ALMA MATER STUDIORUM · UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN MATEMATICA
CICLO XXIX

Settore Concorsuale di afferenza: 01/A2

Settore Scientifico disciplinare: MAT/03

THE GEOMETRY OF BRANCHED COMPLEX PROJECTIVE STRUCTURES ON SURFACES

Presentata da: Lorenzo Ruffoni

Coordinatore Dottorato	Relatore
Chiar.ma Prof.ssa	Chiar.mo Prof.
Giovanna Citti	Stefano Francaviglia

Esame finale 2017

To my parents

Contents

List of Figures

Introduction

Complex projective structures on a surface are geometric structures locally modelled on the geometry of the Riemann sphere $\mathbb{CP}^1 = \mathbb{C} \cup \{\infty\}$ which is defined by its group of automorphisms $\mathrm{PSL}_2\mathbb{C}$ acting by Möbius transformations

$$\mathrm{PSL}_2\mathbb{C} \times \mathbb{CP}^1 \to \mathbb{CP}^1, \begin{pmatrix} a & b \\ c & d \end{pmatrix}, z \mapsto \frac{az+b}{cz+d}$$

More precisely such a structure is the datum of local charts which take values in open sets of $\mathbb{CP}^1$ and are such that the change of coordinates is given by the restriction of a global Möbius transformation $g \in \mathrm{PSL}_2\mathbb{C}$. Given such an atlas, local charts and local change of coordinates can be analytically continued along paths on the surface to obtain respectively a map $dev : \widetilde{S} \to \mathbb{CP}^1$ on the universal cover $\widetilde{S}$ of S and a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ of its fundamental group; dev is called the developing map and ρ the holonomy of the structure, and dev is equivariant with respect to ρ, i.e. the following equation holds

$$dev \circ \gamma = \rho(\gamma) \cdot dev \qquad \forall\, \gamma \in \pi_1(S)$$

Since $\mathrm{PSL}_2\mathbb{C}$ acts holomorphically on $\mathbb{CP}^1$, a complex projective structure has clearly an underlying complex structure. Indeed historically these structures appeared at first in the classical work of Poincaré (see [31]) on the uniformization problem for Riemann surfaces, which consists in realising a Riemann surface as a quotient of an open domain in $\mathbb{CP}^1$ by a discrete group of biholomorphisms. When the genus of the surface is $g \geq 2$ this theory guarantees that one can always choose the domain to be the upper-half plane $\mathcal{H}^+ \subset \mathbb{C} \subset \mathbb{CP}^1$ and the group to be a discrete torsion-free subgroup $\Gamma \subset \mathrm{PSL}_2\mathbb{R}$, a so-called Fuchsian group (see [12, Chapter IV]). Later Bers extended this theory in [6] to consider the problem of *simultaneous* uniformization of two Riemann surfaces: the new problem was to realise a couple of Riemann surfaces as quotient of two disjoint open domains in $\mathbb{CP}^1$ by the *same* discrete group of biholomorphisms. To achieve this one needs to allow for more general discrete subgroups $\Gamma \subset \mathrm{PSL}_2\mathbb{C}$ which are known as quasi-Fuchsian groups and arise as quasi-conformal deformations of Fuchsian groups. In both cases one obtains a complex projective structure on the surface for which the developing map dev is a diffeomorphism with a disk-like subset of $\mathbb{CP}^1$. Maskit was then able a few years later in [27] to obtain many examples of exotic structures by applying to the previously known examples an explicit geometric surgery, known as grafting. This deformation consists in replacing a simple closed curve γ on S with an annulus of the form $(\mathbb{C} \setminus i\mathbb{R}^+)/(z \mapsto \lambda z)$ for some suitable $\lambda \in \mathbb{R}^*$. The peculiarity of the structures obtained in this way is that their developing map is no longer injective, but is surjective onto $\mathbb{CP}^1$. Contrary to what one might expect, these turned out to be quite far from being an isolated curiosity: in [16] Goldman proved that all complex projective structures with quasi-Fuchsian holonomy are actually obtained by this construction.

This thesis is concerned with the study of the moduli space of a class of geometric structures which generalise this classical setting by introducing ramification phenomena into the picture; these are known as branched complex projective structures (BPS in the following) and were at first considered by Mandelbaum in [25] in relation with the study of projective bundles on Riemann surfaces. More precisely we allow for points around which the geometry is not necessarily modelled on $\mathbb{CP}^1$ by a local diffeomorphism, but possibly by a finite branched cover; in other words a local chart can now be a map of the form $z \mapsto z^{k+1}$. This can be interpreted as the introduction of conical singular points of angle $2\pi(k+1)$ for $k \in \mathbb{N}$. This is of course motivated by an extension of the classical uniformization problem for metrics with this type of singularities (see Hitchin [19] and Troyanov [34]); it is a more general framework in which one can study also singular metrics with constant curvature the value of which would be forbidden in the unbranched case by Gauss-Bonnet theorem: for instance flat surfaces (as the one arising from rational billiards) are examples of BPSs.

Another motivation comes from the theory of ODEs on Riemann surfaces; namely let X be a Riemann surface of genus $g \geq 2$ and let us consider the following equation on the holomorphically trivial vector bundle $X \times \mathbb{C}^2$

$$du = Au$$

where $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$ is an $\mathfrak{sl}_2\mathbb{C}$-valued holomorphic 1-form (i.e. an $\mathrm{SL}_2\mathbb{C}$-Higgs field on $X \times \mathbb{C}^2$). The datum of A gives us a monodromy representation $\rho_A : \pi_1(S) \to \mathrm{SL}_2\mathbb{C}$ which encodes the behaviour of the analytic continuation of local holomorphic solutions of the equation. The Riemann-Hilbert problem asks for a characterisation of the representations which occur in this way, and is so far quite open; for instance it is not known if real or discrete representations can be obtain in this way. The study of BPSs has turned out to provide a geometric approach to this analytic problem. First of all Gallo-Kapovich-Marden proved in [13] that a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ can be realised as the holonomy of some BPS with at most one branch point if and only if it is non-elementary (i.e. its image has no finite orbits on $\mathbb{CP}^1$), and one should remark that this condition is generic in the space representations. More recently Calsamiglia-Deroin-Hue-Loray proved that the fact that a representation ρ is realised as the monodromy of an ODE as above is detected by the geometry of a certain slice of the moduli space of BPSs which is defined in terms of ρ itself (see [10]).

A solution of this analytic problem would in turn be interesting from a geometric point of view, since it is linked to the problem of existence of holomorphic curves inside the homogeneous spaces of the form $\mathrm{PSL}_2\mathbb{C}/\Gamma$, where $\Gamma \subset \mathrm{PSL}_2\mathbb{C}$ is a discrete torsion-free subgroup (e.g. one of those used to solve the uniformization problems discussed above). These manifolds arise as the frame bundles for the hyperbolic 3-manifolds of the form $\mathbb{H}^3/\Gamma$, and the non triviality of the problem stems from the fact that they are 3-dimensional complex manifolds which are non quasi-projective in general; for instance Huckleberry and Margulis proved that in general they do not admit complex hypersurfaces (see [20]).

In this work we focus on the geometric study of the branching behaviour of a BPS and on the geometry of certain slices of the moduli space of BPSs obtained by fixing some geometric and combinatorial parameters (e.g. the number of the branch points, the underlying complex structure, or the holonomy representation). This is carried out through the analysis of the interactions between three geometric surgeries which preserve the holonomy of a BPS, whose definition and individual study is the main theme of Chapter 1. One of them is the grafting surgery (1.4.1) which has already been introduced above; the inverse surgery is known as degrafting. The second one is called bubbling (1.4.3) and is a surgery which allows to introduce branching behaviour on a given structure; more precisely it consists in replacing an embedded arc β on S with a disk of the form $\mathbb{CP}^1\setminus\widehat{\beta}$ for some embedded arc $\widehat{\beta}\subset\mathbb{CP}^1$; topologically it is just a connected sum with a sphere. Once it is performed, two new branch points appear at the ends of β, and a whole copy of $\mathbb{CP}^1$ is glued to the surface, so that the developing map becomes suddenly highly surjective and highly non injective. The inverse surgery is known as debubbling. The last deformation is called movement of branch points (1.4.4); it is a continuous deformation which is defined in a neighbourhood of a branch point and consists in a local deformation of the branched local chart. As said above, graftings have been introduced by Maskit in the context of Kleinian groups in [27]. As far as movements are concerned, they were specifically introduced by Tan in [32] for BPSs, but actually analogous local deformations known as Schiffer variations are a classical tool in the study of Riemann surfaces and their moduli spaces (see [30]). On the other hand the bubbling surgery has only been introduced by Gallo-Kapovich-Marden in [13] as a tool to produce ramification; in that work they ask the following question as Problem 12.1.2:

Problem: *Given two BPSs on a surface S with the same holonomy, can one pass from one to the other via a sequence of grafting, bubbling, degrafting and debubbling?*

For unbranched structures the answer to this question is well-known to be positive even without using bubbling and debubbling in the case of quasi-Fuchsian holonomy thanks to Goldman's work in [16]: more precisely he proved that any unbranched complex projective structure with (quasi-)Fuchsian holonomy $\rho:\pi_1(S)\to \mathrm{PSL}_2\mathbb{C}$ is obtained by grafting a finite disjoint collection of simple closed geodesics on the hyperbolic surface $\mathbb{H}^2/\rho(\pi_1(S))$. More recently Baba has extended Goldman's result for the generic case of purely loxodromic holonomy (see [2]), and Calsamiglia-Deroin-Francaviglia have improved Goldman's theorem by showing in [9] that in quasi-Fuchsian holonomy degrafting is not needed and indeed a sequence of two (multi-)graftings is always enough. Our main result on unbranched structures is that on the other hand also grafting is not needed if one uses bubbling and debubbling, and indeed one also gets better bounds on the number of necessary surgeries; more precisely we obtain the following Multi(de)grafting Lemma, which allows us to replace a long sequence of multiple graftings and degraftings by a simple sequence of one bubbling and one debubbling.

Theorem (see 3.6.8)**.** *Let S be a closed, connected and oriented surface of genus $g \geq 2$, let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation and σ an unbranched complex projective structure with holonomy ρ. Then there exist an arc $\beta \subset \sigma$ and an arc $\beta_\rho \subset \sigma_\rho = \mathbb{H}^2/\rho(\pi_1(S))$ such that $Bub(\sigma, \beta) = Bub(\sigma_\rho, \beta_\rho)$.*

As far as the branched case in concerned, the only result available in the literature is the work [8] of Calsamiglia-Deroin-Francaviglia, where they show that the answer to the above question is positive in quasi-Fuchsian holonomy if we add "movements of branch points" to the list of allowed surgeries: what they actually prove is that given a BPS with quasi-Fuchsian holonomy it is always possible to move branch points in a suitable (and quite drastic) way on it so that a bubble appears. The main result of this thesis is the following.

Theorem (see 3.5.6)**.** *Let S be a closed, connected and oriented surface of genus $g \geq 2$ and let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation. Then the space of structures obtained by a bubbling on some unbranched structure with holonomy ρ is a connected, open and dense subspace of full measure in the moduli space $\mathcal{M}_{2,\rho}$ of structures with holonomy ρ and 2 branch points.*

In other words we generically remove the need of moving branch points. As a joint application of the two results (3.6.8 and 3.5.6) we obtain that for a generic couple of BPSs with at most 2 branch points and the same quasi-Fuchsian holonomy it is possible to pass from one to the other in six steps, via a finite sequence of at most three bubblings and three debubblings (see Corollary 3.6.9). This provides an explicit generically positive answer in our setting to the question by Gallo-Kapovich-Marden.

The above theorem relies on a fundamental property of quasi-Fuchsian structures, that is the fact that a quasi-Fuchsian representation preserves a decomposition of $\mathbb{CP}^1$ into a Jordan curve and a couple of disks: for example just consider the $\mathrm{PSL}_2\mathbb{R}$-invariant decomposition $\mathbb{CP}^1 = \mathcal{H}^+ \cup \mathbb{RP}^1 \cup \mathcal{H}^-$ where $\mathcal{H}^\pm$ denote the upper/lower-half plane in $\mathbb{C}$. This decomposition induces a geometric decomposition (2.1) of the surface itself with rich geometric and combinatorial features, which have already been investigated by Goldman in [16] for the unbranched case and by Calsamiglia-Deroin-Francaviglia in [8] in the branched case. The proof of the above theorem consists of two main steps, quite different in nature. The first is a "static" one and is carried out in Chapter 2: it involves an analysis of the properties of this geometric decomposition for a given structure, and culminates in a complete classification of the elementary pieces which can appear in it (see Theorem 2.4.5); in particular it follows that either the two branch points live in the same piece or they live in adjacent pieces. Packing together structures whose geometric decomposition looks the same (in some precise sense), we also obtain a decomposition (2.1.21) of the moduli space $\mathcal{M}_{2,\rho}$; the second step, which is the "dynamic" one, occupies Chapter 3 and consists in understanding what happens when moving branch points on the surface with respect to its geometric decomposition, i.e. what happens when moving a structure around in the moduli space with respect to its decomposition. Roughly speaking, the key observation is that deformations which happen inside a fixed piece of $\mathcal{M}_{2,\rho}$ can be performed in a quantitatively controlled way, so that they preserve the property of "being obtainable by a bubbling on an unbranched structure" (see Theorem 3.5.1). On the other hand the analysis of Chapter 2 can be combined with results from [8] to show that this property is ubiquitous in the moduli space (see Proposition 3.5.5).

The last Chapter 4 is concerned with understanding how a movement of branch points can be related to the problems coming from analysis on Riemann surfaces introduced above. As anticipated, the fact that a representation ρ occurs as the monodromy of some ODE is detected by the geometry of the moduli space $\mathcal{M}_{2g-2,\rho}$ of BPSs with $2g-2$ branch points and holonomy ρ: more precisely Calsamiglia-Deroin-Hue-Loray in [10] have obtained a correspondence between ODEs on a Riemann surface X of genus g with monodromy ρ and embedded holomorphic spheres in $\mathcal{M}_{2g-2,\rho}$, and Calsamiglia-Deroin-Francaviglia asked in [8] for which representations $\pi_2(\mathcal{M}_{2g-2,\rho}) = 0$. In particular, given an ODE, there is a 1-dimensional family of BPSs in this moduli space which enjoy some special features, namely they all have the same underlying complex structure X (tautologically, by construction) and the collection of branch points is a canonical divisor on X (see 4.1.6). More generally we can prove that having a non canonical branching divisor is an obstruction to the existence of deformations preserving both the holonomy and the underlying complex structure; the main result of this chapter is the following.

Theorem (see 4.3.9)**.** *Let S be a closed, connected and oriented surface of genus $g \geq 2$, $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a non elementary representation, $k \leq 2g-2$ and $\sigma \in \mathcal{M}_{k,\rho}$. If the collection of branch points of σ is not a canonical divisor on the underlying Riemann surface X, then any movement of branch points on σ induces a non trivial deformation of X.*

This is the used to prove in 4.3.12 that for $k < 2g-2$ the holonomy fibre $\mathcal{M}_{k,\rho}$ does not contain compact complex submanifolds: this shows a significant relation between the geometry of the branching divisor and the complex geometry of the moduli space. Motivated by the search for partial converses to this theorem in genus $g = 2$, we have included in 4.4 a study of hyperelliptic BPSs (branched complex projective structures on a hyperelliptic Riemann surface for which the hyperelliptic involution is a projective automorphism), through the explicit construction of several geometric examples and some explicit computations with the classical parametrization via meromorphic quadratic differentials (see [25]).

Chapter 1

Branched complex projective structures

In this section we define the structures we are interested in, as well as some classes of deformations thereof; we fix some notations and terminology and prove some basic facts about them.

1.1 First definitions and examples

The focus of this thesis will be on a class of geometric structures on closed, connected and oriented surfaces of genus $g \geq 2$, but we begin with the most general setting, so let S be just an oriented surface, i.e. a 2-dimensional real topological manifold. These structures are locally modelled on the geometry of the Riemann sphere $\mathbb{CP}^1$ and its group of automorphisms $\mathrm{PSL}_2\mathbb{C}$. Even if their study is a traditional topic in surface theory, dating back to the works on the uniformization problem for Riemann surfaces, the study of branched structures was first introduced by Mandelbaum in the 1970s in the series of papers [25], [24] and [26].

Definition 1.1.1. A branched complex projective chart on S is an open subset $U \subset S$ endowed with a finite degree orientation preserving branched covering map $\varphi : U \to \varphi(U) \subseteq \mathbb{CP}^1$ onto an open subset of $\mathbb{CP}^1$, which we consider as a local projective coordinate on U. A point p at which φ is not a local homeomorphism is called a branch point (or critical point) of φ, and $\varphi(p)$ is called a branch value (or critical value). A chart (U, φ) is said to be unbranched if it has no branch points, i.e. if it is a homeomorphism. Two charts (U, φ) and (V, ψ) are compatible if $\exists\, g \in \mathrm{PSL}_2\mathbb{C}$ such that $\psi = g\varphi$ on $U \cap V$.

Definition 1.1.2. A branched complex projective atlas on S is an open cover by branched complex projective charts. A **branched complex projective structure** on S (often abbreviated BPS in the following) is the datum of a maximal branched complex projective atlas.

Notice that a local chart (U, φ) can always be shrunk to ensure that U contains only one branch point p of φ and both U and $\varphi(U)$ are homeomorphic to disks. In particular branch points are isolated, hence in finite number as soon as S is compact. Moreover such a branched cover induces a genuine finite covering map $\varphi : U\setminus\{p\} \to \varphi(U\setminus\{p\})$ between punctured disks, hence it is topologically equivalent to the covering $\mathbb{D}^* \to \mathbb{D}^*, z \mapsto z^{k+1}$ of the punctured unit disk in $\mathbb{C}$, for some integer $k \in \mathbb{N}$, which we call the order $ord(p)$ (or multiplicity) of the branch point p with respect to the chart (U, φ). Since elements of $\mathrm{PSL}_2\mathbb{C}$ are invertible, compatible local charts assign the same order to the same branch point, so that the following definition is well-posed.

Definition 1.1.3. If S is endowed with a BPS, a point $p \in S$ is said to be a branch point of order $k \in \mathbb{N}$ if there is a chart which is a branched cover of degree $k+1$ branching exactly a p. If moreover S is compact, we define its branching divisor to be $div(\sigma) = \sum_{p\in S} ord(p)p$ and we define the order of the structure to be the degree of its branching divisor. We can also specify precise patterns of branching: for a partition $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathbb{N}^n$ we say that a branched projective structure has order λ if $div(\sigma) = \sum_{i=1}^n \lambda_i p_i$. The structure is said to be unbranched if every chart is unbranched, i.e. its branching divisor is empty.

