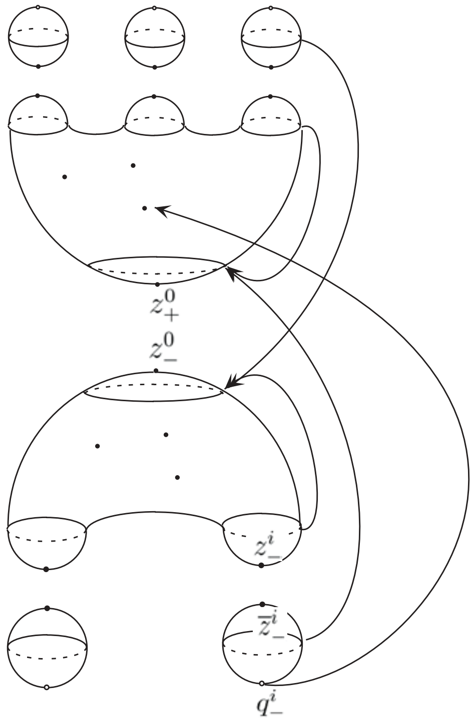

Fig. 9.2. Generalized mating.

To begin the construction of the amalgam, let

$$\mathcal{Z}_\pm = \mathcal{R}^{-1}(z^0_\pm), \quad \mathcal{Q} = \mathcal{Y} - (\mathcal{Z}_- \sqcup \mathcal{Z}_+).$$

Then $\mathcal{R} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ satisfies Embellished Axiom A–6, so that by Proposition 9.2, there exists a canonical embellished map of spheres $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ satisfying Axiom A–7. Let $\mathcal{U}, \mathcal{B}$ be as in Axiom A–7.

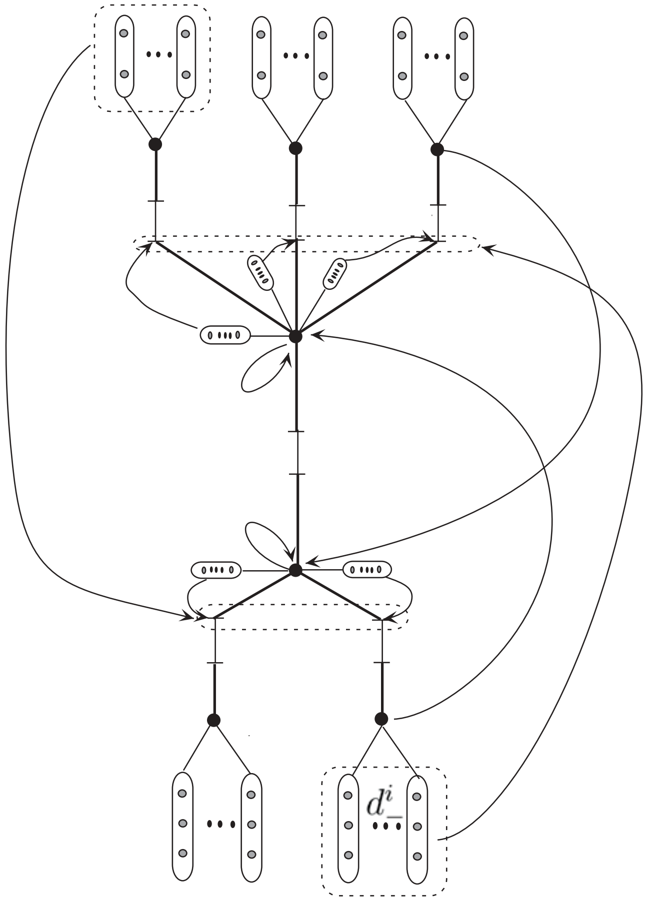

Fig. 9.3. Generalized mating. Grey dots are U-vertices. Solid dots are X-vertices. A solid oval containing vertices indicates a set of vertices which are ends of the tree and are incident to a common X-vertex. If one of the $q^i_\pm$'s is mapped to one of the $z^j_\mp$'s then the picture is slightly less symmetric.

The mapping tree has degree d and vertices $X, B_0 = B_1, U$ given by the connected components of $\mathcal{S}, \mathcal{B}, \mathcal{U}$, respectively, and edges joining

- x and b if and only if $\mathcal{B}_b \subset \mathcal{S}(x)$,
- x and u if and only if $\mathcal{U}_u \subset \mathcal{S}(x)$,
- b and $\overline{b}$ if and only if there exists $z^i_- \in \mathcal{B}_b$ and $\overline{z}^i_- \in \mathcal{B}_{\overline{b}}$, or similarly with sign $+$.

The sets $C, \mathcal{C}$ are empty.

The result looks schematically like Figure 9.3, and trivially satisfies Axioms A–1, A–2. The degree function is defined as the (unsigned) topological degree on the corresponding connected component. There is then an essentially unique way of defining the dynamics map f and projection maps $\pi^{\mathcal{S}}$ so as to satisfy Axiom A–3, and A–9(#s 1-3). The Local and Global Homogeneity Axioms (A–4) follow easily from the definitions and the fact that $d = d_- + d_+ - k$. The Tameness Axiom A–8 holds by default.

At this point we pause to verify Axiom A–9, # 4. If x corresponds to, say, the domain of R_- and $\mathcal{S}_0(x) \cap \mathcal{Q} = \emptyset$ then every critical point of R_- maps to z_-^0, which is impossible. So A–9, #4b holds. If the nonperipheral axiom A–9, #4a fails for some x, then x is preperiodic, and the intersection of the image of $\mathcal{S}_0(x)$ under $\mathcal{R}$ with, say, P_- consists of two points. This forces $|P_-| = 2$. Then R_- is conjugate to $w \mapsto = w^{d_-}$ and by combining (naively) we don't get anything new. Hence we assume $|P_-|, |P_+| > 2$.