Remark 1.1.4. A BPS on S induces a complex structure on it: namely, for a given branched covering space of a complex disk $U \to D$ there is a unique complex structure on U such that the projection is holomorphic. Roughly speaking, in these complex coordinates around a branch point p the local projective chart looks like $z \mapsto z^{k+1}$ for $k = ord(p)$ (see 4.3.1 below for more details). From this point of view, U can be thought of as being a domain inside the Riemann surface of $\sqrt[k+1]{z}$. Moreover this local complex structure is preserved by the change of coordinates of a BPS, simply because $\mathrm{PSL}_2\mathbb{C}$ acts holomorphically on $\mathbb{CP}^1$; therefore we have a well-defined complex structure on S induced by the BPS. This complex structure will be referred to as the underlying complex structure of the BPS. The total angle (with respect to the associated conformal structure) around a branch point of order k is $2\pi(k+1)$.

Remark 1.1.5. A BPS on S can be considered as a generalised $(\mathrm{PSL}_2\mathbb{C}, \mathbb{CP}^1)$-structure, for which the developing map may have critical points, corresponding to branch points. A developing map for such a structure is an orientation preserving smooth map $dev : \widetilde{S} \to \mathbb{CP}^1$ with isolated critical points and equivariant with respect to a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$, which we call the holonomy of the structure. Given such a dev, an atlas is obtained by precomposing dev with local inverses of the universal covering projection. Conversely analytic continuation of local charts and change of coordinates of a given projective atlas give rise to a couple (dev, ρ), which is a standard construction in the theory of geometric structures (see for instance [33, Chapter 3]). Notice that in general these geometric structures are not complete structures (i.e. they are not locally homogeneous symmetric spaces): indeed $\widetilde{S}$ is not even homeomorphic to $\mathbb{CP}^1$ as soon as S is not a sphere.

Understanding which structures have the same underlying complex structure or the same holonomy representation is a problem of major interest, to which we will come back later in 1.3. Let us end this preliminary section with some basic examples, to show how ubiquitous these structures are in the study of surfaces.

Example 1.1.6. Every Riemann surface X admits a non-constant meromorphic function f, which realizes it as a branched cover of $\mathbb{CP}^1$ of finite degree. This endows X with a BPS with trivial holonomy and developing map given by f itself. The pattern of branch points depends on f and has to obey the Riemann-Hurwitz formula

$$\chi(X) = 2deg(f) - \sum_{p \in X} (ord_p(f) - 1)$$

Example 1.1.7. Every surface admits a complete Riemannian metric g of constant curvature $k_g = -1, 0$ or 1, which realises it as a quotient of $\mathbb{S}^2, \mathbb{E}^2$ or $\mathbb{H}^2$ by a group of isometries acting freely and properly discontinuously: this follows directly from Poincaré-Koebe Uniformization theorem (see [12, Chapter IV]). Each one of these three geometries has a conformal embedding in the 1-dimensional complex projective geometry $(\mathrm{PSL}_2\mathbb{C}, \mathbb{CP}^1)$, hence provides an example of an unbranched projective structure, with holonomy landing respectively in the subgroups $\mathrm{PSU}(2)$, $\mathrm{PSO}(2) \ltimes \mathbb{C}$ and $\mathrm{PSL}_2\mathbb{R}$ of $\mathrm{PSL}_2\mathbb{C}$. The value of the curvature depends on the topology of the surface according to Gauss-Bonnet theorem

$$2\pi\chi(S) = \int_S k_g \; dvol_g$$

However if we allow metrics with cone singularities of angle $2\pi(k+1)$ for some $k \in \mathbb{N}$, then we obtain many more examples of BPS. Consider for instance the flat cone structure obtained on a surface of genus g by gluing the sides of a regular Euclidean $4g$-gon. More generally, it follows from the work of Troyanov in [34] that, given $p_1, \dots, p_n \in S$ and $k_1, \dots, k_n \in \mathbb{N}$, there exists a Riemannian metric with cones of angle $2\pi(k_i + 1)$ at p_i and smooth elsewhere, with curvature $-1, 0$ or 1 if the quantity $\chi(S) + \sum_{i=1}^n k_i$ is $< 0, = 0$ or 1 respectively.

Example 1.1.8. If Γ is a finitely generated torsion free Kleinian group (i.e. discrete subgroup of $\mathrm{PSL}_2\mathbb{C}$) with domain of discontinuity $\Omega_\Gamma \subset \mathbb{CP}^1$, then by Ahlfors finiteness theorem (see [22, Theorem 4.108]) $\Omega_\Gamma \to \Omega_\Gamma/\Gamma$ is a (possibly disconnected) covering over a (possibly disconnected) Riemann surface of finite type, which is clearly endowed with a uniform unbranched projective structure with holonomy group Γ. More generally it is shown in [13] that if S is connected, closed and oriented of genus $g \geq 2$ then a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ is the holonomy of a BPS on S with at most one branch point of order 1 if and only if it is non elementary (i.e. its image has no finite orbits on $\mathbb{CP}^1$). Of course also many elementary representations (e.g. many affine ones) arise as holonomy of some BPS, but more branch points are needed.

Example 1.1.9. Let X be a Riemann surface and $\pi : P \to X$ a holomorphic $\mathbb{CP}^1$-bundle with structure group $\mathrm{PSL}_2\mathbb{C}$. Consider a smooth holomorphic codimension 1 foliation $\mathcal{F}$ of P whose leaves are transverse to the fibres of π. Then any holomorphic section $s : X \to P$ which is generically transverse to the foliation can be used to define charts for a BPS on X with holonomy the monodromy of $\mathcal{F}$: namely fix a fibre and compose the section with parallel transport along the leaves; branch points occur at points where $s(X)$ is tangent to $\mathcal{F}$, if any. The basic example is obtained by taking P to be the trivial $\mathbb{CP}^1$-bundle with the natural foliation $\{z = z_0\}_{z_0 \in \mathbb{CP}^1}$; holomorphic sections are just meromorphic functions, and we are back to the first example. More interesting examples can be obtained by choosing an ODE on X, which will be discussed in 4.1.

1.2 Projective maps and orbibranched complex projective structures

In this section we want to define the "right" notion of map between BPSs in order to obtain a reasonable category to work with. The most straightforward definition would be to say that a continuous map $f : (S, \sigma) \to (T, \tau)$ between surfaces endowed with BPSs is a projective map if it is given as the restriction of an element of $\mathrm{PSL}_2\mathbb{C}$ in any couple of local charts. Unfortunately this turns out to be not very useful: since Möbius transformations are biholomorphic, such a map would be holomorphic and without critical points, in particular it would be a genuine covering map, which imposes rigid conditions on the topological type of S and T, regardless of their geometric structures.

Since a biholomorphism of $\mathbb{CP}^1$ is automatically projective, the real difference between a complex atlas and a projective one is that in the first case we allow change of charts to be local biholomorphisms, whereas in the second one we want them to be restrictions of global biholomorphisms: the stress is not on regularity (holomorphic vs algebraic), but on the domain of definition (local vs global). We take therefore the following approach.

Definition 1.2.1. Let $\mathcal{M}(\mathbb{CP}^1) = \{F : \mathbb{CP}^1 \to \mathbb{CP}^1 \mid F \text{ is holomorphic}\}$ be the algebra of holomorphic self-maps of $\mathbb{CP}^1$, also known as the algebra of meromorphic functions on $\mathbb{CP}^1$. The subset of elements invertible with respect to composition is the group $\mathrm{PSL}_2\mathbb{C}$ of automorphisms of $\mathbb{CP}^1$.

Definition 1.2.2. Let (S, σ) and (T, τ) be surfaces endowed with a BPS. We define the set of projective maps between them as

$$Proj(\sigma, \tau) = \{f : S \to T \mid f \text{ is locally the restriction of some } F \in \mathcal{M}(\mathbb{CP}^1)\}$$

We say that $f \in Proj(\sigma, \tau)$ is a projective isomorphism (or projective diffeomorphism) if it is a diffeomorphism.

The following is immediate from the definitions.

Lemma 1.2.3. *Projective maps are holomorphic with respect to the underlying complex structures.*

Denoting by $\mathbb{C}(z)$ the algebra of rational functions in one variable over $\mathbb{C}$, we have this classical fact.

Lemma 1.2.4. *Let $F : \mathbb{CP}^1 \to \mathbb{CP}^1$; then the following are equivalent.*

1. $F \in \mathcal{M}(\mathbb{CP}^1)$

2. $F \in \mathbb{C}(z)$

3. $\exists\, g_1, \dots, g_n \in \mathrm{PSL}_2\mathbb{C} : F = g_1 \dots g_n$

We can obtain the following characterisation, which brings us back to the initial intuition of being projective in the sense of being locally in $\mathrm{PSL}_2\mathbb{C}$.

Lemma 1.2.5. *Let S be a closed surface of genus $g \geq 2$, σ be a BPS on S and $f : S \to S$ be bijective. Then $f \in Proj(\sigma, \sigma)$ if and only if f is locally the restriction of some $g \in \mathrm{PSL}_2\mathbb{C}$.*

Proof. One inclusion is clear since $\mathrm{PSL}_2\mathbb{C} \subset \mathcal{M}(\mathbb{CP}^1)$. For the other one let f be bijective and locally the restriction of some $F \in \mathcal{M}(\mathbb{CP}^1)$. We want to show that actually $F \in \mathrm{PSL}_2\mathbb{C}$. To do this we prove that also f^{-1} is locally represented by a function in $\mathbb{CP}^\infty$, which implies that both local representations are not just holomorphic self-maps of $\mathbb{CP}^1$, but they are indeed biholomorphisms. But notice also that by 1.2.3 f is a biholomorphism of the underlying complex structure; by a classical result of Hurwitz f has finite order. Therefore f^{-1} is a power of f. It is now straightforward to check that powers of f are again locally represented by functions in $\mathcal{M}(\mathbb{CP}^1)$. Indeed let U be a local chart and choose $n \in \mathbb{N}$ and local charts on $f^k(U)$ for $k = 1, \dots, n$. Then we can locally read $f : f^{k-1}(U) \to f^k(U)$ as a map $F_k \in \mathcal{M}(\mathbb{CP}^1)$. Then the composition $F_n \circ \dots \circ F_1 \in \mathcal{M}(\mathbb{CP}^1)$ is exactly how we read $f^n : U \to f^n(U)$. □

This implies in particular that the set of projective bijections of a BPS forms a group under composition.

Definition 1.2.6. If σ is a BPS on S, we denote by $Proj(\sigma)$ the group of projective diffeomorphisms of σ.

This is naturally a subgroup of the group of biholomorphisms of the underlying complex structure by 1.2.3. With respect to the description of BPS as a branched $(\mathrm{PSL}_2\mathbb{C}, \mathbb{CP}^1)$-structure, we have the following characterisation of projective automorphisms.

Lemma 1.2.7. *Let $\sigma = (dev, \rho)$ be a BPS on S and $f \in Diff(S)$. Then $f \in Proj(\sigma)$ if and only if for any lift $\widetilde{f} : \widetilde{S} \to \mathbb{CP}^1$ there exists $g \in \mathrm{PSL}_2\mathbb{C}$ such that*

$$dev \circ \widetilde{f} = g \circ dev$$

Proof. It is enough to observe that local inverses post composing local inverses of the universal cover projection $\widetilde{S} \to S$ with dev provides an atlas of branched complex projective charts on S for σ with respect to which the statement follows by a direct computation. □

According to 1.2.3 projective maps are holomorphic with respect to the underlying complex structure, hence from a topological point of view they are branched covers. It is therefore natural to try to transfer a BPS from a surface to another via a branched cover.

Lemma 1.2.8. *Let $f : T \to S$ be a branched cover and σ a BPS on S. Then there is a natural BPS on τ on T such that $f \in Proj(\tau, \sigma)$. Moreover if f has multiplicity l at some point $p \in T$ and $f(p) \in S$ has order k with respect to the structure σ, then p will have order $l(k+1)$ with respect to τ.*

Proof. We just define an atlas on T by precomposing charts of σ with f. Then f is locally given by the identity of $\mathbb{CP}^1$. The statement about orders is easily checked in local coordinates, where f looks like $z \mapsto z^l$ and a local chart for σ looks like $w \mapsto w^{k+1}$. □

This generalises 1.1.6 to an arbitrary holomorphic map between Riemann surfaces. Notice that the structure on T will be branched as soon as f is, regardless of the fact that σ is branched or not. Of course the reverse direction, i.e. pushing a BPS downstairs with respect to a branched cover is less straightforward. We are interested in branched covering coming from a group action and, as usual, this forces us to consider orbifold points. In the context of more common geometric structures (e.g. complex structures) one can usually forget about this because orbifold charts can be reuniformized to obtain genuine charts; for Riemann surfaces this is usually done by analytic methods as in [12, p. III.7.7], or by algebraic methods as in [29, p. III.3.4]. The former relies on transcendental mappings which are far from being projective (i.e. global), but even the latter is not sufficient in our setting, since the change of coordinates between charts around an orbifold point for the candidate atlas on the quotient surface involves a non trivial product of local expressions of the elements of the local orbifold group of that point, which is locally in $\mathcal{M}(\mathbb{CP}^1)$ (and with non zero derivative on the involved charts) but in general not in $\mathrm{PSL}_2\mathbb{C}$. The best we can get on the quotient surface is a BPS with some special points with orbifold behaviour, which we will shortly define more carefully. We are not going to enlarge the category of structures we consider by including these ones; rather, we use them when needed in order to prove results about BPSs, for instance in the study of hyperelliptic structures in 4.4.

Definition 1.2.9. An **orbifold branched complex projective structure** on S (often abbreviated OBPS in the following) is the datum of a maximal open cover $\{U_i\}_{i\in I}$ of S with finite degree branched covering maps $\pi_i : \widehat{U}_i \to U_i$ on which finite degree branched covering maps $\varphi_i : \widehat{U}_i \to V_i \subseteq \mathbb{CP}^1$ are defined, with the requirement that the needed changes of coordinates are in $\mathrm{PSL}_2\mathbb{C}$.

Of course we can always restrict charts so that in each $\widehat{U}_i$ there is at most one branch point for both π_i and φ_i, and the maps are regular elsewhere. Notice that in particular an OBPS induces an orbifold complex structure on S. As for BPSs, the local degrees of the map involved provide a well defined notion of order of orbibranching at any point.

Definition 1.2.10. If at a point $\widehat{p} \in \widehat{U}_i$ we have $ord_p(\pi_i) = n_i$ and $ord_p(\varphi_i) = m_i+1$, then we call the ratio $\frac{m_i}{n_i}$ the order of the structure at the point $p = \pi_i(\widehat{p})$.

If BPSs should be thought of as a generalisation of complex projective structures where we allow cone points of angle $2\pi(k+1)$ for $k \in \mathbb{N}$, then OBPSs should be thought as a generalisation where we allow cone of angle $2\pi q$ for $q \in \mathbb{Q}$. Namely if a point has order $\frac{m_i}{n_i}$, then the angle at it is $2\pi\frac{m_i+1}{n_i}$. Notice that when $n_i = 1$ (i.e. p is has not an orbifold behaviour), this definition agrees with the one given for BPSs in 1.1.3.

Remark 1.2.11. By the classification of 2-dimensional orbifolds (see for instance [33, Chapter 13]), we know that most of them are good, i.e. are covered by some manifold, the only exceptions being a sphere with at most two cone points and a disk with at most two corner reflectors. We will not be interested in taking covers of such orbifolds, hence we can always view an OBPS on S as a development-holonomy pair (D, ρ) where D is a developing map for a BPS on some manifold cover of S and ρ is a representation of the corresponding automorphism group. For instance we can take the cover to be the orbifold universal cover $\widetilde{S}^{orb}$ and ρ to be a representation of the orbifold fundamental group $\pi_1^{orb}(S)$.

Example 1.2.12. Any constant curvature metric on a surface with cone points with angle of the form $2\pi q$ for $q \in \mathbb{Q}$ clearly endows it with an OBPS, generalising 1.1.7. All good orbifolds are known to admit such a metric, the sign of the curvature depending of course on the orbifold Euler characteristic. These structures arise for instance as quotients of surfaces with smooth constant curvature metrics by finite groups of isometries, and the results in [34] provide many more examples.

As announced above we can prove the following fact.

Lemma 1.2.13. *If σ is a BPS on a closed surface S of genus $g \geq 2$ and $G \subset Proj(\sigma)$, then the quotient surface S/G carries a natural OBPS σ/J, with respect to which the quotient map is projective.*

Proof. The maps $\pi_i : \widehat{U}_i \to U_i$ in the definition are restrictions of the quotient map, and the maps $\varphi_i : \widehat{U}_i \to \mathbb{CP}^1$ are exactly the projective charts of σ. Notice that an automorphism of a BPS is in particular a biholomorphism for the underlying complex structure, therefore such a group is always finite, by a classical result of Hurwitz. Therefore the maps π_i are finite branched covers, as required. We just need to check that the transitions lie in $\mathrm{PSL}_2\mathbb{C}$, but this is guaranteed by the fact that G acts by projective automorphisms. □

Conversely we also have the following result which allows to lift an OBPS to a BPS by suitably branching over the points with orbifold behaviour. Recall that we are not interested in bad orbifolds, so that we always implicitly assume that either S has positive genus, or that the induced orbifold is not a sphere with one or two cone points.

Lemma 1.2.14. *Let $f : T \to S$ be a branched cover between surfaces and let σ be an OBPS on S. Suppose that*

1. *if a point of S has non integer order, then it is a branch value for f,*
2. *if p is a branch point of f of order k and $f(p)$ has non integer order $\frac{m}{n}$, then $n = k$.*

Then T is naturally endowed with a BPS (no orbifold behaviour involved) with respect to which f is projective.

Proof. If p is not a branch point of f, then f is a local diffeomorphism around it, hence we just lift the structure from S. If $f(p)$ has integer order, then a projective chart of σ is already defined around it, hence we just compose it with f as in 1.2.8. The case we are left with is the one in which p is a branch point of f of order k and $f(p)$ has non integer order $\frac{m}{n}$. Let $p \in V_i$ and let $f(p) \in U_i, \pi_i : \widehat{U}_i \overset{m:1}{\to} U_i, \varphi_i : \widehat{U}_i \overset{n:1}{\to} \mathbb{CP}^1$ be as in the definition of OBPS. Since by hypothesis $n = k$, we can lift $f|_{f^{-1}(U_i)}$ with respect to π_i and obtain a diffeomorphism $f_i : V_i \to \widehat{U}_i$. Composing it with φ_i gives a projective coordinate on V_i; notice that p has integer order in the resulting structure. This defines possibly branched local coordinates to $\mathbb{CP}^1$ on the whole T and transition maps are in $\mathrm{PSL}_2\mathbb{C}$ since they come from those of σ. □

Remark 1.2.15. In particular, as soon as the covering is normal, the structure on S is the quotient of the structure on T from the previous lemma by a group of projective automorphisms in the sense of 1.2.13. Indeed S is a quotient of its orbifold universal cover $\widetilde{S}^{orb}$ by the action of its orbifold fundamental group $\pi_1^{orb}(S)$, with $\widetilde{S}^{orb} \to T$ as intermediate covering with a torsion-free group of deck transformations. In other words there is a group G of projective automorphisms of T which $\pi_1^{orb}(S)$ extends by $\pi_1(T)$, i.e.

$$1 \to \pi_1(T) \to \pi_1^{orb}(S) \to G \to 1$$

In particular the holonomy representation for the structure on T is just the restriction of the holonomy representation for the OBPS on S.

1.3 Deformation spaces

We want to study the moduli space of marked branched complex projective structures on a give surface. So we let S be a closed, connected and oriented surface and give the following usual definitions.

Definition 1.3.1. A marked branched complex projective structure on S is a couple (σ, f) where σ is a surface endowed with a BPS and $f : S \to \sigma$ is an orientation preserving diffeomorphism, called the marking. Two marked BPS (σ, f) and (τ, g) are declared to be equivalent if $gf^{-1} : \sigma \to \tau$ is isotopic to a projective isomorphism $h \in Proj(\sigma, \tau)$. We denote by $\mathcal{BP}(S)$ the set of marked branched complex projective structure on S up to this equivalence relation. The subset of unbranched structures is denoted by $\mathcal{P}(S)$.

If (σ, f) is a marked BPS, then by the results in 1.2 we can pullback the BPS via f to obtain a BPS $f^*\sigma$ on S. From this point of view two marked BPS (σ, f) and (τ, g) are equivalent if there exists a diffeomorphism of S which is isotopic to the identity and projective as a map $f : f^*\sigma \to g^*\tau$. It is also straightforward to check that associating to a BPS the underlying complex structure defined in 1.1.4 provides a well defined map $\pi : \mathcal{BP}(S) \to \mathcal{T}(S)$ to the Teichmüller space of S, i.e. the space of marked complex structures on S up to the analogous equivalence relation.

To have yet another perspective on this space, we recall that by 1.1.5 a BPS on S can be seen as a $(\mathrm{PSL}_2\mathbb{C}, \mathbb{CP}^1)$-structure with branch points, therefore a maximal atlas is the same as an equivalence class of development-holonomy pairs (dev, ρ), where $dev : \widetilde{S} \to \mathbb{CP}^1$ is an orientation preserving smooth map with isolated critical points which is equivariant with respect to a representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ and two such pairs (dev_1, ρ_1) and (dev_1, ρ_1) are equivalent if $\exists\ g \in \mathrm{PSL}_2\mathbb{C}$ such that $dev_2 = g dev_1$ and $\rho_2 = g\rho_1 g^{-1}$. The group $\mathrm{Diff}_0(S)$ of diffeomorphisms of S isotopic to the identity acts on the set of such equivalence classes of pairs by precomposition on the first factor (and the trivial action on the second). The quotient is still the set $\mathcal{BP}(S)$ defined above. From this description we obtain an easy way to put a topology on $\mathcal{BP}(S)$, namely we put the compact-open topology on the set of development-holonomy pairs and then take the quotient topologies with respect to the action of $\mathrm{PSL}_2\mathbb{C}$ and $\mathrm{Diff}_0(S)$. Moreover, by composing the holonomy of a marked BPS with the map induced by the marking on fundamental groups, we also get a well defined holonomy map $hol : \mathcal{BP}(S) \to \chi(S)$ to the $\mathrm{PSL}_2\mathbb{C}$-character variety of $\pi_1(S)$, which is by definition the GIT quotient $\chi(S) = Hom(\pi_1(S), \mathrm{PSL}_2\mathbb{C})//\mathrm{PSL}_2\mathbb{C}$ of the representation variety by conjugation.

We have so far defined two maps from our deformation space $\mathcal{BP}(S)$ to some other moduli spaces which are well known in the literature:

$$\mathcal{T}(S) \xleftarrow{\pi} \mathcal{BP}(S) \xrightarrow{hol} \chi(S)$$

The study of these maps is of major interest in the understanding of branched complex projective structures

Remark 1.3.2. Already at the level of a BPS as a maximal atlas, there is no well defined map to the representation variety $Hom(\pi_1(S), \mathrm{PSL}_2\mathbb{C})$, because the holonomy of such a structure is only defined up to conjugation, i.e. as a map to the character variety. However in the following we will need to choose a representation in its conjugacy class, i.e. to fix a representative representation of the holonomy of a structure. First notice that in general if H is a subgroup of a group G with trivial centralizer in G, then any subgroup of G conjugate to H has trivial centralizer, so that it makes sense to say that a structure has holonomy with trivial centralizer. If σ is such a structure, then for any choice of representative representation ρ there is a unique developing map for σ which is equivariant with respect to that representation (of course up to an isotopy of S). This is seen as follows: suppose that dev_1 and dev_2 are developing maps for σ equivariant with respect to ρ. Then $\exists\ g \in \mathrm{PSL}_2\mathbb{C}$ such that $dev_2 = g dev_1$ and for any $\gamma \in \pi_1(S)$ and any $x \in \widetilde{S}$ we have

$$\rho(\gamma) g dev_1(x) = \rho(\gamma) dev_2(x) = dev_2(\gamma.x) = g dev_1(\gamma.x) = g\rho(\gamma) dev_1(x)$$

so that $(\rho(\gamma)g)^{-1} g\rho(\gamma)$ is an element of $\mathrm{PSL}_2\mathbb{C}$ fixing every point of $dev_1(\widetilde{S})$. Since a developing map has isolated critical points, there is some point of $\widetilde{S}$ at which it is a local diffeomorphism, hence its image has non empty interior. But a Möbius transformation fixing more than three points is the identity of $\mathbb{CP}^1$, hence $(\rho(\gamma)g)^{-1} g\rho(\gamma) = id$. This means that g is in the centralizer of the image of ρ, which is assumed to be trivial, so that $g = id$ and the two developing maps coincide.

The study of the subspace $\mathcal{BP}(X) = \pi^{-1}(X)$ of $\mathcal{BP}(S)$ made of BPSs with a given underlying complex structure $X \in \mathcal{T}(S)$ is well established in the literature, both in the unbranched case and in the branched case: in the former an analytic parametrization via the Schwarzian derivative is available, from which it follows that the map π has the structure of a holomorphic affine bundle (see [11] for a good survey); in the branched case some singularities appear, but a parametrization of fibres via meromorphic differentials with a given branching divisor has been established in [25]; we will review this case later in 4.4.1. On the other hand, we will be mostly interested in subspaces of $\mathcal{BP}(S)$ obtained by fixing the holonomy representation.

1.3.1 Holonomy fibres and their strata

Here we focus on subspaces of $\mathcal{BP}(S)$ given by the fibres of the holonomy map $hol : \mathcal{BP}(S) \to \chi(S)$. We remark that the character variety is an algebraic variety which can exhibit quite wild behaviour, and is not smooth in general. To avoid difficulties which are unnecessary with respect to what we will need in the following, here we restrict to the setting of non elementary representations of the fundamental group of a closed, connected and oriented surface S of genus $g \geq 2$. These can be characterised as representations whose image does not have finite orbits on $\mathbb{CP}^1$, but see the Appendix for details. They are in particular irreducible, therefore the locus they define in $\chi(S)$ turns out to be a smooth subset in the Euclidean topology.

Definition 1.3.3. For a given $\rho \in \chi(S)$ we define the **holonomy fibre** to be

$$\mathcal{M}_\rho = \{\sigma \in \mathcal{BP}(S) \mid hol(\sigma) = \rho\}$$

and for $k \in \mathbb{N}$ and λ a partition of k we also define

$$\mathcal{M}_{k,\rho} = \{\sigma \in \mathcal{BP}(S) \mid ord(\sigma) = k, hol(\sigma) = \rho\}$$

$$\mathcal{M}_{\lambda,\rho} = \{\sigma \in \mathcal{BP}(S) \mid ord(\sigma) = \lambda, hol(\sigma) = \rho\}$$

where the order of a structure is the one defined in 1.1.3. We call the **principal stratum** of $\mathcal{M}_{k,\rho}$ the subspace given by the partition $\lambda = (1, \dots, 1)$, i.e. the one in which all branch points are simple ($z \mapsto z^2$ in suitable local coordinates).

When ρ is a non elementary representation the results in the appendix of [8] imply that the space $\mathcal{M}_{k,\rho}$ carries a natural structure of complex manifold of dimension k and that the subspace determined by a partition λ of length n is a complex submanifold of dimension n. In particular the principal stratum is a connected, dense and open complex submanifold of $\mathcal{M}_{k,\rho}$. For more details see 1.4.22 below and the comments after it. We find it convenient to extend the notation as follows: for any subspace $\mathcal{X}$ of $\mathcal{M}_{k,\rho}$ and λ a partition of k we define

$$\mathcal{X}_\lambda = \{\sigma \in \mathcal{X} \mid ord(\sigma) = \lambda\}$$

A useful feature of this topology is that local neighbourhoods of a structure σ inside the whole $\mathcal{M}_{k,\rho}$ or inside a specified stratum $\mathcal{M}_{\lambda,\rho}$ are easily described in terms of concrete geometric surgeries on σ which preserve the holonomy of the structure; the same holds for standard ways to jump from $\mathcal{M}_{k,\rho}$ to $\mathcal{M}_{k+2,\rho}$ and all these deformations are described in the next section.

1.4 Surgeries on a branched complex projective structure

In this section we describe how to perform some geometric operations on a given BPS to produce other structures; these deformations preserve the marking and the holonomy representation of the structure, hence they are well defined as operations on the deformation spaces $\mathcal{BP}(S)$ and $\mathcal{M}_\rho$. However they do not preserve the underlying complex structure in general, as will be proved in 4.3. Throughout this section let S be an oriented surface and $\sigma \in \mathcal{BP}(S)$.

1.4.1 Grafting

The first surgery consists in replacing a simple closed curve with a large annulus endowed with a projective structure determined by the structure we begin with; in particular it is topologically trivial. It was first introduced by Maskit in [27] to produce examples of projective structures with discrete holonomy and surjective developing map; here we review it mainly to fix terminology and notation.

Definition 1.4.1. Let $\gamma \subset S$ be a simple closed curve, and let (dev, ρ) be a development-holonomy pair defining σ. We say that γ is graftable with respect to σ if $\rho(\gamma)$ is loxodromic and γ is injectively developed, i.e. the restriction of the developing map to any of its lifts $\widetilde{\gamma} \subset \widetilde{S}$ is injective.

Since the developing map is ρ-equivariant, if γ is graftable then a developed image of it is an embedded arc in $\mathbb{CP}^1$ joining the two fixed points of $\rho(\gamma)$. Notice that $\rho(\gamma)$ acts freely and properly discontinuously on $\mathbb{CP}^1 \setminus \overline{dev(\widetilde{\gamma})}$ and the quotient is a Hopf annulus, i.e. an annulus endowed with a complete unbranched complex projective structure.

Definition 1.4.2. Let $\sigma \in \mathcal{BP}(S)$ and $\gamma \subset S$ be graftable with respect to σ. For any lift $\widetilde{\gamma}$ of γ we cut $\widetilde{S}$ along it and $\mathbb{CP}^1$ along $\overline{dev(\widetilde{\gamma})}$ and glue them together equivariantly via the developing map. This gives us a simply connected surface $\widetilde{S}'$ to which the action $\pi_1(S) \curvearrowright \widetilde{S}$ and the map $dev : \widetilde{S} \to \mathbb{CP}^1$ naturally extend, so that the quotient gives rise to a new structure $\sigma' \in \mathcal{BP}(S)$. We call this structure the **grafting** of σ along γ and denote it by $Gr(\sigma, \gamma)$.

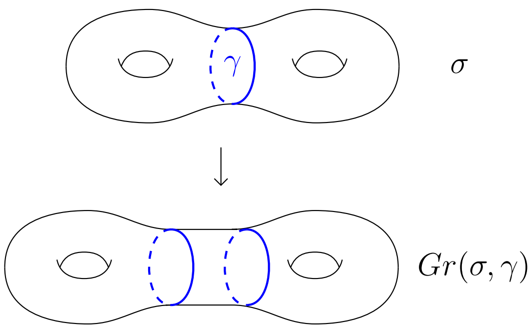

Figure 1.1: Grafting a surface

Notice that $\sigma \setminus \gamma$ projectively embeds in $Gr(\sigma, \gamma)$ and that the complement is the annulus $A_\gamma = (\mathbb{CP}^1 \setminus \overline{dev(\widetilde{\gamma})})/\rho(\gamma)$, which we call the **grafting annulus** associated to γ. This construction can of course be extended to perform simultaneous graftings on a disjoint collection of graftable curves. It is also possible to attach an integer weight $n \in \mathbb{N}$ to a graftable curve and to perform an n-fold grafting along it by gluing not just one copy of $\mathbb{CP}^1 \setminus \overline{dev(\widetilde{\gamma})}$ but n copies of it, attached in a chain of length n along their boundaries. The corresponding region in the surface is a chain $A_\gamma = \cup_{k=1}^n A_\gamma^k$ of n copies of the annulus $(\mathbb{CP}^1 \setminus \overline{dev(\widetilde{\gamma})})/\rho(\gamma)$, which we call the **grafting region** associated to $n\gamma$, and we reserve the term grafting annulus for the individual A_{γ_i}'s. This generalisation allows to perform a grafting along any graftable multicurve; we call this operation **multigrafting**.

Example 1.4.3. The main example of the grafting construction consists in grafting simple geodesics. On a genuine hyperbolic surface every simple essential curve γ is graftable, since the holonomy is purely hyperbolic and the developing map is globally injective (actually a global diffeomorphism $dev : \widetilde{S} \to \mathbb{H}^2$). Since the holonomy is also real, it preserves $\mathbb{RP}^1$ inside $\mathbb{CP}^1$, therefore the grafting annulus A_γ carries a natural decomposition in a couple of annuli coming from the quotient of the upper-half plane with the developed image of γ removed, a couple of simple closed curves l_R, l_L coming from the quotient of $\mathbb{RP}^1$ with the fixed points of $\rho(\gamma)$ removed and a negative annulus coming from the lower-half plane. The boundary of A_γ consists of a couple of closed geodesic γ_R^+, γ_L^+ coming from γ and developing to the positive part of the invariant axis of $\rho(\gamma)$.

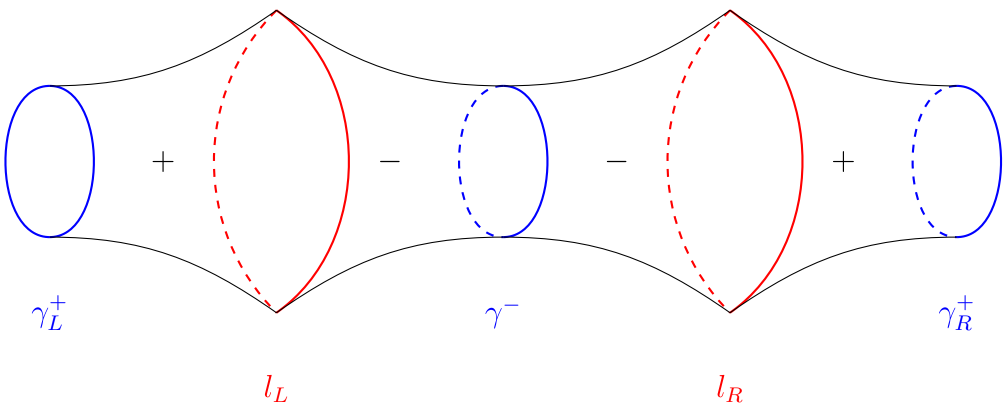

Figure 1.2: Grafting annulus from a geodesic on a hyperbolic surface

Notice that this surgery preserves the holonomy and does not introduce any new branch point, so that it induces a map

$$Gr : \mathcal{M}_{\lambda,\rho} \times \mathcal{GMC}(S) \to \mathcal{M}_{\lambda,\rho}$$

where $\mathcal{GMC}(S)$ is the set of graftable multicurves on S and λ is any partition of some natural number. Notice that for any structure σ and any graftable curve γ on it the structure $Gr(\sigma, \gamma)$ has surjective developing map onto $\mathbb{CP}^1$; in any case $Gr(\sigma, \gamma)$ is nor a complete structure, neither a uniform one, i.e. it is not isomorphic to the quotient of an open domain of $\mathbb{CP}^1$ by a discrete subgroup of $\mathrm{PSL}_2\mathbb{C}$.

Of course if a structure displays a grafting region A, it is possible to degraft it, i.e. to remove the interior of A and glue the resulting boundaries together, which results in a new structure on the same surface with a distinguished graftable curve. The main result in [16] implies that any two unbranched structures on the same surface with the same quasi-Fuchsian holonomy (see 2.1.2 below for definitions) are related by a sequence of grafting and degrafting along suitable multicurves. This was improved in [9, Theorem 1.1], where it is shown that actually a sequence of two multigraftings is always enough.

1.4.2 Cut and paste

In this section we show how to put a BPS on the connected sum of two surfaces, i.e. we give a complete proof of [8, Lemma 2.8]. In particular taking the connected sum with a sphere gives what is called a bubbling, which will be the subject of a following section. We need some preliminary results.

Lemma 1.4.4. *Let $g \in \mathrm{PSL}_2\mathbb{C}, g \neq id$ have two fixed points $z^{\pm} \in \mathbb{CP}^1$ and $U \subseteq \mathbb{CP}^1$ be a g-invariant (i.e. $gU = U$) connected simply-connected*[1] *open set with $z^{\pm} \in U$. Then $U = \mathbb{CP}^1$.*

Proof. Having two fixed points, g can be either elliptic or loxodromic.