Choose orientations of annular edges A_0 arbitrarily (Axiom A–5) and use the data to choose annulus maps $\mathcal{F}^{\mathcal{A}}$ and projections $\pi^{\mathcal{A}}$ satisfying Axioms A–10 through A–12. Choose arbitrarily a topological gluing (in our present notation, this induces the bijection $z_{\pm}^i \leftrightarrow \overline{z}_{\pm}^i$; the signs "+" and "-" have been already used for different purposes).

The critical gluing κ is defined as $\kappa(q) = q_{\pm}^i$ if $\mathcal{F}(q) = z_{\pm}^i$. We use this to define the new mapping scheme (Q, τ, ω) as in §3.2. The Topological Compatibility Axiom A–13 is clear, so after choosing arbitrarily the irrelevant missing disk maps $\mathcal{F}^{\mathcal{U}}$ we may appeal to Proposition 3.3 and conclude that we have now in hand a well-defined branched covering $F : (S^2, Q) \to (S^2, Q)$. To verify that F is indeed embellished, we need only check that the mapping scheme (Q, τ, ω) is complete to verify Axiom A–14, i.e. that $F^{-1}(F(Q)) = Q$. However, this is clear, since an element $p \in F(Q)$ is either (i) in $P_- \sqcup P_+$, or (ii) one of the q_-^i or q_+^i which is also the image under F of some other point in Q. Case (ii) occurs if and only if the corresponding point z_-^i or z_+^i is in P_- or P_+. In either case, one counts preimages and finds that all $d = d_- + d_+ - k$, with multiplicities, are accounted for. So $F^{-1}(F(Q)) = Q$ and we are done.

How many such maps F can we produce? By Corollary 4.6, it is enough to count the number of classes of annulus maps. This is determined by the common degree k and by the set of local degrees at the $z_{\pm}^i$'s. One easily finds, applying Theorem 7.1, that this number is at most $l - l/k$, where l is the least common multiple of k and these degrees. Note that if $l = k$, which is the case for matings, then this bound specializes to $k - 1$, which is the bound for matings.

9.4 Integral Lattès examples

Let $\Lambda = \omega_1\mathbb{Z} \oplus \omega_2\mathbb{Z} \subset \mathbb{C}$ be a lattice and denote by $\mathbb{C}/\Lambda$ the quotient torus. The involution $w \mapsto -w$ descends to an involution on the torus with quotient $\mathbb{P}^1$. The map $w \mapsto nw$ on $\mathbb{C}$ commutes with this involution and preserves the lattice, hence descends to a well-defined degree n^2 map on $\mathbb{P}^1$ called an *integral Lattès map*. Different choices of lattice yield (quasiconformally) conjugate maps on $\mathbb{P}^1$.

Taking $n = 2$, the resulting degree four map $R : \mathbb{P}^1 \to \mathbb{P}^1$ may be topologically described as in Figure 9.4, from which it is clear that the multicurve $\Gamma = \{\gamma_1, \gamma_2\}$ is invariant with linear map $\mathbb{R}^\Gamma$ given by the identity matrix. The associated orbifold is the $(2, 2, 2, 2)$ orbifold, and is Euclidean.

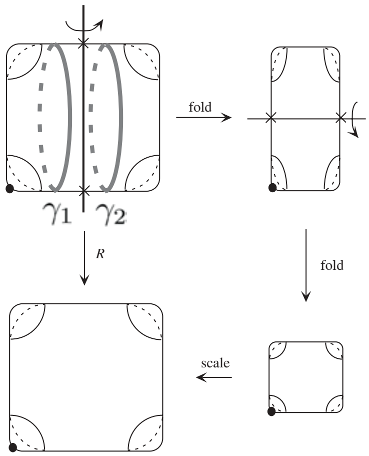

Fig. 9.4. The map on the sphere induced by $w \mapsto 2w$ on the torus can be thought of as the composition of a horizontal fold, a vertical fold, and a rescaling. The unique periodic point in the postcritical set is highlighted as a black dot. There are six critical points, four in the centers of the edges of the pillowcase, one in the front center, and one in the back center. The postcritical set consists of the four "corners" of the pillowcase. Note that the two thick vertical gray curves form an invariant multicurve.

The decomposition of the sphere and the mapping tree is illustrated in Figure 9.5. The set $\mathcal{U}$ is empty. The dynamics map $f : T \to T$ acts on X-vertices sending x_1, x_0 to x_0 by degree two and x_2 to x_1 by degree four. With respect to the indicated orientations, f reverses orientations on connected components of A_1 which lie between x_2 and x_1. There are two missing annuli $\mathcal{C}_1, \mathcal{C}_2$.

After decomposing, the resulting map of spheres $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ is in fact rational. The unique periodic component $\mathcal{S}(x_0)$ of $\mathcal{S}$ has period one and on this component the map $\mathcal{F}$ is holomorphically conjugate to $z \mapsto z^2 - 2$.

Note, however, that we could have taken Γ to be the pair of horizontal curves instead, and the resulting analysis would yield the same results. Hence, the map R can be decomposed into the same set of sphere maps in two distinct ways.