Suppose that g is loxodromic; then $z^{\pm}$ are respectively the attracting and repelling fixed points. Suppose $U \neq \mathbb{CP}^1$ and let $p \notin U$; then $K = \{p\}$ is a compact subset of $\mathbb{CP}^1 \setminus \{z^-\}$. Since U is an open neighbourhood of z^+, by the convergence property of loxodromic transformations we can find some $n \in \mathbb{N}$ such that $g^n(p) \in U$. But g-invariance of U would imply that $p \in U$, which contradicts the choice of p.

Now suppose that g is elliptic. Then $\mathbb{CP}^1 \setminus \{z^{\pm}\}$ is foliated by simple closed curves c_i such that each c_i is the union of the g-orbits of its points and this foliation is isomorphic to the foliation of $S^1 \times [0, 1]$ by curves of the form $k_t = \{(\theta, t) \mid \theta \in S^1\}$ for $t \in [0, 1]$.

Since U is connected and open, we can join z^- to z^+ by a continuous embedded path γ; we have that $g\gamma$ is contained in U, by g-invariance of U. Moreover since $\pi_1(U) = 1$, we can find a homotopy $H : [0, 1] \times [0, 1] \to U, (s, t) \mapsto H_s(t)$ relative to $z^{\pm}$ from γ to $g\gamma$.

By construction each path H_s must intersect[2] each c_i. Fix some c_i and let p be the first intersection between γ and c_i; then gp is the first intersection between $g\gamma$ and c_i, and c_i is divided in two subarcs a_i and b_i with endpoints p and gp. We claim that at least one of these two subarcs is contained in the image of H (and thus in U).

Indeed, if we could find a point A on a_i and a point B on b_i not contained in the image of H, then H would realize a homotopy between γ and $g\gamma$ in $\mathbb{CP}^1 \setminus \{A, B\}$. But these curves are not homotopic in that space by construction: namely following γ and $g\gamma$ we obtain a path in the annulus $\mathbb{CP}^1 \setminus \{A, B\}$ which winds around at least once. Since both a_i and b_i contain a fundamental domain for the action of g

[1] Needed only in the elliptic case to avoid trivial counterexamples.

[2] We cannot in general take a γ which intersects each γ_i exactly once. Or hope that γ and $g\gamma$ intersect each other only at $z^{\pm}$.

on the leaf c_i, we can conclude that the whole leaf c_i is contained in U, since it is g-invariant. Since $\mathbb{CP}^1 \setminus \{z^\pm\}$ is foliated by such leaves and by hypothesis $z^\pm \in U$ we conclude $U = \mathbb{CP}^1$. □

Lemma 1.4.5. *Let S and T be compact surfaces with one boundary component whose interior is endowed with a BPS; let $U \subseteq S$ and $V \subseteq T$ be collar neighbourhoods of the boundaries and $f : U \to V$ be a homeomorphism which is a projective diffeomorphism of their interiors. Let X_f be the surface obtained as $X_f := {S \sqcup T}/{\sim_f}$, where $\partial S \ni x \sim f(x) \in \partial T$. Then there exists a unique BPS on X_f compatible with those on S and T. Moreover, if $g : U \to V$ is another homeomorphism which is projective on the interiors and such that $g^{-1}f$ extends to a homeomorphism of S which is projective in the interior, then $X_g := {S \sqcup T}/{\sim_g}$ is projectively diffeomorphic to X_f.*

Proof. We just have to check that the only possible choice works. Define an atlas for X_f by putting together all the charts from S and T. The change of coordinates between an S-chart and a T-chart is exactly given by f, which is projective by hypothesis.

Let g be another gluing map such that $g^{-1}f$ extends to a projective diffeomorphism h of S. Then the map

$$H : X_f \to X_g \ , \ H(x) := \begin{cases} h(x) & x \in S \\ x & x \in T \end{cases}$$

is well defined by construction of X_f and X_g and establishes a projective diffeomorphism between them. □

Lemma 1.4.6. *Let S and T be surfaces with a BPS, let $U \subseteq S$ and $V \subseteq T$ be open subsets, let $\gamma : [0,1] \to U$ and $\eta : [0,1] \to V$ be embedded continuous arcs. Then there is at most one projective diffeomorphism $f : U \to V$ such that $f\gamma = \eta$.*

Proof. In local charts such a map is given by Möbius transformations and has fixed behaviour on more than three points (actually infinitely many); then f is uniquely determined in every chart. □

Definition 1.4.7. Let S and T be be closed surfaces with a BPS, $\gamma \subset U \subset S$ and $\eta \subset V \subset T$ two simple continuous arcs contained in open connected simply connected neighbourhoods in S and T respectively and let $f : U \to V$ be a projective diffeomorphim such that $f(\gamma) = \eta$. We cut S and T along the two arcs and compactify the resulting surfaces to get surfaces with boundary S', T' (for instance take the metric completion of $S \setminus \gamma$ and $T \setminus \eta$ with respect to any hyperbolic metric on them for which the end has trivial holonomy). Notice that f can be uniquely extended to the boundary components in a consistent way. Then we define X_f to be $X_f := {S' \sqcup T'}/{\sim_f}$ (notations as in 1.4.5). We will call X_f the **cut and paste** of S and T along the arcs γ and η.

Lemma 1.4.8. *(the BPS on a cut and paste) Let S and T be surfaces with a BPS, $H : [0,1] \times [0,1] \to S, H(s,t) = \gamma_s(t)$ and $K : [0,1] \times [0,1] \to T, K(s,t) = \eta_s(t)$ be isotopies relative to endpoints of embedded continuous arcs. Let U and V be open* *connected simply connected neighbourhoods of $Im(H)$ and $Im(K)$ respectively such that $\forall s \in [0,1]$ there exists a projective diffeomorphism $f_s : (U, \gamma_s) \to (V, \eta_s)$. Let X_s be the cut and paste of S and T along γ_s and η_s. The all the X_s are projectively diffeomorphic.*

Proof. By lemma 1.4.5, X_s carries a unique BPS compatible with those of S and by lemma 1.4.6 the gluing maps f_s are uniquely determined. We claim that $f_s = f_0$ for any s. We consider $g := f_s f_0^{-1} : U \to U$. By construction g is a projective transformation which fixes two points (namely $\gamma_0(0)$ and $\gamma_0(1)$) with an invariant connected simply connected open set (i.e. U). By lemma 1.4.4, since U cannot be a whole $\mathbb{CP}^1$, g must be the identity, i.e. $f_s = f_0$. But then g extends trivially to an automorphism of S (namely the identity), hence $X_s \cong X_0$ by lemma 1.4.5. □

Let us conclude this section with a few remarks. From the topological point of view, the cut and paste is a connected sum; in the case one of the two surfaces involved is a sphere it is a topologically trivial operation, hence it preserves the holonomy. Two new branch points of order 1 are created at the endpoints of the arcs used to perform the construction. It is true that the isomorphism class of the resulting BPS depends only the relative isotopy class of the arcs only if at each time of the isotopy we are still able to continuously define a projective isomorphism between uniform neighbourhoods; on the other hand independent isotopies of the two arcs can result in non isomorphic structures, already when one of the two surfaces is a sphere, as we will show in the next section.

1.4.3 Bubbling

We apply the construction of the previous paragraph in the special setting in which T is a sphere endowed with the standard projective structure of $\mathbb{CP}^1$ (the only unbranched one), in order to fix terminology and notation. In the following σ is a BPS on S.

Definition 1.4.9. Let $\beta \subset S$ be a simple arc, and let (dev, ρ) be a development-holonomy pair defining σ. We say that β is bubbleable if it is injectively developed, i.e. the restriction of the developing map to any of its lifts $\widetilde{\beta} \subset \widetilde{S}$ is injective. An isotopy β_t of such an arc relative to endpoints is said to be a bubbleable isotopy if for any time β_t is still bubbleable.

Definition 1.4.10. If β is a bubbleable arc, let U be a connected simply connected neighbourhood of β. Then the restriction of a developing map dev establishes a projective diffeomorphism $dev : (\widetilde{U}, \widetilde{\beta}) \to (dev(\widetilde{U}), dev(\widetilde{\beta})) \subset \mathbb{CP}^1$ for any lift. We can perform a cut and paste of $\widetilde{S}$ and a copy of $\mathbb{CP}^1$ along $\widetilde{\beta}$ and $dev(\widetilde{\beta})$ for any lift in an equivariant way to obtain a new BPS on $\widetilde{S}$. The action of $\pi_1(S)$ on $\widetilde{S}$ clearly extends to an action by covering transformation on it, using the action via the holonomy representation inside the glued copies of $\mathbb{CP}^1$, so that we get a well defined BPS on S. We call this structure the **bubbling** of σ along β and denote it by $Bub(\sigma, \beta)$.

An immediate consequence of the previous discussion is the following.

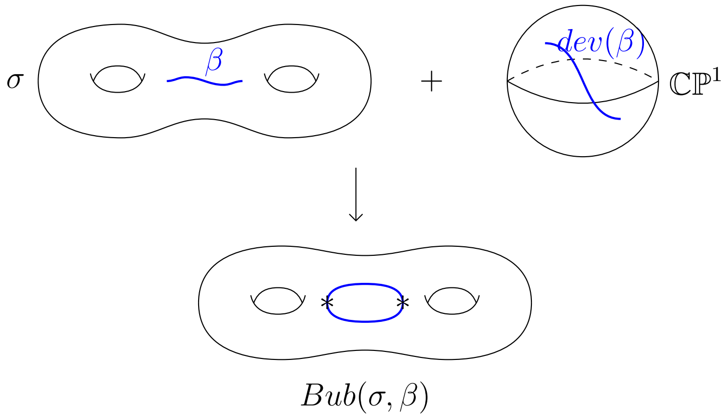

Figure 1.3: Bubbling a surface

Corollary 1.4.11. *The bubbling of a BPS along a bubbleable arc carries a unique BPS, whose isomorphism class depends only on the bubbleable isotopy class of the arc.*

Remark 1.4.12. The above result is not true if we allow non bubbleable isotopies, i.e. isotopies which at some time produce an embedded path whose developed image is not embedded. An explicit example is given below in 1.4.14.

The following lemma gives some more details: we can indeed always do some very small deformations; in other words, a bubbleable arc always admits bubbleable isotopies.

Lemma 1.4.13. *Let σ be a BPS on S. Let $\gamma \subset S$ be an embedded arc having embedded developed image and not going through branch points (except possibly at its endpoints). Let $H : [-1,1] \times [0,1] \to S$ be any isotopy relative to endpoints such that $H(0,.) = \gamma$. Then $\exists\ \varepsilon > 0$ such that $\forall\ s \in]-\varepsilon, \varepsilon[$ the arc $\gamma_s = H(s,.)$ is still injectively developed.*

Proof. By contradiction let $s_n \to 0$ be a sequence in $[-1,1]$ such that $\gamma_n := \gamma_{s_n}$ is not injectively developed, i.e. $\forall\ n, \exists\ t_n, t'_n \in [0,1]$ such that $dev\widetilde{\gamma_n}(t_n) = dev\widetilde{\gamma_n}(t'_n)$. By compactness we can extract convergent subsequences of t_n and t'_n and take t, t' as the limits. By continuity of H in the couple (s,t), we can pass simultaneously to the limit and get $dev\widetilde{\gamma}(t) = dev\widetilde{\gamma}(t')$. Since γ is known to be injectively developed, necessarily $t = t'$. But then we have

$$\widetilde{\gamma_n}(t_n) \to \widetilde{\gamma}(t) = \widetilde{\gamma}(t') \leftarrow \widetilde{\gamma_n}(t'_n)$$

thus for large n we have points very closed to $\widetilde{\gamma}(t)$ and such that $dev\widetilde{\gamma_n}(t_n) = dev\widetilde{\gamma_n}(t'_n)$ (by the hypothesis of contradiction). If $\widetilde{\gamma}(t)$ is not a branch point, then dev is locally injective around it. If $\widetilde{\gamma}(t)$ is a branch point then we can assume it is the initial point of γ (i.e. $t = 0$). Even in this case we still can get an absurd: let us say $\gamma(0)$ is a branch point of order k, then a neighbourhood of it decomposes in $k+1$ sectors of angle 2π; since γ_n is an isotopy of γ, for n large enough its initial segment belongs to the same sector which contains the initial segment of γ; but the developing map is injective on each of these sectors. ☐

In other words the set $\{s \in [-1,1] \mid \gamma_s \text{ is bubbleable }\}$ is open. Once more, there is no reason why it should be closed, and indeed it is not in general. The following is an explicit example of two non isomorphic structures obtained as bubbling on two arcs which are isotopic relative to endpoints, but via an isotopy which does not keep the developed arc embedded at every time. The same technique was already used in [9] to produce a similar example in the case of grafting.

Example 1.4.14. To describe this example we will use some of the terminology and basic facts about structures with quasi-Fuchsian holonomy which are introduced in the next Chapter 2. Let S be a genus 2 surface and fix some hyperbolic structure σ_ρ on it. Let γ be an oriented closed geodesic which disconnects S into two one-holed tori. Let η be an oriented embedded geodesic arc on S with one endpoint x on γ and orthogonally intersecting γ only in x; call y the other endpoint, which we assume to be on the right of γ. We want to perform a grafting of σ_ρ along γ and then show how to perform two different bubbling on $Gr(\sigma_\rho, \gamma)$ along two different extensions of η. As observed in 1.4.3, on $Gr(\sigma_\rho, \gamma)$ we have two distinguished curves $\gamma^\pm$ coming from γ and bounding the grafting annulus A_γ. We also have two marked points $x^\pm \in \gamma^\pm$ coming from the point x, and an arc coming from η, which we still denote by the same name, which starts at $x^+ \in \gamma^+$ orthogonally and moves away from the annulus.

There is a natural way to extend η by analytic continuation to an embedded arc reaching the other point $x^- \in \gamma^-$: namely consider the extension of the developed image of η (which is a small geodesic arc in the upper half-plane) to a great circle $\widehat{\eta}$ on $\mathbb{CP}^1$. This gives an embedded arc on $Gr(\sigma_\rho, \gamma)$ which is geodesic (outside the real curve), but which is not injectively developed: its developed image goes twice over the developed image of η. Therefore we can not bubble on it.

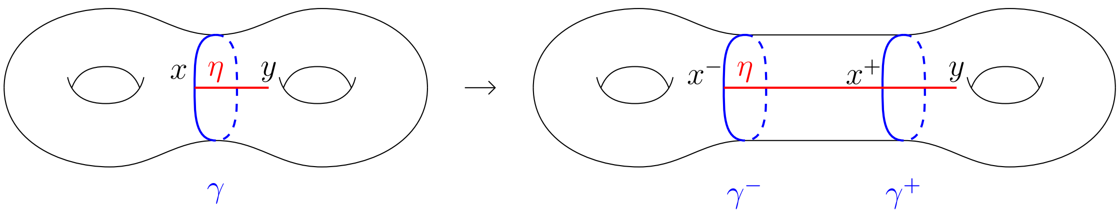

Figure 1.4: Analytic extension of η in $Gr(\sigma_\rho, \gamma)$

To obtain bubbleable arcs we slightly perturb this construction; in $\mathbb{CP}^1$ consider an embedded arc which starts at the developed image $\widehat{x}$ of x and ends at the developed image $\widehat{y}$ of y, but leaves $\widehat{x}$ with a small angle θ with respect to $\widehat{\eta}$, stays close to it, and reaches $\widehat{y}$ with angle θ on the other side, crossing $\widehat{\eta}$ just once at some point in the lower-half plane (see left side of Picture 1.5). This arc can be chosen to sit inside a fundamental domain for $\rho(\gamma)$, so that it gives an embedded arc inside the grafting annulus A_γ of $Gr(\sigma_\rho, \gamma)$ starting at x^-, reaching γ^+ at a point z^+ close to x^+ and ending at y. Changing the value of θ in some small interval $]-\varepsilon, \varepsilon[$ we obtain a family of embedded arcs α_θ in $Gr(\sigma_\rho, \gamma)$ which are isotopic relative to the endpoints x^-, y and are all injectively developed, except $\alpha_0 = \eta$.

Fix now some small θ and consider the BPS obtained by bubbling along $\alpha_{\pm\theta}$,

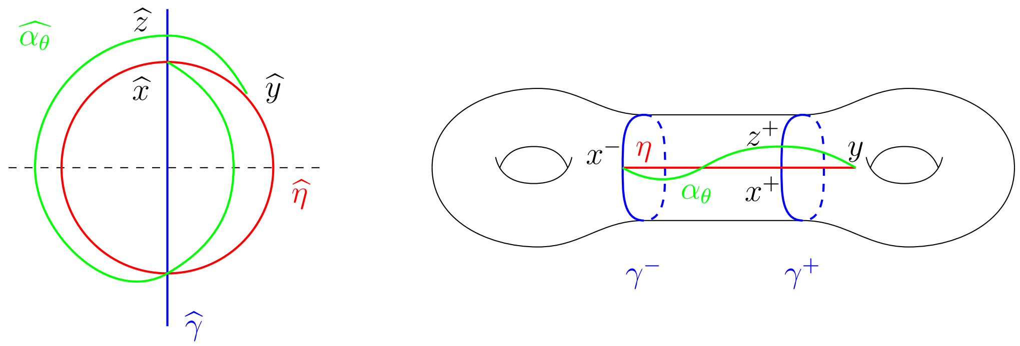

Figure 1.5: The bubbleable arc α_θ in $\mathbb{CP}^1$ and $Gr(\sigma_\rho, \gamma)$

i.e. $\sigma_\pm = Bub(Gr(\sigma_\rho, \gamma), \alpha_{\pm\theta})$. We now proceed to show that these two BPSs are not isomorphic: they can be distinguished by looking at the configuration of certain curves, which we now define. The first curve we need is the analytic continuation of γ^+ into the positive part: namely we extend it inside the bubble by following its developed image. The result is a curve which hits the real curve and reaches x^-, and we still denote it by γ^+. The other curve is the unique geodesic δ between x^- and y inside the bubble, which develops isometrically onto the developed image of the original geodesic segment η of σ_ρ. Notice that the whole construction can be made in such a way that this is indeed the shortest geodesic between its endpoints in the positive region. Indeed, let δ' be a shorter geodesic in the positive region between the same pair of points. Since it is a CAT(-1) space, we can look at the unique geodesic representative ω in the free homotopy class of $\delta'\delta^{-1}$, whose length is bounded by twice the length of δ. Therefore it is enough to choose on σ_ρ the points x, y to be very close with respect to the systole of σ_ρ.

Now we look at the tangent space at x^-. The tangent vector to γ^+ at x^- sits on the right or on the left of the tangent vector to δ (with respect to the underlying orientation of S) depending on the fact that we look at σ_+ or at σ_-. But a projective isomorphism between the two structures should be orientation preserving at x^-. Notice that consistently saying that a tangent direction comes to the left or right of another requires a check of the amplitude of the angles involved; but the angle we are interested in is clearly less the 2π (since both curves enter the bubble), i.e. less then half the total angle at x^-.

Following the ideas of 1.4.13, we can also prove the following lemma, which similarly allows to perform local deformations of an injectively developed path in an injectively developed way.

Lemma 1.4.15. *Let σ be a BPS on S. Let $\gamma : [0, 1] \to S$ be an embedded arc having embedded developed image and not going through branch points (except possibly at its endpoints). Then there exists an injectively developed subset $U \subseteq S$ such that $\gamma \subset U$ and $\gamma(]0, 1[) \subset int(U)$.*

Proof. Let us choose a lift $\widetilde{\gamma}$ and work in the universal cover for simplicity. If γ does not go through any branch point at all, then we just consider a sequence of nested open neighbourhoods $U_{n+1} \subsetneq U_n$ of γ and prove that for n large enough U_n

is injectively developed. Assume by contradiction that $\forall\, n \in \mathbb{N}$ we can find a couple of distinct point $x_n, y_n \in U_n$ such that $dev(\widetilde{x}_n) = dev(\widetilde{y}_n)$. By compactness of S these sequences subconverge to $x, y \in \gamma$. Since the path is injectively developed we get $x = y$ and since it does not go through branch points we reach an absurd, exactly as in the proof of 1.4.13. If one endpoint (let us say it is $\gamma(0)$) of γ is a branch point of order k, then clearly every set containing it in its interior is not injectively developed. Once again we observe that nevertheless a sufficiently small neighbourhood Ω of $\gamma(0)$ decomposes as a disjoint union of injectively developed sectors $A_1, \dots, A_{k+1}$, and that an initial segment of γ belongs to one of them (let us say it is A_1); so we can simply pick a sequence of nested sets $V_{n+1} \subsetneq V_n$ such that for every $n \in \mathbb{N}$ and for every $\varepsilon > 0$ we have that $\gamma(0) \in V_n$, $V_n \cap \Omega \subsetneq A_1$ and V_n contains $\gamma(]\varepsilon, 1-\varepsilon[)$ in its interior, and apply the previous argument to obtain a couple of sequences $x_n \neq y_n \in V_n$ converging to $x = y \in \gamma$. The non trivial case to discuss is the case in which the limit is a branch point, i.e. $x = y = \gamma(0)$; by construction of V_n, for n large enough the points x_n, y_n must fall inside A_1, which is injectively developed, hence we reach a contradiction exactly as in 1.4.13. □

Let us get back to the description of the properties of this particular case of cut and paste. As noticed before, this surgery preserves the holonomy and introduces a couple of simple branch points corresponding to the endpoints of the bubbling arc; thus it induces a map

$$Bub : \mathcal{M}_{\lambda,\rho} \times \mathcal{BA} \to \mathcal{M}_{\lambda+(1,1),\rho}$$

where $\mathcal{BA}$ is the set of bubbleable arcs on S, λ is any partition of some natural number $k \in \mathbb{N}$ and $\lambda + (1,1)$ is the partition of $k+2$ obtained appending $(1,1)$ to λ. Notice that for any structure σ and any bubbleable arc β on it the structure $Bub(\sigma, \beta)$ has surjective developing map onto $\mathbb{CP}^1$; in any case it is nor a complete structure, neither a uniform one.

Once we have performed a bubbling, we see a subsurface of S homeomorphic to a disk and isomorphic to $\mathbb{CP}^1$ cut along a simple arc, the isomorphism being given by any determination of the developing map itself. It is useful to be able to recognise this kind of subsurface, since there is an obvious way to remove it and lower the branching order by 2; therefore we find it convenient to give the following definition.

Definition 1.4.16. A **bubble** on $\sigma \in \mathcal{BP}(S)$ is an embedded closed disk $B \subset S$ whose boundary decomposes as $\partial B = \beta' \cup \{x, y\} \cup \beta''$ where $\{x, y\}$ are simple branch points of σ and β', β'' are embedded injectively developed arcs which overlap once developed; more precisely there exist a determination of the developing map on B which injectively maps β', β'' to the same simple arc $\widehat{\beta} \subset \mathbb{CP}^1$ and restricts to a diffeomorphism $dev : int(B) \to \mathbb{CP}^1 \setminus \widehat{\beta}$.

Of course if β is a bubbleable arc on σ with endpoints x, y, then the structure $Bub(\sigma, \beta)$ displays a subsurface which is a bubble. As said above, the nice things about bubbles is that they provide a standard way to lower the branching order of a structure.

Definition 1.4.17. Given a bubble B on a structure σ, we can define a new structure $Deb(\sigma, B)$ by removing $int(B)$ and collapsing ∂B to a single arc β. This structure is called the debubbling of σ with respect to B.

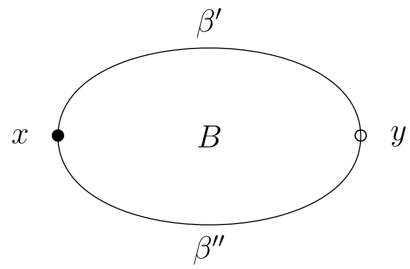

Figure 1.6: A bubble

Of course the two operations of bubbling and debubbling are one the inverse of the other. Of course it is possible to perform simultaneous bubbling on any collection of disjoint bubbleable arcs.

Notice that BPS obtained by bubblings from unbranched structures have by definition an even number of branch points and surjective developing map. As a consequence branched hyperbolic structures with an even number of branch points do not arise as bubblings, as their developing map takes value only in the upper-half plane. By the work of [34] these structures exist on every Riemann surface of genus $g \geq 3$. This example was already mentioned in [8].

Even in the case of surjective developing map, it is not at all clear in general whether a BPS with an even number of branch points is a bubbling. For instance consider any BPS σ, perform a bubbling, then perform another bubbling on an arc intersecting the bubble. Of course the introduction of the second bubble destroys the first one, thus it is not clear a priori that this new structure is still obtainable as a simultaneous bubbling of σ over a couple of disjoint bubbleable arcs. From this point of view the grafting operation is more stable: it was proved in [9, Proposition 3.3] that a structure obtained by two consecutive graftings on transverse multicurves can actually be obtained by a single multigrafting on the same underlying structure. The following question is posed as Problem 12.1.2 in [13]

Problem: given $\sigma_1, \sigma_2 \in \mathcal{M}_\rho$, is it possible to pass from one to the other using the operations of grafting, degrafting, bubbling and debubbling?

The main result in [8] provides a positive answer in the case of quasi-Fuchsian holonomy (see next Chapter 2 for definitions) and if an additional surgery, to which we dedicate the next section, is allowed.

1.4.4 Movements of branch points

In this section we introduce deformations obtained by replacing a chart of a given atlas by a new chart. This construction works for any kind of branched (G, X)-structures on a manifold, i.e. structures for which charts are possibly branched covers over open subsets of X and change of coordinates are given by restriction of elements of G, but we stick to the case of BPSs on a surface.

We need to fix some notations. Let S be a compact surface endowed with a

BPS $\sigma \in \mathcal{BP}(S)$. Let $\mathcal{U} = \{(U, \varphi)\}$ be an atlas of projective charts defining σ. It is not restrictive to assume that each chart contains at most one branch point and that pairwise intersections of chart domains do not contain branch points. By compactness we can also suppose that $\mathcal{U}$ is finite. Let $p \in S$ and $(A, d) \in \mathcal{U}$ be a chart with $p \in A$. By the previous assumptions we have that $A \setminus \cup_{(U,\varphi)\in\mathcal{U}, U\neq A} U$ is a neighbourhood of p with non empty interior. In particular we can find an open set B such that $p \in B \subset A$ and $U \cap A \subseteq A \setminus B$ for any chart domain $U \neq A$. Now consider a map $d' : A \to \mathbb{CP}^1$ such that $d' = d$ on $A \setminus B$ and which is a branched cover onto its image, with finite degree and finitely many critical values on $\mathbb{CP}^1$. Then let $\mathcal{U}' = \mathcal{U} \setminus \{(A, d)\} \cup \{(A, d')\}$.

Lemma 1.4.18. *$\mathcal{U}'$ is a branched complex projective atlas on S.*

Proof. By construction elements of $\mathcal{U}'$ provide a cover of S by branched complex projective charts. We need to check the change of coordinates. So let (U, φ) be a chart with $U \cap A \neq \varnothing$; since $\mathcal{U}$ was a legitimate atlas, there exists some $g_U \in \mathrm{PSL}_2\mathbb{C}$ such that $g_U \circ \varphi = d$ on $U \cap A$. But $U \cap A \subseteq A \setminus B$ and $d' = d$ on $A \setminus B$, hence $g_U \circ \varphi = d'$ on $U \cap A$, which implies that the change of coordinates between φ and d' is still given by the restriction of $g_U \in \mathrm{PSL}_2\mathbb{C}$. □

Definition 1.4.19. We say that the BPS σ' defined by the maximal atlas containing $\mathcal{U}'$ is a BPS obtained from σ by a **local deformation at** p.

It turns out that performing a local deformation at a non branched point without changing the codomain of the local chart results in an isomorphic structure.

Lemma 1.4.20. *If p is not a branch point for σ, nor for d', and if $d(A) = d'(A)$, then the BPS σ' obtained by local deformation at p is isomorphic to σ.*

Proof. Let us keep using the above notations, and let us define a map $F : (S, \mathcal{U}) \to (S, \mathcal{U}')$ to be $F(x) = x$ for $x \notin A$ and $F(x) = d'^{-1}d(x)$ for $x \in A$. Notice that this is well defined because $d(A) = d'(A)$, because the local charts at a non branched point are diffeomorphisms and because $d = d'$ on $A \setminus B$. Then F is a diffeomorphism which is read as the identity in every chart: the unique non trivial computation being the one relative to the chart (A, d) in the domain and (A, d') in the codomain:

$$d' \circ F \circ d^{-1} = d' \circ d'^{-1} \circ d \circ d^{-1} = id$$

In particular it is projective. □

Examples of this kind of trivial deformations are obtained as follows. In the above notations, let $f : d(A) \to d(A)$ be a diffeomorphism compactly supported inside $d(B)$; then define $d' = f \circ d$. If p is not a branch point for d then it is not a branch point for d' and the lemma applies. On the other hand when p is a branch point, even this simple kind of local deformation turns out to provide a rich deformation theory. For instance one can replace the branched cover $d : B \to d(B)$ with any other branched cover $d' : B' \to d(B)$ taken from a suitable space of deformations of it, which we now define, following [8].

Definition 1.4.21. Let $\mathbb{D}$ denote the closed unit disk in $\mathbb{C}$ and let U be a closed disk. Consider a smooth branched cover $\pi : U \to \mathbb{D}$ of degree m with critical values in the interior of $\mathbb{D}$ and a diffeomorphism $f : \partial\mathbb{D} \to \partial U$ such that $\pi \circ f(z) = z^m$ for $z \in \partial\mathbb{D}$. We consider the set $Def(U, \pi, f) = \{(U', \pi', f') \mid U' \text{ is a closed disk}, \pi' : U' \to \mathbb{D} \text{ is a smooth branched cover of degree } m \text{ without critical values on } \partial\mathbb{D}, f : \partial U \to \partial U' \text{ is a diffeomorphism such that } \pi = \pi' \circ f' \text{ on } \partial U\}$. We impose an equivalence relation on $Def(U, \pi, f)$ by declaring $(U_1, \pi_1, f_1) \sim (U_2, \pi_2, f_2)$ if and only if the diffeomorphism $f_2 \circ f_1^{-1} : \partial U_1 \to \partial U_2$ can be extended to a diffeomorphism $F : U_1 \to U_2$ such that $\pi_1 = \pi_2 \circ F$. The quotient $\mathcal{H}(\pi) = {Def(U, \pi, f)}/{\sim}$ is called the **Hurwitz space** of the (marked) branched cover $\pi : U \to \mathbb{D}$.

In [8, Lemma 12.7] the space $\mathcal{H}(\pi)$ is shown to be in bijection with on open subset of a $(m-1)$-dimensional complex vector space of polynomials, which endows it with the structure of a complex manifold of dimension $m-1$. Moreover, by sending a branched cover to the divisor given by its critical values we obtain a map $Crit : \pi \to Sym^{m-1}\mathbb{D}$ which realises the Hurwitz space as a holomorphic branched cover over a $(m-1)$-dimensional complex manifold.

Now, if $\sigma \in \mathcal{M}_{k,\rho}$, $p \in \sigma$ is a branch point of order m and (A, d) is a local chart at p for σ, we can consider the Hurwitz space $\mathcal{H}(d)$ (notice that up to a projective transformation we can always assume that charts take value in the unit disk $\mathbb{D}$). For any other branched cover $d' \in \mathcal{H}(d)$ we can perform a local deformations in the sense of 1.4.19, the diffeomorphisms of the boundaries of the disks being used to perform the gluing in a well-defined way. This is explained in detail in [8, §12.5]. There it is also shown that performing local deformations parametrised by Hurwitz spaces at all branch points of σ provides a parametrisation of a full neighbourhood of σ in $\mathcal{M}_{k,\rho}$ as soon as the holonomy representation is complicated enough. More precisely they prove the following.

Theorem 1.4.22. *Let S be a closed, connected and oriented surface of genus $g \geq 2$ and $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ a non elementary representation. Then $\mathcal{M}_{k,\rho}$ is a complex manifold of dimension k, with an atlas valued in suitable products of Hurwitz spaces.*

Indeed their proof shows that for any partition λ of k the stratum $\mathcal{M}_{\lambda,\rho}$ is a smooth complex submanifold of dimension equal to the length of λ. In particular the principal stratum $\mathcal{M}_{(1,\dots,1),\rho}$ is an open dense submanifold, and the inclusion $\mathcal{M}_{(1,\dots,1),\rho} \hookrightarrow \mathcal{M}_{k,\rho}$induces a bijection on the sets of connected components. Also notice that all of this is consistent with 1.4.20, since in that case $d = 1$.

We want now to give a description of these deformations in terms of geometric surgeries on the surface. To simplify the discussion we consider only the deformations that preserve the structure of the branching divisor of σ, i.e. the partition λ defining the stratum σ lives in. These come from elements of the Hurwitz space which are mapped by the map $Crit : \pi \to Sym^{m-1}\mathbb{D}$ inside the same generalized diagonal defined by λ, so that we can reduce the situation to a lower-dimensional unbranched cover; in other words, these deformations do not involve any collapsing or splitting of branch points. We will come back to these more general deformations at the end of this section (see 1.4.27).

Let $\sigma \in \mathcal{M}_{\lambda,\rho}$. By the above discussion, the deformations keeping the structure inside its stratum come in a n-dimensional family, where n is the length λ, i.e. the number of branch points of σ; in other words we have to consider a 1-dimensional family of local deformations at each point. We want to define explicit local deformation which are geometrically easy to visualise and control and which realise this 1-dimensional family of deformations.

So let $p \in \sigma$ be a branch point of order m and let us choose an atlas $\mathcal{U} = \{(U, \varphi)\}$ defining σ with the properties used in the definition of local deformations above (see 1.4.19), namely let (A, d) be a local chart at p and B an open set such that $p \in B \subset A$ and $U \cap A \subseteq A \setminus B$ for any chart domain $U \neq A$. Now pick any point $q \in d(B)$ and let φ_q be an isotopy of $d(A)$ with compact support inside $d(B)$, such that that $t \mapsto \varphi_q(d(p), t)$ describes a continuous embedded curve $\widehat{\mu}$ from $d(p)$ to q. Then $d_q : A \to d(A), d_q = \varphi_q \circ d$ gives us an element in $\mathcal{H}(d)$, with which we can perform a local deformation, and if we take the preimage of $\widehat{\mu}$ in S with respect to d, we obtain a collection of $m+1$ embedded paths $\mu = \{\mu_1, \dots, \mu_{m+1}\}$ which meet at p, are disjoint otherwise and end at points $q_1, \dots, q_{m+1}$ such that $d(q_i) = q$.

Definition 1.4.23. We call this local deformation at p determined by the choice of q a **movement of branch point** at p towards q. We also refer to μ as the embedded twin $(m+1)$-pod of the movement and say that the deformation is a movement of branch point along μ. We will denote the resulting structure by $Move(\sigma, \mu)$.

Since a neighbourhood of $d(p)$ is 1-dimensional, by the above discussion this gives a full description of the local deformations at p which preserve the structure of the branching divisor. In other words, if $\sigma \in \mathcal{M}_{\lambda,\rho}$, then these deformations account for a full neighbourhood of σ inside $\mathcal{M}_{\lambda,\rho}$.

Remark 1.4.24. The choice of q is actually the only parameter in the deformation, so we can always choose the isotopy φ_q to be a straight-line isotopy on its support, i.e. we can always choose the curve $\widehat{\mu}(t) = \varphi_q(d(p), t)$ to be the straight-line segment from $d(p)$ to q. We will find this useful in the following chapters, but do not really need it here.

It was already noticed in [32] (for simple branch points) and [8] (in the general case) that the structure $Move(\sigma, \mu)$ can be obtained by the following cut-and-paste construction on S: cut S along μ to open a star-like buttonhole, then glue a side to the other adjacent side to close it (see Picture 1.7 for the case of a simple branch point, i.e $m = 1$). This gives a deformation of the structure around p such that the new chart is the branched cover d_q defined above. Notice that we can reverse the deformation by considering the inverse isotopy φ_q^{-1} or, equivalently, by operating a cut-and-paste construction on the induced collection of arcs μ'. The above cut-and-paste construction is actually defined without any reference to local charts, local isotopies or Hurwitz spaces: it just needs any configuration of arcs which satisfies the following definition.

Definition 1.4.25. If p is a branch point of σ, then an **embedded twin n-pod** based at p is a collection $\gamma = \{\gamma_1, \dots, \gamma_n\}$ of n embedded paths which meet at p, are disjoint outside p and injectively develop to an embedded arc $\widehat{\gamma} \subset \mathbb{CP}^1$, i.e. there exists a determination of the developing map (i.e. a local chart) d at p which maps

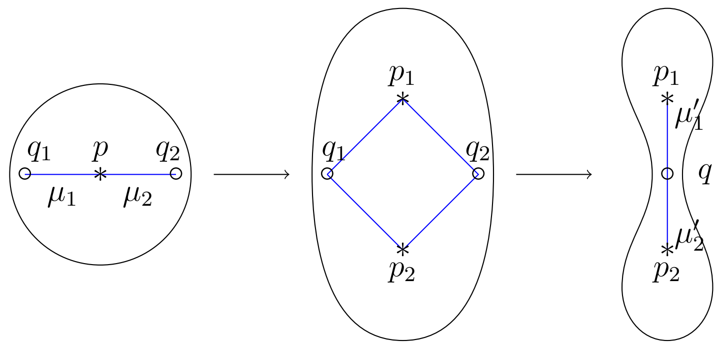

Figure 1.7: A movement of branch point

diffeomorphically γ_i to $\widehat{\gamma}$. When $n = 2$ we also call $\gamma = \{\gamma_1, \gamma_2\}$ an **embedded twin pair**.