Moreover, the Twist Theorem [8.2] applied to the vertical multicurve Γ shows, after a brief but straightforward calculation, that the mapping class group contains a rank-one free abelian subgroup generated by performing

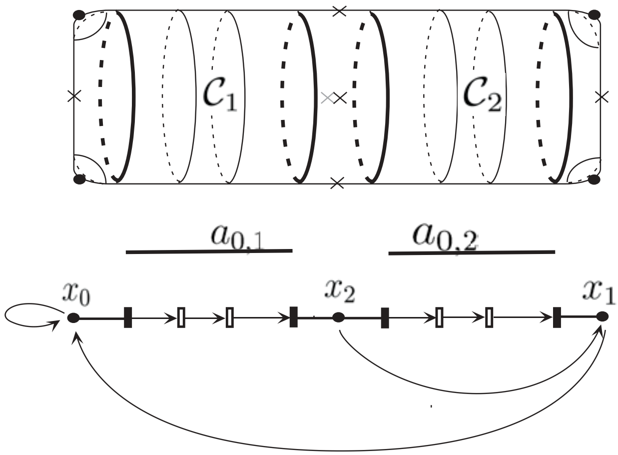

Fig. 9.5. Structure of a Lattès example. There are two annuli in $\mathcal{A}_0$ bounded by the bold curves lying above the annular edges labelled $a_{0,1}, a_{0,2}$. Each of these annuli has two preimages in $\mathcal{A}_1$, both mapping by degree two, with those on the left preserving and those on the right reversing orientation of edges. Note that one obtains the same picture using annuli which are horizontal.

double Dehn twists, in the same directions, simultaneously about γ_1 and γ_2. Of course, the same holds true in the vertical directions as well.

10

Canonical Decomposition Theorem

In this section, we consider the structure of *postcritically finite* branched coverings from the sphere to itself. We give, in §10.2 below, a structure theorem which says that *any* such map is, in a *canonical* fashion, decomposable into "pieces", each of which is of one of three possible types:

1. (elliptic case) a homeomorphism of spheres,
2. (parabolic case) covered by a homeomorphism of planes, or
3. (hyperbolic, rational case) equivalent to a rational map of spheres.

We make these notions precise in the next two subsections.

10.1 Cycles of a map of spheres, and their orbifolds

Let $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ be an embellished map of spheres for which $\mathcal{Y}$ is finite. Let X denote the set of connected components of $\mathcal{S}$ and let $f : X \to X$ be the induced map. Let $\deg : X \to \mathbb{N}$ be given by $\deg(x) = \deg(f, x)$.

We can think of the triple $(X, f, \deg)$ as a directed graph K with directed edges $x \to f(x)$ labelled, or weighted, by the degree $\deg(x)$.

Definitions and notation.

- A *component* of $\mathcal{F}$ is a connected component of K, viewed as an undirected graph.
- A *cycle* of $\mathcal{F}$ is a cycle of K. Note that each component contains a unique cycle, since f is a function from X to X.
- For a cycle ζ we set
$$\mathcal{S}^\zeta = \cup_{x\in\zeta}\mathcal{S}(x)$$
and
$$\mathcal{Y}^\zeta = \mathcal{S}^\zeta \cap \mathcal{Y}.$$

K.M. Pilgrim: LNM 1827, pp. 105–109, 2003.

Thus

$$\mathcal{F}^\zeta = \mathcal{F}|\mathcal{S}^\zeta : (\mathcal{S}^\zeta, \mathcal{Y}^\zeta) \to (\mathcal{S}^\zeta, \mathcal{Y}^\zeta)$$

is again an embellished map of spheres.

- The *orbifold* $\mathcal{O}^\zeta$ associated to a cycle ζ has, as its underlying topological space, the space $\mathcal{S}^\zeta$. The weight $\nu(p)$ at $p \in \mathcal{S}^\zeta$ is given by the least common multiple, over all iterated preimages p' of *the restricted map* $\mathcal{F}^\zeta$, of the local degree of $\mathcal{F}^\zeta$ at p'. Note that any singular points (where $\nu(p) > 1$) must lie in $\mathcal{Y}$. We let $\mathcal{O}^\zeta(x)$ denote the orbifold given by the xth connected component of $\mathcal{O}^\zeta$.
- The *Euler characteristic* of a cycle ζ is given by

$$\sum_{x\in\zeta}\left(2 - \sum_{p\in\mathcal{S}(x)}\left(1 - \frac{1}{\nu(p)}\right)\right).$$

When the cycle has period one, this generalizes the usual definition of the orbifold associated to a branched covering, cf. [DH3] and [McM6]. Note that the Euler characteristic of a cycle is the sum of the Euler characteristics over all components.

It will turn out (Lemma 10.1 below) that, although the numerical values of the Euler characteristics of connected orbifolds $\mathcal{S}(x)$ may vary as x varies within ζ, the sign of this value is constant in the cycle. Thus:

- It makes sense to speak of a cycle as being *elliptic, parabolic, hyperbolic* according as the Euler characteristic of the cycle is positive, zero, or negative.

Lemma 10.1.