The structure obtained by the above cut-and-paste surgery along an $(m+1)$-pod γ at a branch point of order m is still denoted by $Move(\sigma, \gamma)$ and called the movement of branch point at p along γ. Since, as already observed, the structures defined in 1.4.23 fill a full neighbourhood of local deformations of σ at p, we are not really introducing any new deformations here, just a different point of view. In the first definition the stress is on the point we want to move to, in the second one it is on the path we want to use for the journey. To safely exploit this second point of view, we need to check the dependence on the choice of the $(m+1)$-pod among all the $(m+1)$-pods with the same endpoints, similarly to what we have done for the bubbling operation.

We will now restrict to the case of a simple branch point (i.e. $m = 1$), since we will need the second point of view only in this case. So let us consider a simple branch point p on a BPS σ and two embedded twin pairs $\mu = \{\mu_1, \mu_2\}, \nu = \{\nu_1, \nu_2\}$ based at p. By the above discussion, if μ and ν are entirely contained in a chart around p and developing to arcs $\widehat{\mu}, \widehat{\nu}$ with the same endpoints, then $Move(\sigma, \mu) = Move(\sigma, \nu)$. Notice that in this case μ_i and ν_i are isotopic relatively to their endpoints through an injectively developed isotopy. However we can more generally consider μ and ν ending the same points without being contained inside a single chart; in this case the isotopy class of the arcs in the embedded twin pair we use might be relevant, and different choices might give rise to non isomorphic structures. We can prove the following reassuring result.

Lemma 1.4.26. *Let $\sigma \in \mathcal{M}_{k,\rho}$ and let p be a simple branch point. Let $\mu = \{\mu_1, \mu_2\}$ and $\nu = \{\nu_1, \nu_2\}$ be embedded twin pairs based at p with the same endpoints, and let q_i be the common endpoint of μ_i and ν_i for $i = 1, 2$. Suppose that there exists an isotopy $H : [0,1] \times [-1,1] \to S$ such that*

1. *H is an isotopy from μ to ν i.e. $H(0,t) = \mu_1(t)$ for $t \in [-1,0]$, $H(0,t) = \mu_2(t)$ for $t \in [0,1]$, $H(1,t) = \nu_1(t)$ for $t \in [-1,0]$ and $H(1,t) = \nu_2(t)$ for $t \in [0,1]$*
2. *H is relative to endpoints, i.e. $H(s,-1) = q_1$, $H(s,0) = p$, and $H(s,1) = q_2$ for all $s \in [0,1]$*

3. *H is an isotopy of embedded twin pairs, i.e.* $\alpha^s = \{\alpha_1^s = H(s, [-1,0]),\ \alpha_2^s = H(s, [0,1])\}$ *is an embedded twin pair for all* $s \in [0,1]$

Then $Move(\sigma, \mu) = Move(\sigma, \nu)$.

Proof. First of all notice that each path α_i^s appearing in an embedded twin pair α^s is in particular an embedded arc which is injectively developed and goes through exactly one branch point, which is p. Therefore we can pick an injectively developed set U_i^s containing $\alpha_i^s \setminus \{p\}$ in its interior as in 1.4.15. We can choose this set in such a way that $U^s = U_1^s \cup U_2^s$ is an open neighbourhood of α^s: for instance we can take U_1^s such that its developed image is an open neighbourhood of the developed image of α^s, then pull it back via the developing map, so that U^s is the domain of a local projective chart which simply branches at p and contains the whole embedded twin pair α^s.

The sets U^s provide an open cover of $Im(H)$; by compactness we extract a finite subcover indexed by some $s_0 = 0, s_1, \ldots, s_N = 1$. Up to taking an intermediate finite subcover between $\{U^{s_0}, \ldots, U^{s_N}\}$ and $\{U^s \mid s \in [0,1]\}$ we can assume that the local chart U^{s_i} contains not only α^{s_i} but also $\alpha^{s_{i\pm1}}$. Then we conclude by observing that $\alpha^{s_0} = \mu$ and $\alpha^{s_N} = \nu$ and that, as remarked above, the results in the Appendix of [8] imply that $Move(\sigma, \alpha^{s_i}) = Move(\sigma, \alpha^{s_{i+1}})$, because α^{s_i} and $\alpha^{s_{i+1}}$ are contained in a single local chart. □

Let us conclude this section with the following observation, which is actually not needed in the sequel.

Remark 1.4.27. The above description works also for local deformations which do not preserve the partition λ encoding the structure of the branching divisor. For instance by a local deformation we can split a branch point of order m into lower-order branch points. This can be done by choosing a point q in a local coordinate (A, d) around p and an embedded twin n-pod strictly contained in the embedded twin $(m+1)$-pod at p induced by q, for some $2 \leq n < m$. Of course this does not only depend on the choice of q, but also on that of the n-pod, which can be done in $\binom{m}{n}$ ways, corresponding to the fact that the map $Crit : \mathcal{H}(d) \to Sym^m d(A)$ genuinely branches along generalised diagonals. As a result we can not write the new chart as the old one postcomposed with an isotopy, as we did in the discussion above, but we need to pick a branched cover in the Hurwitz space $\mathcal{H}(d)$ with a different branching structure. This point of view extends the deformations defined in [32] for simple branch points to higher order branch points.

Chapter 2

Combinatorics in quasi-Fuchsian holonomy

Throughout this chapter we are interested in structures on S whose holonomy preserves a decomposition of the model space $\mathbb{CP}^1$ into two disks separated by a Jordan curve. The key feature of such a representation is the presence of a canonical decomposition of the surface into pieces which carry complete hyperbolic structures with ideal boundary.

2.1 Geometric decomposition of a quasi-Fuchsian BPS and real decomposition of holonomy fibres

Let S be a closed, connected and oriented surface of genus $g \geq 2$.

Definition 2.1.1. A **Fuchsian** (respectively **quasi-Fuchsian**) **group** is a subgroup of $\mathrm{PSL}_2\mathbb{C}$ whose limit set is $\mathbb{RP}^1$ (respectively a Jordan curve) on $\mathbb{CP}^1$.

We refer to the appendix (see 5.2) for more background on (quasi-)Fuchsian subgroups and a collection of equivalent definitions. In particular we are interested in the fact that a finitely generated quasi-Fuchsian group Γ preserves a well defined decomposition $\mathbb{CP}^1 = \Omega^+_\Gamma \cup \Lambda_\Gamma \cup \Omega^-_\Gamma$ of the Riemann sphere into a pair of disks $\Omega^\pm_\Gamma$ and a Jordan curve Λ_Γ, i.e. the two components of the domain of discontinuity and the limit set of Γ. When Γ is Fuchsian this is the decomposition $\mathbb{CP}^1 = \mathcal{H}^+ \cup \mathbb{RP}^1 \cup \mathcal{H}^+$, where $\mathcal{H}^\pm$ denote the upper and lower-half plane in $\mathbb{C}$, which allows us to distinguish between Ω^+_Γ and Ω^-_Γ (see 5.2.6 in the appendix for more details).

Definition 2.1.2. A faithful representation $\rho : \pi_1(S) \hookrightarrow \mathrm{PSL}_2\mathbb{C}$ is a **Fuchsian** (respectively **quasi-Fuchsian**) **representation** if its image is a Fuchsian (respectively quasi-Fuchsian) subgroup and there exists an orientation preserving ρ-equivariant diffeomorphism $f : \widetilde{S} \to \Omega^+_{\rho(\pi_1(S))}$. A structure $\sigma \in \mathcal{BP}(S)$ is said to be Fuchsian or quasi-Fuchsian when its holonomy is.

In our setting ρ is an isomorphism between $\pi_1(S)$ and a subgroup of $\mathrm{PSL}_2\mathbb{C}$; since S is assumed to be closed, the first one is finitely generated, hence the above discussion applies and this definition is well posed. We will adopt the notation $\Omega^\pm_\rho = \Omega^\pm_{\rho(\pi_1(S))}$ and $\Lambda_\rho = \Lambda_{\rho(\pi_1(S))}$.

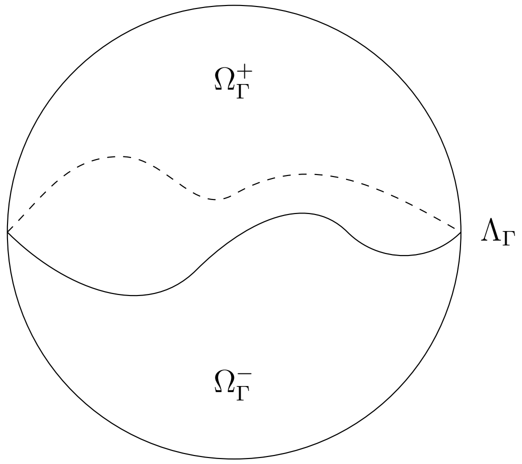

Figure 2.1: Geometric decomposition of $\mathbb{CP}^1$ under a quasi-Fuchsian group Γ.

Notice that the action on $\Omega^\pm_\rho$ admits an invariant complete hyperbolic metric $d^\pm_\rho$, since the action is conjugated to the action of a Fuchsian group on $\mathbb{H}^2$. We can therefore obtain an extended metric d_ρ on $\mathbb{CP}^1$ by considering the path metric associated to $d^\pm_\rho$: a point in one disk has infinite distance from any point of the other disk. Given a quasi-Fuchsian representation ρ, by definition we have an orientation preserving ρ-equivariant diffeomorphism $f : \widetilde{S} \to \Omega^+_\rho$. This descends to an orientation preserving diffeomorphism $F : S \to \Omega^+_\rho / Im(\rho)$, giving us a (marked) unbranched complete hyperbolic structure on S with holonomy ρ; we can use it as a base point in the moduli space $\mathcal{M}_\rho$, so we give it a special name.

Definition 2.1.3. If S is a closed, connected and oriented surface and $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ is a quasi-Fuchsian representation, then $\sigma_\rho = \Omega^+_\rho / Im(\rho)$ is called the **uniformizing structure** for ρ.

More generally, if $dev : \widetilde{S} \to \mathbb{CP}^1$ is a developing map for a BPS on S with quasi-Fuchsian holonomy ρ, then the decomposition of the Riemann sphere induced by ρ can be pulled back via the dev to obtain a decomposition of $\widetilde{S}$. Since the developing map is $(\pi_1(S), \rho)$-equivariant, this decomposition is $\pi_1(S)$-invariant and thus descends to a decomposition of the surface into possibly disconnected subsurfaces $\sigma^+, \sigma^\mathbb{R}$ and σ^-, defined as the subset of points developing to $\Omega^+_\rho, \Lambda_\rho$ and Ω^-_ρ respectively.

Definition 2.1.4. We will call $S = \sigma^+ \cup \sigma^\mathbb{R} \cup \sigma^-$ the **geometric decomposition** of S with respect to the BPS defined by the pair (dev, ρ); we will call $\sigma^\pm$ the positive/negative part of S and $\sigma^\mathbb{R}$ the real curve of S.

We already observe that, despite their apparent symmetry, the positive and negative part play very different role in the geometry of σ, because of the special role played by Ω^+_ρ in the definition 2.1.2 of quasi-Fuchsian representation. This is a phenomenon already exploited by Goldman in the unbranched case ([16]), which we will explore in the branched case below.

Notice that a priori the decomposition of the surface depends not only on the representation, but also on the choice of a developing map. However this ambiguity

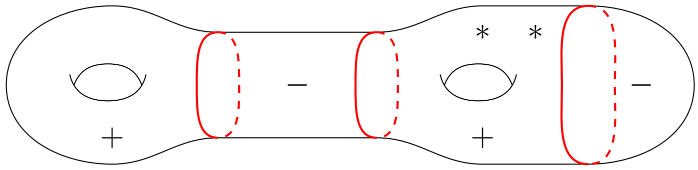

Figure 2.2: Geometric decomposition of a quasi-Fuchsian structure with $g = 2 = k$.

can be fixed by choosing a representation ρ in its conjugacy class, which is possible by the following easy observations.

Lemma 2.1.5. *A quasi-Fuchsian group* Γ *has trivial centraliser.*

Proof. A centralising element should commute with every element in Γ, hence have the same set of fix points of any element of Γ. But quasi-Fuchsian groups are non elementary, hence they contain a couple of loxodromic transformations whose fix points are different, and an element of $\mathrm{PSL}_2\mathbb{C}$ has at most 2 fix points. □

It has already been observed in 1.3.2 that this is enough to assign a unique developing map to each representative representation in the conjugacy class giving the holonomy of the structure.

Corollary 2.1.6. *If* $\{(dev, \rho)\}$ *is an equivalence class of development-holonomy pairs defining a quasi-Fuchsian BPS, then for each representation* ρ_0 *in the conjugacy class there is a unique developing map* dev_0 *in its equivalence class which is* ρ_0*-equivariant.*

Proof. Two different development-holonomy pairs for the same structure differ by a Möbius transformation, i.e. they are of the form (dev, ρ) and $(gdev, g\rho g^{-1})$ for some $g \in \mathrm{PSL}_2\mathbb{C}$. Since the centralizer of a quasi-Fuchsian representation is trivial, we have that $\rho = g\rho g^{-1}$ implies $g = 1$, hence $(dev, \rho) = (gdev, g\rho g^{-1})$. □

As a result, the decomposition of S depends only on the structure $\sigma = \{(dev, \rho)\}$ and not on the choice of particular representatives. In particular many combinatorial properties of the geometric decomposition (such as the number and type of components, the adjacency pattern,...) are well defined; these invariants are the main focus of the following section. Also notice that, as already observed by [16], the key feature of quasi-Fuchsian structures is that the pieces of the induced decomposition carry (possibly branched) geometric structure, namely hyperbolic outside the real curve and real projective on the real curve, as established by the following result.

Lemma 2.1.7. *If* S *is endowed with a BPS* σ*, then* $\sigma^{\pm}$ *is a finite union of subsurfaces carrying complete hyperbolic metrics with cone points of angle* $2\pi(k + 1)$ *(*$k \in \mathbb{N}$*) corresponding to branch points of order* k *of the BPS, and* $\sigma^{\mathbb{R}}$ *is a finite 1-dimensional CW-complex on* S*; moreover if branch points are not on the real curve, then* $\sigma^{\mathbb{R}}$ *is a finite union of simple closed curves with a (*$\mathrm{PSL}_2\mathbb{R}, \mathbb{RP}^1$*)-structure.*

This is proved in [8, §3], together with more detailed results about the geometric properties of the components of this decomposition. Moreover this motivates the following terminology.

Definition 2.1.8. If S is endowed with a quasi-Fuchsian BPS σ, a connected component C of $\sigma \setminus \sigma^{\mathbb{R}}$ will be called a geometric component of the decomposition; a connected component C of $\sigma^{\pm}$ will be called a positive/negative component. A connected component of $\sigma^{\mathbb{R}}$ will be called a real component.

Notice that the components of the real curve can be canonically oriented by declaring that they have positive regions on the left and negative regions on the right. Some examples are in order.

Example 2.1.9. A hyperbolic structure on S is an example of an unbranched projective structure with Fuchsian holonomy. Any developing map is a diffeomorphism with the upper-half plane $\mathcal{H}^+$. The induced decomposition is $\sigma^+ = S, \sigma^- = \varnothing = \sigma^{\mathbb{R}}$. Hence there is only one geometric component, which is the whole surface.

Example 2.1.10. If we graft a hyperbolic surface along a simple closed geodesic we obtain an example of an unbranched projective structure with Fuchsian holonomy with surjective and non injective developing map to $\mathbb{CP}^1$. As anticipated in 1.4.3 there will be a negative geometric annulus bounded by two essential simple closed real curves and two or one positive geometric components, depending on the fact that the geodesic we use is separating or not.

The main result in [16] claims that every unbranched structure with quasi-Fuchsian holonomy arises via a multigrafting of the uniformizing hyperbolic structure; one of the key observations is the fact that geometric components of an unbranched structure can not be simply connected, i.e. they can not be disks. This completely fails for branched structures as the following easy example shows.

Example 2.1.11. If we bubble a hyperbolic surface along a simple arc we obtain an example of a branched projective structure with Fuchsian holonomy with surjective and non injective developing map to $\mathbb{CP}^1$. There will be a negative geometric disk bounded by a contractible simple closed real curve and one positive geometric component containing the two branch points. The same type of geometric decomposition is obtained by grafting a simple closed geodesic on a hyperbolic surface and then bubbling along an arc crossing transversely the grafting region.

The location of branch points with respect to the geometric decomposition is of course something we want to care about in the following, therefore we introduce the following definitions.

Definition 2.1.12. Let $\sigma \in \mathcal{M}_\rho$. A branch point of σ is said to be geometric (respectively real) if it belongs to $\sigma^{\pm}$ (respectively to $\sigma^{\mathbb{R}}$). The structure is said to be **geometrically branched** if all its branch points are geometric; it is said to be **really branched** if it has some real branch point, and **purely really branched** if all branch points are real. We will denote by $\mathcal{M}_{k,\rho}^{\mathbb{R}}$ the subspace of really branched structures of $\mathcal{M}_{k,\rho}$.

The space $\mathcal{M}^{\mathbb{R}}_{k,\rho}$ could be naturally decomposed by looking at how many branch points actually are on the real curve and how they are grouped. The minimal stratum is given by structures with a single branch point of order k on the real curve, which has real dimension 1, and the maximal stratum is given by structures with only one simple branch point on the real curve, which has real codimension 1 in $\mathcal{M}_{k,\rho}$. Up to a very small movement of branch points, we can always assume that the branch points do not belong to the real curve $\sigma^{\mathbb{R}}$; more precisely, $\mathcal{M}^{\mathbb{R}}_{k,\rho}$ has real codimension 1 inside $\mathcal{M}_{k,\rho}$. We will come back to the study of certain structures with points on the real curve in Section 2.4.2, and now we focus on geometrically branched structures. Under the hypothesis that no branch point belongs to the real curve, some index formulae are available, which give a relation between the geometry and the topology of the pieces of the geometric decomposition. This approach was already exploited by Goldman in [16] for unbranched structures and extended to the general case in [8]. We now introduce the terminology to state the formula; see [8, §3-4] for more details.