1. *The sign of $\mathcal{O}^\zeta(x)$ is independent of $x \in \zeta$.*
2. *$\mathcal{O}^\zeta$ is elliptic if and only if $\mathcal{F}^\zeta$ is a homeomorphism.*
3. *$\mathcal{O}^\zeta$ is parabolic if and only if $\mathcal{F}^\zeta$ is a covering map of orbifolds. In this case, $\mathcal{F}^\zeta$ lifts under the orbifold universal covering map $\pi : \mathbb{R}^2 \times \zeta \to \mathcal{O}^\zeta$ to a homeomorphism $\widetilde{\mathcal{F}}^\zeta : \mathbb{R}^2 \times \zeta \to \mathbb{R}^2 \times \zeta$.*

Proof: (Cf. [DH3], Prop. 9.1)

1. Let $\widehat{\mathcal{O}}^\zeta$ be the orbifold with the same underlying space $\mathcal{S}$ and with weight

$$\widehat{\nu}(p) = \frac{\nu(\mathcal{F}^\zeta(p))}{\deg(\mathcal{F}^\zeta, p)}.$$

Then $\mathcal{F} : \widehat{\mathcal{O}}^\zeta \to \mathcal{O}^\zeta$ is a covering map of orbifolds. Write the cycle ζ as $x_0 \mapsto x_1 \mapsto ... \mapsto x_{n-1}$. and let d_i be the degree of $\mathcal{F}^\zeta$ on $\mathcal{S}(x_i)$. Then

$$\chi(\widehat{\mathcal{O}}^\zeta(x_i)) = d_i \cdot \chi(\mathcal{O}^\zeta(x_{i+1})) \leq \chi(\mathcal{O}^\zeta(x_i)).$$

Applying this inequality around the cycle yields **1.**

2. In particular for each i we have

$$(d-1)\chi(\mathcal{O}^\zeta(x_i)) \leq 0 \tag{10.1}$$

where $d = \Pi_{i=1}^n d_i$ is the product of the degrees around the cycle. Thus if the orbifold is elliptic, then $d = 1$ and so $\mathcal{F}^\zeta$ is a homeomorphism. Conversely, if $\mathcal{F}^\zeta$ is a homeomorphism, then in fact there are no singular points at all, the orbifold is actually a manifold, and each component has Euler characteristic $\chi(S^2) = 2$.

3. $\mathcal{O}^\zeta$ is parabolic if and only if $d > 1$ and equality holds in Inequality (10.1). Equivalently, $\widehat{\mathcal{O}}^\zeta(x_i) = \mathcal{O}^\zeta(x_i)$ and so $\mathcal{F}^\zeta$ is a (self) covering map of the orbifold $\mathcal{O}^\zeta$ to itself.

In this case $\mathcal{F}^\zeta$ lifts to a homeomorphism $\widetilde{\mathcal{F}}^\zeta$ of the disjoint union of the universal covers of the $\mathcal{O}^\zeta(x_i)$, i.e. to n copies of the plane, which commutes with the (orbifold) covering transformations.

□

10.2 Statement of Canonical Decomposition Theorem

Theorem 10.2 (Canonical decomposition theorem). *Let $F : (S^2, Q) \to (S^2, Q)$ be a postcritically finite embellished branched covering. Then F decomposes along a canonically determined, possibly empty multicurve Γ^c into an embellished map $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ of spheres such that every cycle of $\mathcal{F}$ is either*

1. *elliptic, i.e. a homeomorphism,*
2. *parabolic, i.e. covered by a homeomorphism of a set of Euclidean planes to itself, or*
3. *hyperbolic and rational, i.e. equivalent to an embellished postcritically finite rational map with hyperbolic orbifold.*

Furthermore, Γ^c is empty if and only if the map F itself, considered as a cycle of length one, is one of the three types above.

In the second case above, the cycle may or may not be equivalent to a rational map of spheres.

Since Γ^c is canonical, we have under the same hypotheses

Corollary 10.3 ($\mathrm{Mod}(F)$ **reduces along** Γ^c**).** *For each $\alpha \in \mathrm{Mod}(G)$, $\alpha(\Gamma^c) = \Gamma^c$, up to isotopy relative to Q.*

10.3 Proof of Canonical Decomposition Theorem

10.3.1 Characterization of rational cycles with hyperbolic orbifold

The Douady-Hubbard-Thurston characterization of postcritically finite rational maps with hyperbolic orbifold as branched coverings generalizes immediately to the setting of embellished maps of spheres.

First, it is enough to consider only a single cycle, since an invariant complex structure on a cycle can always be lifted to a holomorphic map between a preperiodic component and its image.

Next, to an embellished map of spheres with a single cycle, one considers an induced map $\sigma_{\mathcal{F}}$ on a product $\mathcal{T}$ of Teichmüller spaces. The analytic properties of the map are identical to those considered in [DH3] in the case when the orbifold is hyperbolic: the map (iterated $2n$ times, where n is the length of the cycle) is contracting, though not uniformly so. Existence of a fixed point is equivalent to being equivalent to a rational map of spheres. Choose arbitrarily a basepoint $\tau_0 \in \mathcal{T}$, and let $\tau_i = \sigma_{\mathcal{F}}^{\circ i}(\tau_0)$. If a fixed point does not exist, then the Riemann surfaces determined by τ_i develop short geodesics , and the same arguments provide as before a topological obstruction in the form of an invariant multicurve Γ such that the leading eigenvalue of $\mathcal{F}_\Gamma : \mathbb{R}^\Gamma \to \mathbb{R}^\Gamma$ is at least one. The enlargement of the postcritical set to a larger, forward-invariant finite set in no way alters the character of the proof or the statement of the results.

In [Pil2] the Douady-Hubbard characterization was refined to show that in fact not only do short geodesics develop, but that the curves whose lengths converge to zero form an invariant multicurve and a topological obstruction; see 1.2.

Precisely the same arguments yield an analogous theorem for maps of spheres.