Definition 2.1.13. Let σ be a geometrically branched BPS and l be a real component on it. Let $p \in \Lambda_\rho$ be a fix point of $\rho(l)$ and $\widetilde{l}$ is any lift of l. The **index** of the induced real projective structure on l is the integer $I(l) = \#\{dev^{-1}_{|_{\widetilde{l}}}(p)\}$.

The index of a real component can be thought as a degree of the restriction of the developing map to it, as a map with values in the limit set of ρ, and it can a priori assume any value. However if $\rho(l)$ is trivial then the index must be strictly positive: this follows by the classification of $\mathbb{RP}^1$-structures on S^1 given in [8, Proposition 3.2], which we recall for future reference.

Lemma 2.1.14. *Two unbranched $\mathbb{RP}^1$-structures on an oriented circle with non elliptic holonomy are isomorphic if and only if they have the same index and conjugated holonomy. The only case which does not occur is the case of index 0 and trivial holonomy.*

Definition 2.1.15. For a quasi-Fuchsian representation ρ let E_ρ be the induced flat $\mathbb{RP}^1$-bundle on S. For any subsurface $i : C \hookrightarrow S$ we denote by ρ_C the restriction of ρ to $i_*\pi_1(C)$. For any component $l \subset \partial C$ we define a section $s_\rho : l \to E_\rho|_l$ by choosing the flat section passing through a fixed point of $\rho(l)$. Then the **Euler class** eu of ρ_C is defined to be the Euler class of the bundle $E_{\rho_C} = E_\rho|_C$ with respect to this choice of boundary sections.

Finally we say that a subsurface $C \subset S$ is incompressible if the inclusion is injective on fundamental groups; equivalently if all the boundary curves are essential (i.e. not nullhomotopic) in S. The following index formulae hold.

Theorem 2.1.16. *([8, pp. 4.1-5]) Let ρ be a quasi-Fuchsian representation and let $\sigma \in \mathcal{M}_\rho$ be geometrically branched. Let $C \subset \sigma^\pm$ be a geometric component containing k_C branch points (counted with multiplicity) and with $\partial C = \{l_1, \dots, l_n\}$. Then*

$$\pm eu(\rho_C) = \chi(C) + k_C - \sum_{i=1}^{n} I(l_i)$$

Moreover if C is incompressible (e.g. $C = S$) then $eu(\rho_C) = \chi(C)$.

Under the same hypothesis of 2.1.16 the following can be deduced

Corollary 2.1.17. *If $k^{\pm}$ denotes the number of positive/negative branch points of σ, then $2\chi(\sigma^-) = k^+ - k^-$.*

In particular in quasi-Fuchsian holonomy there is always an even number of branch points, so that $\mathcal{M}_{2k+1,\rho}$ are all empty. On the other hand the spaces $\mathcal{M}_{2k,\rho}$ admit a decomposition into pieces defined by the combinatorial properties of the geometric decomposition induced by the structures they contain, as we now explain.

Definition 2.1.18. Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched and let $S = \sigma^{\pm} \cup \sigma^{\mathbb{R}}$ be the induced geometric decomposition. The **combinatorics** of σ is defined to by a finite labelled 1-dimensional CW-complex $\mathfrak{c}_\sigma$ with a vertex for any geometric component and such that two vertices are joined by an edge if the corresponding components are adjacent along a real curve. Every vertex is labelled by the sign, the Euler characteristic and the number of branch points of the corresponding geometric component, and every edge is labelled by the index of the corresponding real curve.

For instance the following picture shows the combinatorics for the geometric decomposition of Picture 2.2.

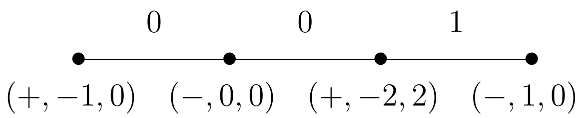

Figure 2.3: Combinatorics for the geometric decomposition of Picture 2.2.

In this complex there are no loops, but bigons may occur, so $\mathfrak{c}_\sigma$ is not a simplicial in general. The information encoded in $\mathfrak{c}_\sigma$ is enough to recover the Euler numbers of every component, by the above index formula 2.1.16; it is also enough to determine the topology of the pieces that occur in the geometric decomposition and how to glue them to reconstruct S. On the other hand not all labelled graphs are allowed: for instance the vertex labels must sum to the total Euler characteristic and the total number of branch points of S. The following is immediate from the definitions.

Lemma 2.1.19. *Let $\sigma \in \mathcal{M}_{k,\rho}$. Moving a branch point inside its own component does not change the geometric decomposition (as a collection of subsurfaces of the surface S); in particular it does not change the associated combinatorics $\mathfrak{c}_\sigma$.*

Definition 2.1.20. Let $\mathfrak{c}$ be a finite labelled graph of the type occurring in 2.1.18. We define the subspace of structures having a fixed combinatorics equal to $\mathfrak{c}$ as $\mathcal{M}^{\mathfrak{c}}_{k,\rho} = \{\sigma \in \mathcal{M}_{k,\rho} \mid \mathfrak{c}_\sigma = \mathfrak{c}\}$.

By the above observations we have that $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is an open smooth complex submanifold of $\mathcal{M}_{k,\rho}$, which is possibly disconnected (see 2.1.22 below for more details). The complement of the union of all the subspaces of the form $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is given by the subspace of really branched structures, which we have denoted by $\mathcal{M}^{\mathbb{R}}_{k,\rho}$.

Definition 2.1.21. We will refer to the decomposition $\mathcal{M}_{k,\rho} = \mathcal{M}^{\mathbb{R}}_{k,\rho} \cup \bigcup_{\mathfrak{c}} \mathcal{M}^{\mathfrak{c}}_{k,\rho}$ of the holonomy fibre as the **real decomposition** of $\mathcal{M}_{k,\rho}$; any connected component of $\mathcal{M}_{k,\rho} \setminus \mathcal{M}^{\mathbb{R}}_{k,\rho}$ will be called a piece or cell of the real decomposition of $\mathcal{M}_{k,\rho}$.

We conclude with the following remark which shows that spaces with fixed combinatorics $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ can fail to be connected in a quite drastic way.

Remark 2.1.22. If the information encoded in a combinatorics $\mathfrak{c}$ is enough to reconstruct the topology of the pieces of the geometric decomposition and the gluing pattern, it is not enough to determine them *as subsurfaces* of the marking surface S. Indeed, the collection Θ_σ of the isotopy classes of the essential real curves of σ is another natural invariant of a BPS which takes values in the set of isotopy classes of multicurves[1] of S, which we denote by $\mathcal{MC}(S)$. This is well-defined for geometrically branched structures, and is invariant under small local deformations of σ by 2.1.19. Therefore we obtain a map

$$\mathcal{M}_{k,\rho} \setminus \mathcal{M}^{\mathbb{R}}_{k,\rho} \to \mathcal{MC}(S), \sigma \mapsto \Theta_\sigma$$

which is constant on the pieces of the real decomposition; in particular it can be used to distinguish connected components of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$. In the unbranched case combinatorics have almost trivial labelling, $\mathcal{M}^{\mathbb{R}}_{0,\rho}$ is empty, $\mathcal{M}_{0,\rho}$ is discrete, and Goldman's result (see [16]) implies that this map is bijective; on the other hand we can construct unbranched structures with the same combinatorics. For instance we can simply graft the uniformizing structure σ_ρ along non isotopic non separating simple closed geodesics γ_1, γ_2 to obtain two unbranched projective structures σ_1, σ_2 with the same combinatorics $\mathfrak{c}$ (a bigon with the same labellings in both cases) but such that $\Theta_{\sigma_1} = \gamma_1 \neq \gamma_2 = \Theta_{\sigma_2}$. Bubbling these structures in a suitable way (for instance along an arc which simply crosses one of the real curves and is disjoint from the other one) gives infinitely many examples of structures with the same combinatorics $\mathfrak{c}$ but different sets of isotopy classes of real curves for any even number of branch points; such structures belong to different components of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$.

2.2 Locating branch points

As shown above, BPSs generally can have geometric disks (hence compressible components) in their geometric decomposition: we can just perform a bubbling on a hyperbolic surface to see them appear. We now show how to control the presence of compressible geometric components and to use them to locate branch points with respect to the geometric decomposition. Throughout this section, we only care about geometrically branched structures, i.e. those with no points on the real curve. We begin by noticing that even if the structure has branch points, nevertheless unbranched components are quite well behaved. In the following $\sigma \in \mathcal{M}_{k,\rho}$ will denote a quasi-Fuchsian BPS with total branching order k on a surface S.

Lemma 2.2.1. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. If $C \subset \sigma^\pm$ is unbranched component then either it is a disk or it is incompressible.*

Proof. We know that unbranched disks can occur. If C is not a disk and is not branched, then it carries a complete hyperbolic structure such that the index of each boundary component is zero; by 2.1.14 we know that it can not have trivial

[1] Recall that a multicurve is a finite collection of disjoint essential simple closed curves.

holonomy. But quasi-Fuchsian representations are in particular injective, hence this implies that each boundary component must be essential in the surface S, hence C is incompressible. □

We are now ready to see the first manifestation of the asymmetry between positive and negative regions hinted at before, and which is a consequence of the special role played by Ω_ρ^+ in the definition 2.1.2 of quasi-Fuchsian representation.

Corollary 2.2.2. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. If C is an unbranched and negative component, then either C is a disk or C is an incompressible annulus.*

Proof. By 2.2.1 if C is not a disk then C is incompressible, so that we may apply the index formula and get $-\chi(C) = -eu(\rho_C) = \chi(C) + k_C - \sum_{l \subset \partial C} I(l) = \chi(C)$, hence $\chi(C) = 0$. □

Notice that the same strategy gives a useless identity in the case of a positive component. In an unbranched structure all geometric component are non simply connected ([16]) and carry complete hyperbolic metrics, hence all the real components have index 0 by the following easy dichotomy.

Lemma 2.2.3. *Let σ be a complete hyperbolic structure on a surface S. Then $\sigma \cong \mathbb{H}^2/\Gamma$ for some Fuchsian group Γ. Moreover we have the following dichotomy*

1. *if $\Gamma = 1$, then S is a disk and the ideal boundary of σ has index 1 and trivial holonomy*
2. *if $\Gamma \neq 1$ then S is not a disk and every component of the ideal boundary has index 0 and hyperbolic holonomy.*

Proof. For complete Riemannian structures the developing map is a covering map, and $\mathbb{H}^2$ is simply connected, so that the developing map is actually an isometry. If $\Gamma \neq 1$, let $\gamma \in \Gamma \setminus \{1\}$; then the covering $\mathbb{H}^2 \to \mathbb{H}^2/\Gamma$ factors through $\mathbb{H}^2/ < \gamma >$. Each real curve of this structure develops to a fundamental domain of γ, which does not contain its limits points, hence it has index 0. □

In a BPS the appearance of disks (or more generally of real components of positive index) can be used to locate branched points.

Lemma 2.2.4. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let l be a real component and C, C' be the components of $\sigma^\pm$ which are adjacent along l. If $I(l) \geq 1$ then*

1. *at most one of C, C' is a disk*
2. *any non disk component is branched*
3. *at least one of C, C' is branched*

Proof. To prove 1) observe that C, C' can not both be disks, otherwise we would get an embedded sphere in S (indeed the same is true even if $I(l) = 0$). To get 2) observe that being unbranched and having ideal boundary with positive index implies indeed being a disk with boundary index 1, by Lemma 2.2.3. Then 3) follows by putting 1 and 2 together. □

Corollary 2.2.5. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. The (unique) component adjacent to a disk of the geometric decomposition of σ is branched.*

Proof. The boundary of a disk has always strictly positive index. Since S is not a sphere, the adjacent component can not be a disk, therefore it is branched. □

Example 2.2.6. It may happen that none of the two components adjacent to a real component with positive index is a disk. For instance graft the uniformizing structure σ_ρ along a geodesic γ, and then perform a bubbling on a small simple arc which has one endpoint inside the negative component of $Gr(\sigma_\rho, \gamma)$, crosses one of the real components, and then reaches a point in the adjacent positive component; see Picture 2.4, left side. The resulting structure will have an essential real component of index 1 such that both adjacent components are non disks and are both branched.

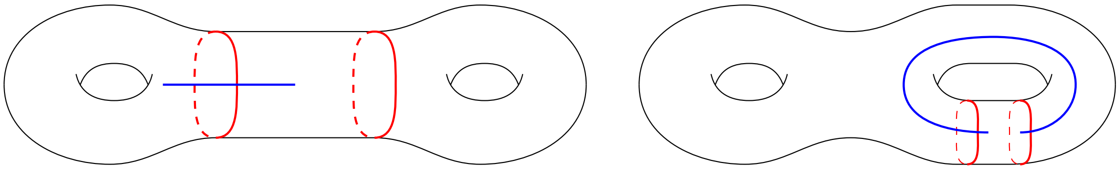

Figure 2.4: Examples for 2.2.6 and 2.2.7.

Example 2.2.7. There are structures with negative components with essential boundary with index 0 which are nevertheless branched; the easiest example is obtained as follows (see Picture 2.4, right side). Graft a non separating geodesic γ on the uniformizing structure σ_ρ, then bubble $Gr(\sigma_\rho, \gamma)$ along an arc with endpoints inside the negative annulus but which is not itself contained inside the negative annulus. The resulting structure has one real component, a positive unbranched incompressible component and a negative incompressible component containing both branch points.

From the above results, in particular we obtain a bound on the number of branch points contained inside a disk of the geometric decomposition of a quasi-Fuchsian structure $\sigma \in \mathcal{M}_{k,\rho}$.

Proposition 2.2.8. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let $D \subset \sigma^\pm$ be a disk with branching order k_D on a quasi-Fuchsian BPS σ with total branching order $k \geq 2$. Then $k_D \leq k - 2$.*

Proof. By 2.2.5 we already know that a disk can not contain all the branching. So we assume by contradiction that it has branching order $k_D = k - 1$. The boundary of D is a real component l of index $I(l) = k_D + 1 = k$. Let C be the component adjacent to D; then we know it is branched by 2.2.5, so $k_C \geq 1$. Indeed $k_D = k - 1$ implies that $k_C = 1$. The boundary of C a priori could contain also m more non essential boundary components and n essential ones. Notice that all components of $\sigma^\pm$ different from C, D are unbranched, simply because $C \cup D$ contains all the branching.

Therefore if $l' \neq l$ is a non essential component of ∂C, then $\rho(l') = id$ hence $I(l') \geq 1$ by 2.1.14, and then by 2.2.4 the geometric component after it must be an unbranched disk D' and l' must have index 1. Let $l, l'_1, \ldots, l'_m$ be the non essential components of ∂C, $D, D'_1, \ldots, D'_m$ the corresponding disks; then $I(l) = k$ but $I(l'_i) = 1, k_{D'_i} = 0$ for $i = 1, \ldots, m$. On the other hand, if l'' is an essential boundary component, then the geometric component after it is a non simply connected complete hyperbolic surface, hence $I(l'') = 0$.

Now observe that the subsurface $E = C \cup D \cup D'_1 \cup \cdots \cup D'_m$ has essential boundary by construction, hence it is incompressible. Therefore the index formula 2.1.16 gives us that

$$eu\left(\rho_{|_{\pi_1(E)}}\right) = \chi(E) = \chi(C) + \chi(D) + \sum_{i=1}^{m} \chi(D'_i) = \chi(C) + 1 + m$$

On the other hand, since disks have trivial Euler class, we obtain, again by 2.1.16, that

$$eu\left(\rho_{|_{\pi_1(E)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) + eu\left(\rho_{|_{\pi_1(D)}}\right) + \sum_{i=1}^{m} eu\left(\rho_{|_{\pi_1(D'_i)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) =$$

$$= \pm\left(\chi(C) + k_C - I(l) - \sum_{i=1}^{m} I(l'_i) - \sum_{j=1}^{m} I(l''_j)\right) = \pm(\chi(C) + 1 - k - m)$$

where the sign depends on the sign of C (hence of that of D). We are now going to compare the two expressions for the Euler class of E. If $C \subset \sigma^+$ then we get $2m + k = 0$ which is absurd since $m \geq 0, k \geq 2$. If $C \subset \sigma^-$ then we get $2\chi(C) = k - 2 \geq 0$. But C can not be a disk, hence $\chi(C) = 0$, i.e. C is an annulus. Its boundary consists of l and another curve l' homotopic to it; so l' is non essential too, hence of positive index. The component adjacent to l' can not be a disk, otherwise S would have genus $g = 0$, hence it must be branched; but by construction all branch points live in $C \cup D$, so we have a contradiction. □

Notice that so far C could be either positive or negative. Indeed, by performing suitable bubbling, we can find structures with either positive or negative disks, either branched or not. We recall the following useful lemma, which was proved in [8, Lemma 10.3] for the positive part (but the proof is exactly the same for the negative part too).

Lemma 2.2.9. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. If all branch points live in σ^+ and $C \subset \sigma^+$ is a branched component with n adjacent disks, then $k_C = 2n$. If all branch points live in σ^- and $C \subset \sigma^-$ is a branched component then $k_C = -2\chi(C)$.*

Proof. Suppose all branch points live in the positive part or in the negative part, and let C be a branched component. The hypothesis implies that all components adjacent to C are unbranched, therefore by 2.2.1 and 2.2.3 a boundary real component of C has index 0 if it is essential or index 1 if it is non essential, and in the second case it bounds a disk. Let $l_1, \ldots, l_n$ be the non essential boundary components of C and let $D_1, \ldots, D_n$ be the adjacent disks. We introduce the subsurface

$E = C \cup D_1 \cup \cdots \cup D_n$, which is clearly incompressible. By 2.1.16 and the remark that disks have trivial Euler class we obtain

$$\chi(C) + n = \chi(E) = eu\left(\rho_{|_{\pi_1(E)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) + \sum_{i=1}^{n} eu\left(\rho_{|_{\pi_1(D_i)}}\right) =$$

$$\pm\left(\chi(C) + k_C - \sum_{i=1}^{n} I(l_i)\right) = \pm\left(\chi(C) + k_C - n\right)$$

from which the statement follows. $\square$

The previous results were concerned with the localisation of branch points with respect to different components, i.e. they clarify which components have to be branched. The following results aim to specify the location of a branch point inside a branched component. We start with the following definition.