Theorem 10.4 (Canonical obstructions). *Let $\mathcal{F} : (\mathcal{S}, \mathcal{Y}) \to (\mathcal{S}, \mathcal{Y})$ be an embellished map of spheres consisting of a single cycle. Suppose $\mathcal{Y}$ is finite, and that the orbifold of $\mathcal{F}$ is hyperbolic.*

Let Γ_c denote the set of all homotopy classes of nonperipheral, simple closed curves γ in $\mathcal{S} - \mathcal{Y}$ such that the hyperbolic length of γ on τ_i converges to zero under as $i \to \infty$. Then Γ_c is independent of τ_i. Moreover:

1. *If Γ_c is empty, then $\mathcal{F}$ is equivalent to a rational map.*
2. *Otherwise, Γ_c is an invariant multicurve for which the leading eigenvalue $\lambda(\mathcal{F}, \Gamma_c)$ is ≥ 1, and hence is a canonically defined obstruction to the existence of a rational map of spheres equivalent to $\mathcal{F}$.*

The characterization of embellished maps with parabolic orbifold, however, is a bit more subtle. The addition of other marked points causes potential complications. We do not pursue such a classification here.

10.3.2 Conclusion of proof

Suppose F itself is neither a homeomorphism, nor covered by a homeomorphism of the plane, nor has hyperbolic orbifold and is equivalent to a rational map.

Then the orbifold of F is hyperbolic, and F is not equivalent to a rational map.

By the Canonical Obstructions Theorem, (Thm. 10.4), there is a canonical topological obstruction Γ^0 which is an F-invariant multicurve.

By the Decomposition Theorem [5.1], decomposing F along Γ^0 yields a map of spheres $\mathcal{F}^1 : (\mathcal{S}^1, \mathcal{Y}^1) \to (\mathcal{S}^1, \mathcal{Y}^1)$. By the Uniqueness of Decompositions Theorem [6.1], the combinatorial class of $\mathcal{F}^1$ depends only on the equivalence class of pairs (F, Γ^0).

We now proceed by induction. At the inductive step, we have an F-invariant multicurve Γ^n. Decomposing F along Γ^n produces a map of spheres $\mathcal{F}^n : (\mathcal{S}^n, \mathcal{Y}^n) \to (\mathcal{S}^n, \mathcal{Y}^n)$ whose combinatorial class depends only on the equivalence class of pair (F, Γ^n). Either conclusions (1)-(3) of the Theorem hold, or else there exists a cycle ζ of $\mathcal{F}^n$ not equivalent to a rational map and whose orbifold is hyperbolic.

Apply the Canonical Obstructions theorem [10.4] to the map $\mathcal{F}^n$ on the cycle ζ to obtain a nonempty $(\mathcal{F}^n)^\zeta$-invariant multicurve Γ^ζ_{n+1}. This multicurve pulls back under $\mathcal{F}^n$ to a canonically defined invariant multicurve $\Gamma^{\mathcal{S}}_{n+1}$ for $\mathcal{F}^n$. By Proposition 3.6, *Invariant multicurves persist*, the image Γ_{n+1} of $\Gamma^{\mathcal{S}}_{n+1}$ under the inclusion $\mathcal{S}^n_0 \hookrightarrow S^2 - Q$ is well-defined and, up to isotopy, is disjoint from Γ^n. By the same proposition,

$$\Gamma^{n+1} = \Gamma_{n+1} \sqcup \Gamma^n$$

is again an F-invariant multicurve.

Thus, the cardinality of the Γ^n's is strictly increasing in n. Since the number of elements of Γ^n can be at most $|Q| - 3$, the inductive process must conclude at some point, i.e. at some stage, no further decomposition is possible. That is, all cycles of $\mathcal{F}^n$ have the properties listed in the Theorem. □

We conjecture that in fact the procedure stops at $n = 1$, i.e. that after decomposing F along the canonical obstruction Γ_c picked out by the iterative algorithm on Teichmüller space, every cycle with hyperbolic orbifold is equivalent to a rational map. Theorem 1.3 is evidence in favor of this conjecture.

References

[Ahm] Davood Ahmadi. *Dynamics of Certain Rational Maps of Degree Two.* PhD thesis, University of Liverpool, 1991.

[AF] Lluís Alsedà and Núria Fagella. Dynamics on Hubbard trees. *Fund. Math.* **164**(2000), 115–141.

[Bar1] Krzysztof Barański. From Newton's method to exotic basins. I. The parameter space. *Fund. Math.* **158**(1998), 249–288.

[Bar2] Krzysztof Barański. From Newton's method to exotic basins II. Bifurcation of the Mandelbrot-like sets. *Fund. Math.* **168**(2001), 1–55.

[BGN] L. Bartholdi, R. Grigorchuk, and V. Nekrashevych. From fractal groups to fractal sets. In *Fractals in Graz 2001*, page 25–118 Birkhäuser, 283, 2002. arXiv:math.GR/0203244v1.

[Ber] J. Bernard. *Dynamique des perturbations d'un example de Lattès.* PhD thesis, Univ. Paris-Sud Orsay, 1994.

[BFH] Ben Bielefeld, Yuval Fisher, and John Hubbard. The Classification of Critically Preperiodic Polynomials as Dynamical Systems. *Journal of the American Mathematical Society* **5**(1992).

[Bir] J. Birman. *Braids, Links, and Mapping Class Groups.* Princeton University Press Annals of Mathematics Studies 82, 1975.

[BL1] A. Blokh and G. Levin. An inequality for laminations, Julia sets and "growing trees". *Ergodic Theory Dynam. Systems* **22**(2002), 63–97.

[BL2] A. Blokh and G. Levin. On dynamics of vertices of locally connected polynomial Julia sets. *Proc. Amer. Math. Soc.* **130**(2002), 3219–3230 (electronic).

[BS] Philip L. Bowers and Kenneth Stephenson. Uniformizing dessins and Belyi maps via circle packing. *to appear, AMS Memoirs.*

[BD] Bodil Branner and Adrien Douady. Surgery on complex polynomials. In *Holomorphic Dynamics: Proceedings, Mexico.* Springer-Verlag Lecture Notes in Mathematics No. 1345, 1986.

[BF] Bodil Branner and Núria Fagella. Extensions of homeomorphisms between limbs of the Mandelbrot set. *Conform. Geom. Dyn.* **5**(2001), 100–139 (electronic).

[BBL+] Eva Brezin, Rosemary Byrne, Joshua Levy, Kevin Pilgrim, and Kelly Plummer. A census of rational maps. *Conform. Geom. Dyn.* **4**(2000), 35–74 (electronic).

K.M. Pilgrim: LNM 1827, pp. 111–116, 2003.

[Bro] David Brown. *Using Spider Theory to explore parameter spaces.* PhD thesis, Cornell University, 2001.

[CFP1] J. W. Cannon, W. J. Floyd, and W. R. Parry. Sufficiently rich families of planar rings. *Ann. Acad. Sci. Fenn. Math.* **24**(1999), 265–304.

[CFP2] J. W. Cannon, W. J. Floyd, and W. R. Parry. Finite subdivision rules. *Conform. Geom. Dyn.* **5**(2001), 153–196 (electronic).

[CFP3] J. W. Cannon, W. J. Floyd, and W. R. Parry. Constructing rational maps from subdivision rules. *to appear, Conform. Geom. Dyn. (American Math Society electronic journal).*

[Can] James W. Cannon. The combinatorial Riemann mapping theorem. *Acta Math.* **173**(1994), 155–234.

[CP] Michel Coornaert and Athanase Papadopoulos. *Symbolic dynamics and hyperbolic groups*, volume 1539 of *Lecture Notes in Mathematics.* Springer-Verlag, Berlin, 1993.

[Cui] Guizhen Cui. Geometrically finite rational maps with given combinatorics. Manuscript, 2000.

[CJS] Guizhen Cui, Yunping Jiang, and Dennis Sullivan. Dynamics of geometrically finite rational maps. Manuscript, 1996.

[DR] R. L. Devaney and M. Moreno Rocha. Geometry of the antennas in the Mandelbrot set. *Fractals* **10**(2002), 39–46.

[Dev] Robert L. Devaney. The fractal geometry of the Mandelbrot set. 2. How to count and how to add. *Fractals* **3**(1995), 629–640. Symposium in Honor of Benoit Mandelbrot (Curaçao, 1995).

[Dou] A. Douady. Algorithms for computing angles in the Mandelbrot set. In *Chaotic dynamics and fractals (Atlanta, Ga., 1985)*, pages 155–168. Academic Press, Orlando, FL, 1986.

[DH1] A. Douady and John Hubbard. Etude Dynamique des Polynomes Complexes, I, II. *Publications Mathematique Orsay* (1984/5).

[DH2] A. Douady and John Hubbard. On the Dynamics of Polynomial-Like Mappings. *Ann. Scient. Ec. Norm. Sup.* 4^e **série t. 18**(1985), 287–343.

[DH3] A. Douady and John Hubbard. A Proof of Thurston's Topological Characterization of Rational Functions. *Acta. Math.* **171**(1993), 263–297.

[Dou] Adrien Douady. A Description of Compact Sets in $\mathbb{C}$. In L. R. Goldberg and A. V. Phillips, editors, *Topological Methods in Modern Mathematics: Proceedings of the Milnor Symposium, Stony Brook, NY.* Publish or Perish, 1993.

[Eme] Nathaniel D. Emerson. Dynamics of polynomials with disconnected Julia sets. *Discrete and continuous dynamical systems* **9**(2003), 801–834.

[EKT] Adam Epstein, Linda Keen, and Charles Tresser. The set of maps $F_{a,b}\colon x \mapsto x + a + (b/2\pi)\sin(2\pi x)$ with any given rotation interval is contractible. *Comm. Math. Phys.* **173**(1995), 313–333.

[EY] Adam Epstein and Michael Yampolsky. Geography of the cubic connectedness locus: intertwining surgery. *Ann. Sci. École Norm. Sup. (4)* **32**(1999), 151–185.

[Eps1] Adam L. Epstein. Counterexamples to the quadratic mating conjecture. Manuscript in preparation.

[Eps2] Adam Lawrence Epstein. Bounded hyperbolic components of quadratic rational maps. *Ergodic Theory Dynam. Systems* **20**(2000), 727–748.

[Fri] David Fried. Finitely presented dynamical systems. *Ergodic Theory Dynam. Systems* **7**(1987), 489–507.

[Gan] F. R. Gantmacher. *Theory of Matrices.* Chelsea, 1959.

[GM] Lisa R. Goldberg and John Milnor. Fixed points of polynomial maps. II. Fixed point portraits. *Ann. Sci. École Norm. Sup. (4)* **26**(1993), 51–98.

[Haï] Peter Haïssinsky. Modulation dans l'ensemble de Mandelbrot. In *The Mandelbrot set, theme and variations*, pages 37–65. Cambridge Univ. Press, Cambridge, 2000.

[Hea] Janet Head. *The combinatorics of Newton's method for cubic polynomials.* PhD thesis, Cornell University, 1992.

[Hen] Christian Henriksen. Lack of combinatorial rigidity in degree 3. *to appear, Trans. AMS.*

[Hub] J. H. Hubbard. Local connectivity of Julia sets and bifurcation loci: three theorems of J. C. Yoccoz. In L. R. Goldberg and A. V. Phillips, editors, *Topological Methods in Modern Mathematics: Proceedings of the Milnor Symposium, Stony Brook, NY.* Publish or Perish, 1993.

[IP] Clemens Inninger and Franz Peherstorfer. A new simple class of rational functions whose Julia set is the whole Riemann sphere. *J. London Math. Soc. (2)* **65**(2002), 453–463.

[JS] William H. Jaco and Peter B. Shalen. Seifert fibered spaces in 3-manifolds. *Mem. Amer. Math. Soc.* **21**(1979), viii+192.

[Joh] Klaus Johannson. On the mapping class group of simple 3-manifolds. In *Topology of low-dimensional manifolds (Proc. Second Sussex Conf., Chelwood Gate, 1977)*, pages 48–66. Springer, Berlin, 1979.

[Kam1] Atsushi Kameyama. Julia sets of postcritically finite rational maps and topological self-similar sets. *Nonlinearity* **13**(2000), 165–188.

[Kam2] Atsushi Kameyama. The Thurston equivalence for postcritically finite branched coverings. *Osaka J. Math.* **38**(2001), 565–610.

[Kam3] Atsushi Kameyama. On Julia sets of postcritically finite branched coverings. I. Coding of Julia sets. *J. Math. Soc. Japan* **55**(2003), 439–454.

[Kam4] Atsushi Kameyama. On Julia sets of postcritically finite branched coverings. II. S^1-parametrization of Julia sets. *J. Math. Soc. Japan* **55**(2003), 455–468.

[Kam5] Atsushi Kameyama. Coding and tiling of Julia sets for subhyperbolic rationalmaps. *to appear, Erg. Thy. & Dynam. Sys.*

[Kau] Virpi Kauko. Trees of visible components in the Mandelbrot set. *Fund. Math.* **164**(2000), 41–60.

[Kel] Karsten Keller. *Invariant factors, Julia equivalences and the (abstract) Mandelbrot set*, volume 1732 of *Lecture Notes in Mathematics.* Springer-Verlag, Berlin, 2000.

[Kiw] Jan Kiwi. Rational laminations of complex polynomials. In *Laminations and foliations in dynamics, geometry and topology (Stony Brook, NY, 1998)*, volume 269 of *Contemp. Math.*, pages 111–154. Amer. Math. Soc., Providence, RI, 2001.

[Lav] P. Lavaurs. Une Description Combinatoire de L'Involution Definie par M sur les Rationels a Denominateur Impair. *C. R. Acad. Sci.* **303**(1986), 143–146.

[Lev] Silvio Levy. *Critically Finite Rational Maps.* PhD thesis, Princeton University, 1986.

[Luo] Jiaqi Luo. *Combinatorics and holomorphic dynamics: captures, matings, Newton's method.* PhD thesis, Cornell University, 1995.

[Lyu] Mikhail Lyubich. The quadratic family as a qualitatively solvable model of chaos. *Notices Amer. Math. Soc.* **47**(2000), 1042–1052.

[Mak1] P. M. Makienko. Modular groups of certain rational functions. *Dokl. Akad. Nauk SSSR* **310**(1990), 793–794.

[Mak2] Peter M. Makienko. Unbounded components in parameter space of rational maps. *Conform. Geom. Dyn.* **4**(2000), 1–21 (electronic).

[Mar] Albert Marden. The geometry of finitely generated kleinian groups. *Ann. of Math. (2)* **99**(1974), 383–462.

[Mas] B. Maskit. *Kleinian Groups.* Springer-Verlag, 1988.

[McM1] Curt McMullen. Automorphisms of Rational Maps. In *Holomorphic Functions and Moduli.* MSRI Publications, 1986.

[McM2] Curt McMullen. Families of rational maps and iterative root-finding algorithms. *Annals of Math.* **125**(1987), 467–493.

[McM3] Curt McMullen. Iteration on Teichmüller space. *Invent. math.* **99**(1990), 425–454.

[McM4] Curt McMullen. Rational maps and Kleinian groups. In *Proceedings of the International Congress of Mathematicians, Vol. I, II (Kyoto, 1990)*, pages 889–899, Tokyo, 1991. Math. Soc. Japan.

[McM5] Curtis T. McMullen. The classification of conformal dynamical systems. In *Current developments in mathematics, 1995 (Cambridge, MA)*, pages 323–360. Internat. Press, Cambridge, MA, 1994.

[McM6] Curtis T. McMullen. *Complex dynamics and renormalization.* Princeton University Press, Princeton, NJ, 1994.

[McM7] Curtis T. McMullen. Self-similarity of Siegel disks and Hausdorff dimension of Julia sets. *Acta Math.* **180**(1998), 247–292.

[MT] J. Milnor and W. Thurston. On iterated maps of the interval. In *Dynamical systems.* Springer Verlag Lecture Notes 1342, 1988.

[Mil1] John Milnor. Self-Similarity and Hairiness in the Mandelbrot Set. In Martin C. Tangora, editor, *Computers in Geometry and Topology.* Marcel Dekker Lecture Notes in Pure and Applied Mathematics Volume 114, 1989.

[Mil2] John Milnor. Hyperbolic components in spaces of polynomial maps. *SUNY Stonybrook Institute for Mathematical Sciences Preprint* (1992/3).

[Mil3] John Milnor. Geometry and dynamics of quadratic rational maps. *Experiment. Math.* **2**(1993), 37–83. With an appendix by the author and Tan Lei.

[Mil4] John Milnor. *Dynamics in one complex variable.* Friedr. Vieweg & Sohn, Braunschweig, 1999. Introductory lectures.

[Mil5] John Milnor. Pasting together Julia sets: a worked out example of mating. at www.math.sunysb.edu/∼jack/PREPR/index.html, 2000.

[Pet] Carsten Lunde Petersen. No elliptic limits for quadratic maps. *Ergodic Theory Dynam. Systems* **19**(1999), 127–141.

[PL] Kevin Pilgrim and Tan Lei. Rational maps with disconnected Julia set. *Astérisque* (2000), xiv, 349–384. Géométrie complexe et systèmes dynamiques (Orsay, 1995).

[Pil1] Kevin M. Pilgrim. Dessins d'enfants and Hubbard trees. *Ann. Sci. École Norm. Sup. (4)* **33**(2000), 671–693.

[Pil2] Kevin M. Pilgrim. Canonical Thurston obstructions. *Adv. Math.* **158**(2001), 154–168.

[Pil3] Kevin M. Pilgrim. An algebraic formulation of Thurston's combinatorial equivalence. *to appear, Proc. AMS.*, 2002.

[Pil4] Kevin M. Pilgrim. A Hurwitz-like characterization of Thurston combinatorial classes. to appear, Osaka J. Math.

[Pil5] Kevin M. Pilgrim. *Cylinders for Iterated Rational Maps.* PhD thesis, University of California at Berkeley, May, 1994.

[PT] Kevin M. Pilgrim and Lei Tan. Combining rational maps and controlling obstructions. *Ergodic Theory Dynam. Systems* **18**(1998), 221–245.

[Poi1] Alfredo Poirier. On Post-Critically Finite Polynomials Part One: Critical Portraits. *SUNY Stony Brook Institute for Mathematical Sciences Preprint* (1993/5).

[Poi2] Alfredo Poirier. On Post-Critically Finite Polynomials Part Two: Hubbard Trees. *SUNY Stony Brook Institute for Mathematical Sciences Preprint* (1993/7).

[Ree1] Mary Rees. Realization of matings of polynomials as rational maps of degree two. 1986.

[Ree2] Mary Rees. Components of Degree Two Hyperbolic Rational Maps. *Inventiones Mathematicae* **100**(1990), 357–382.

[Ree3] Mary Rees. A Partial Description of the Parameter Space of Rational Maps of Degree Two: Part I. *Acta Math.* **168**(1992), 11–87.

[Ree4] Mary Rees. A partial description of the parameter space of rational maps of degree two. II. *Proc. London Math. Soc. (3)* **70**(1995), 644–690.

[Ree5] Mary Rees. Views of parameter space: topographer and resident. url: http://www.liv.ac.uk/∼ maryrees/maryrees.homepage.html, 2002.

[Sch] Dierk Schleicher. Rational parameter rays of the Mandelbrot set. *Astérisque* (2000), xiv–xv, 405–443. Géométrie complexe et systèmes dynamiques (Orsay, 1995).

[Shi1] M. Shishikura. On the Quasiconformal Surgery of Rational Functions. *Ann. Sci. Éc. Norm. Sup.* **20**(1987).

[Shi2] Mitsuhiro Shishikura. Surgery of complex analytic dynamical systems. In *Dynamical systems and nonlinear oscillations (Kyoto, 1985)*, volume 1 of *World Sci. Adv. Ser. Dynam. Systems*, pages 93–105. World Sci. Publishing, Singapore, 1986.

[Shi3] Mitsuhiro Shishikura. Trees associated with the configuration of Herman rings. *Ergodic Theory Dynam. Systems* **9**(1989), 543–560.

[Shi4] Mitsuhiro Shishikura. On a theorem of M. Rees for matings of polynomials. In *The Mandelbrot set, theme and variations*, pages 289–305. Cambridge Univ. Press, Cambridge, 2000.

[SL] Mitsuhiro Shishikura and Tan Lei. A family of cubic rational maps and matings of cubic polynomials. *Experiment. Math.* **9**(2000), 29–53.

[Sti] James R. P. Stimson. *Degree two rational maps with a periodic critical point.* PhD thesis, University of Liverpool, 1993.

[Sul] Dennis Sullivan. Quasiconformal Homeomorphisms in Dynamics, Topology, and Geometry. *Proceedings of the International Congress of Mathematicians, Kyoto, Japan* (1990).

[Tan1] Tan Lei. Accouplements des polynômes complexes. *C. R. Acad. Sc. Paris Série I* **302**(1986), 635–638.

[Tan2] Tan Lei. Matings of quadratic polynomials. *Ergodic Theory Dynamical Systems* **12**(1992), 589–620.

[Tan3] Tan Lei. Branched coverings and cubic Newton maps. *Fund. Math.* **154**(1997), 207–260.

[Tan4] Tan Lei, editor. *The Mandelbrot set, theme and variations*, volume 274 of *London Mathematical Society Lecture Note Series.* Cambridge University Press, Cambridge, 2000.

[Tan5] Tan Lei. On pinching deformations of rational maps. Université de Cergy-Pontoise preprint, September 2000.

[Thu1] W. P. Thurston. Hyperbolic structures on 3-manifolds I: Deformations of acylindrical manifolds. *Annals of Math.* **124**(1986), 203–246.

[Thu2] William P. Thurston. On the Combinatorics and Dynamics of Iterated Rational Maps. *Princeton University and IAS Preprint.* 1985.

[Wit] Ben Wittner. *On the Bifurcation Loci of Rational Maps of Degree Two.* PhD thesis, Cornell University, 1988.

[YZ] Michael Yampolsky and Saeed Zakeri. Mating Siegel quadratic polynomials. *J. Amer. Math. Soc.* **14**(2001), 25–78 (electronic).

[Yin] Yin Yongchen. The topology of Julia sets for polynomials. *Sci. China Ser. A* **45**(2002), 1020–1024.

Index