Definition 2.2.10. Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let $C \subset \sigma^{\pm}$ be a geometric component and $l \subset \partial$ a real component in its boundary. We call the **peripheral geodesic** of l in C the unique geodesic representative γ in the free homotopy class of l. The **end** of l in C is the connected component E_l of $C \setminus \gamma$ which has l in its boundary.

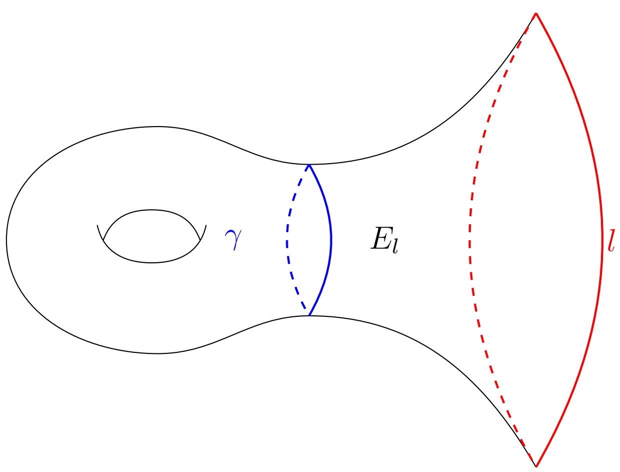

Figure 2.5: A geometric component with one end.

In Picture 1.2 we have shown the geometric decomposition of a grafting annulus for Fuchsian structure: it is the union of a negative annulus (entirely made of the two ends relative to its real boundaries, joined along the peripheral geodesic) and two positive ends coming from the adjacent positive component(s). We insist on the fact that in our terminology the grafting annulus contains also these two positive ends, and properly contains the negative annulus.

It is shown in [8, §3.3] that ends are embedded open annuli and ends associated to different real components are disjoint. However peripheral geodesics are not necessarily embedded, i.e. the closure of an end is just an immersed annulus.

Lemma 2.2.11. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let $C \subset \sigma^{\pm}$ be a geometric component, $l \subset \sigma^{\mathbb{R}} \cap \partial C$ a real component in its boundary and γ the peripheral geodesic of l in C. If γ self-intersects at a point p, then p must be a branch point.*

Proof. A geodesic always makes an angle of at least π at any of its points; therefore if it self-intersects at p then the total angle at p must be strictly larger than 2π. □

Notice that the converse does not hold, i.e. an embedded peripheral geodesic can go through branch points. For instance let β a bubbleable geodesic arc on a hyperbolic surface σ_ρ; then $Bub(\sigma_\rho, \beta)$ has a trivial real component whose peripheral geodesic coincides with the boundary of the bubble, hence is embedded and goes through both branch points.

Lemma 2.2.12. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let $C \subset \sigma^{\pm}$ be a geometric component, $l \subset \sigma^{\mathbb{R}} \cap \partial C$ a real component in ∂C and γ the peripheral geodesic of l in C. Assume that $I(l) > 0$ and that C is not a disk. Then C is branched and at least one of its branch points lives in the closure of the end of l in C.*

Proof. By 2.2.4 we already know that C must be branched. Assume by contradiction that the closure of the end of C relative to l does not contain any branch point. Then it is an embedded annulus homeomorphic to $S^1 \times [0, 1[$ with an unbranched hyperbolic structure with one geodesic boundary and one real boundary of positive index, which is absurd: namely if the holonomy around the annulus is loxodromic then the real boundary can not have positive index, and conversely if the holonomy is trivial the geometric boundary can not be geodesic without the introduction of branch points. □

Of course branch points can live in the interior of an end: take one of the structures of 2.2.6 such that the bubbling arc is entirely contained in the grafting annulus. On the other hand notice that end closures are not necessarily disjoint: for instance let γ be a geodesic on σ_ρ and consider a geodesic bubbleable arc β on $Gr(\sigma_\rho, \gamma)$ with one endpoint x on the boundary of the grafting annulus A_γ and entering the convex core of the adjacent component. Then $\sigma = Bub(Gr(\sigma_\rho, \gamma), \beta)$ has a positive component C containing both branch points; ∂C contains two real boundaries (there is one more if γ is non separating) whose peripheral geodesics come respectively from the geodesic boundary of A_γ and from β, and they intersect at x, even if each end closure is an embedded closed annulus.

2.3 Fuchsian models and the divisor map

Let S be a closed, connected and oriented surface of genus $g \geq 2$, and let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation. In [8, § 7] the authors introduce a standard way to associate to any component C of a geometric decomposition a complete hyperbolic surface C_{Fuchs} together with a branched covering map $D_C : C \to C_{Fuchs}$ which is a local isometry outside branch points. We exploit this construction to define a natural map on the geometric part of the moduli space $\mathcal{M}_{k,\rho}$, i.e. on $\mathcal{M}_{k,\rho} \setminus \mathcal{M}_{k,\rho}^{\mathbb{R}}$.

Definition 2.3.1. Let $\sigma \in \mathcal{M}_{k,\rho}$ be a geometrically branched structure and $C \subset \sigma^{\pm}$ be any geometric component. Choose a lift $\widetilde{C}$ of C to $\widetilde{S}$. Let G_C be the stabiliser of $\widetilde{C}$ inside $\pi_1(S)$ in the action by deck transformations. We define the **Fuchsian model** of C to be the complete unbranched hyperbolic surface $C_{Fuchs} = \Omega_\rho^{\pm}/\rho(G_C)$. The restriction of the developing map $dev_{|_{\widetilde{C}}} : \widetilde{C} \to \Omega_\rho^{\pm}$ is $(G_C, \rho(G_C))$-equivariant, hence descends to a map $D_C : C \to C_{Fuchs}$ which is a branched cover, locally isometric outside branch points.

As observed in [8, Lemma 7.5], if C is incompressible then $G_C = \pi_1(C)$ and thus the component is diffeomorphic to its own Fuchsian model. However, notice that D_C does not extend continuously to the real boundaries, not even in the incompressible case.

Lemma 2.3.2. *Let $\sigma_1, \sigma_2 \in \mathcal{M}_{k,\rho}$ be geometrically branched. If $\mathfrak{c}_{\sigma_1} = \mathfrak{c}_{\sigma_2}$, then there is a bijective correspondence between the pieces of the geometric decomposition of σ_1 and the one of σ_2 such that corresponding geometric components have diffeomorphic Fuchsian model.*

Proof. It was already observed that the combinatorics of a structure completely determines the topological type of the pieces occurring in the geometric decomposition. If there are more components of the same topological type (e.g. if there are several grafting annuli), then we can use the adjacency pattern of the combinatorics to construct a bijection between the two sets of geometric components. In particular, corresponding components induce isomorphic subgroups in $\pi_1(S)$. Since quasi-Fuchsian representations are injective and since the Fuchsian model of a subsurface $i : C \hookrightarrow S$ depends only on its sign and on the restriction of the representation to $i_*(\pi_1(C))$, the statement follows. □

The following notion is thus well defined.

Definition 2.3.3. The Fuchsian model of a combinatorics $\mathfrak{c}$ is the space

$$S_{\mathfrak{c}} = \prod_{C \in \mathfrak{c}^0} Sym^{k_C}\left(C_{Fuchs}\right)$$

where the product is taken over the vertices of $\mathfrak{c}$ (i.e. geometric components) and k_C is the number of branch points contained in C as usual. We also define the divisor map for $\mathfrak{c}$: let us denote by $div(\sigma)$ the branching divisor of σ and define

$$\mathcal{D}_{\mathfrak{c}} : \mathcal{M}_{k,\rho}^{\mathfrak{c}} \to S_{\mathfrak{c}}, \sigma \mapsto \prod_{C \in \mathfrak{c}^0} D_C(div(\sigma) \cap C)$$

by sending a structure to the image of its branching divisor via the maps D_C.

It is understood that the symmetric power of order 0 of a space is a point, so that unbranched components do not actually show up in the above product, and there are at most k of them. Since symmetric products of 1-dimensional complex manifolds are smooth complex manifolds, $S_{\mathfrak{c}}$ carries actually the structure of a connected complex manifold of dimension k, the same dimension of $\mathcal{M}_{k,\rho}$.

Remark 2.3.4. As already observed in 2.1.22, the space $\mathcal{M}_{k,\rho}^{\mathfrak{c}}$ is not connected in general, as there might be structures in $\mathcal{M}_{k,\rho}^{\mathfrak{c}}$ with non isotopic collections of real curves. When this happens, Fuchsian models of corresponding components are diffeomorphic (by 2.3.2) but not isometric in general. For a concrete example consider the construction in 2.1.22: pick two non isotopic non separating simple closed geodesics γ_1, γ_2 on the uniformizing structure σ_ρ along, graft along them and the bubble along an arc crossing once a real curve and disjoint from the other. The resulting structures have corresponding negative annuli A_1, A_2 whose Fuchsian models are respectively the Hopf annuli $\mathbb{H}^2/\rho(\gamma_1)$ and $\mathbb{H}^2/\rho(\gamma_2)$ which are very likely to be non isometric (i.e. non biholomorphic), since ρ is injective. However this clearly does not happen for two structures belonging to the same piece of the real decomposition, since their geometric decompositions consist of the same subsurfaces of S. Nevertheless even when Fuchsian models are isometric the developing maps are different, hence they induce different maps $D_{C,1}, D_{C,2} : C \to C_{Fuchs}$ for any geometric component $C \subset S$. As a consequence of these observation, we see that the Fuchsian model $S_{\mathfrak{c}}$ as a complex manifold actually depends on the choice of a piece of the real decomposition, whereas as a real manifold it only depends on the combinatorics of the piece (by 2.3.2). Since here we are only interested in the topology of $\mathcal{D}_{\mathfrak{c}} : \mathcal{M}_{k,\rho}^{\mathfrak{c}} \to S_{\mathfrak{c}}$, we will consider it as a topological space associated only to the combinatorics, to avoid more complicated discussion and notations.

With the caveat of the above remark, the map $\mathcal{D}_{\mathfrak{c}}$ defines a continuously-valued function on the moduli space outside the locus $\mathcal{M}_{k,\rho}^{\mathbb{R}}$ of really branched structures. Since $\mathcal{M}_{k,\rho}$ is connected, a discrete invariant would not be very interesting. Let us now discuss some examples.

Example 2.3.5. Let $\sigma \in \mathcal{M}_{2,\rho}$ be obtained as a standard bubbling of the uniformizing structure σ_ρ. The associated combinatorics $\mathfrak{c}$ has a negative disk D, a non essential real curve and index 1 and a positive component C with both branch points. In the above notations $G_C = \pi_1(S)$ and $G_D = 1$, so that $C_{Fuchs} = \sigma_\rho$ and $D_{Fuchs} = \mathcal{H}^-$. The divisor map for this combinatorics is a map $\mathcal{D}_{\mathfrak{c}} : \mathcal{M}_{2,\rho}^{\mathfrak{c}} \to Sym^2(\sigma_\rho)$. Tipically the target space has non trivial topology in every degree of homotopy and homology; for instance in genus $g = 2$ it is canonically identified with a blow-up of $Pic^2(\sigma_\rho)$, the moduli space of line bundles of degree 2 over the Riemann surface underlying σ_ρ; its exceptional divisor is given by the canonical locus and provides a non homotopically trivial sphere.

Example 2.3.6. Let $\sigma \in \mathcal{M}_{2,\rho}$ be obtained by grafting the uniformizing structure σ_ρ and then bubbling it along an arc with one endpoint of different signs (as in the left side of Picture 2.4). Then σ has the property that every component of the geometric decomposition contains at most one branch point. The divisor map takes value in the product of two complete unbranched hyperbolic surfaces, which is an aspherical space.

We now proceed to investigate the topological structure of the divisor map.

Definition 2.3.7. In the above notations, let Δ_C be the generalised diagonal of $Sym^{k_C}(C_{Fuchs})$, i.e. the closed subspace of points with at least two coinciding coordinates. We then define the simple part of the Fuchsian model

$$\mathcal{S}S_{\mathfrak{c}} = \prod_{C \in \mathfrak{c}^0} Sym^{k_C} \left(C_{Fuchs} \right) \setminus \Delta_C$$

and the subspace of simply developed structures $\mathcal{SM}^{\mathfrak{c}}_{k,\rho} = \mathcal{D}^{-1}_{\mathfrak{c}} \left(\mathcal{S}S_{\mathfrak{c}} \right) \subset \mathcal{M}^{\mathfrak{c}}_{k,\rho}$.

Of course if a combinatorics is such that all branched components have exactly one simple branch point, then $\Delta_C = \varnothing$, so that $\mathcal{S}S_{\mathfrak{c}} = S_{\mathfrak{c}}$ and $\mathcal{SM}^{\mathfrak{c}}_{k,\rho} = \mathcal{M}^{\mathfrak{c}}_{k,\rho}$. More generally we have that:

Lemma 2.3.8. *$\mathcal{S}S_{\mathfrak{c}}$ is an open dense connected submanifold of $S_{\mathfrak{c}}$ and $\mathcal{SM}^{\mathfrak{c}}_{k,\rho}$ is an open dense connected submanifold of $\mathcal{M}^{\mathfrak{c}}_{(1,\dots,1),\rho}$.*

Proof. The first part is classical, and follows from the fact that the generalised diagonals are defined by complex equations, so they have complex codimension at lest 1. Let σ be a BPS with a non simple branch point. Then the divisor map sends σ inside Δ_C by construction. Therefore a simply developed structure belongs necessarily to the principal stratum. On the other hand if a structure is in the principal stratum but is not simply developed, then there are at least two points inside the same component which develop to the same point in the Fuchsian model of that component. If we want this condition to be preserved, then once we choose how to perform a local deformation of one of them the deformation of the other is uniquely determined. In other words the space of non simply developed structures has complex dimension at least 1. □

Example 2.3.9. If β is a bubbleable arc on a hyperbolic surface σ_ρ, then $\sigma = Bub(\sigma_\rho, \beta)$ is simply developed, because the developing map for σ_ρ is injective. More precisely if C is the positive branched component of σ, then we already know that $C_{Fuchs} = \sigma_\rho$ and actually D_C maps the boundary of the bubble to β itself, i.e. D_C is a debubbling map.

Remark 2.3.10. Conversely, there exist simply branched structure which are not simply developed, i.e structures contained in the principal stratum whose image via the divisor map does not land in the simple part of the Fuchsian model. The easiest example is obtained by bubbling a grafting annulus along an arc having endpoints on a couple of twin points on the boundary of the annulus. A more interesting example can be obtained by bubbling a hyperbolic surface and then moving both branch points inside the bubble along suitable arcs, without crossing the real curve (i.e without changing the combinatorics $\mathfrak{c}$ and actually staying inside the same component of $\mathcal{M}^{\mathfrak{c}}_{2,\rho}$). This means that the divisor map provides a tool to prove that a structure can not be obtained as a bubbling of a hyperbolic surface, which is more refined that the invariant given by the combinatorics.

To proceed in the study of the structure of $\mathcal{D}_{\mathfrak{c}}$ we will extensively use the following lemma which is just a reformulation of [8, Lemma 6.1-2]. Recall that an embedded twin pair is what we need to perform a movement of branch points.

Lemma 2.3.11. *Let $\sigma \in \mathcal{M}_{k,\rho}$ be geometrically branched. Let $p \in C$ be a branch point of σ and μ_1, μ_2 be a couple of geodesics starting at σ and contained in C such that $D_C(\mu_1) = D_C(\mu_2)$ is a properly embedded geodesic in C_{Fuchs}. Then μ_1, μ_2 are* *an embedded twin pair. In particular we can always move the branch point p at a distance which is at least the injectivity radius of C_{Fuchs}.*

Theorem 2.3.12. *The restriction of the divisor map to any connected component $\mathcal{X}$ of $\mathcal{SM}^{\mathfrak{c}}_{k,\rho}$ is a covering map $\mathcal{D}_{\mathfrak{c}} : \mathcal{X} \to \mathcal{S}S_{\mathfrak{c}}$.*

Proof. First of all recall that each branched geometric component has its own Fuchsian model and the Fuchsian model $S_{\mathfrak{c}}$ is a direct product of symmetric products thereof, hence we can assume that in $\mathfrak{c}$ there is only one branch component C. This is just to simplify notation and discussion.

Let $\sigma \in \mathcal{X}$ and $div(\sigma) = \sum_{i=1}^k x_i$ and $\widehat{x_i} = D_C(x_i)$. Let $d_{i,j} = d_{C_{Fuchs}}(\widehat{x_i}, \widehat{x_j})$ and let $d^* = \min\{d_{i,j}\}$. For any $0 < \varepsilon < d^*$ we consider the metric balls $B_\varepsilon(x_i) \subset C$ and $B_\varepsilon(\widehat{x_i}) \subset C_{Fuchs}$ and observe that D_C maps the former to the latter. Moreover $\widehat{x_i} \in B_\varepsilon(\widehat{x_j})$ if and only if $i = j$, thus $\prod_{i=1}^k B_\varepsilon(\widehat{x_i})$ projects diffeomorphically to a neighbourhood V_ε of $\mathcal{D}_{\mathfrak{c}}(\sigma)$ in $\mathcal{S}S_{\mathfrak{c}} = Sym^k(C_{Fuchs}) \setminus \Delta_C$. Also notice that moving branch points inside the balls $B_\varepsilon(x_i)$ provides a full parametrisation of a neighbourhood U_ε of σ inside $\mathcal{X}$. We claim that $\mathcal{D}_{\mathfrak{c}} : U_\varepsilon \to V_\varepsilon$ is a diffeomorphism. Injectivity is clear and surjectivity can be proved as follows: for any point $\sum_{i=1}^k y_i \in V_\varepsilon$ choose a geodesic segment μ^i joining $\widehat{x_i}$ to y_i inside $B_\varepsilon(\widehat{x_i})$. By 2.3.11 branch points on σ can be moved along the embedded twin pairs $D_C^{-1}(\mu^i)$; the resulting structure $\sigma' \in U_\varepsilon$ is such that $\mathcal{D}_{\mathfrak{c}}(\sigma') = \sum_{i=1}^k y_i$. This proves that the map is a local diffeomorphism.

We can use the same idea to prove that it is also surjective. Let $\sum_{i=1}^k y_i \in \mathcal{S}S_{\mathfrak{c}}$ be a generic point. We can find a path $\gamma : [0,1] \to \mathcal{S}S_{\mathfrak{c}}$ with $\gamma(0) = \mathcal{D}_{\mathfrak{c}}(\sigma) = \sum_{i=1}^k \widehat{x_i}$ and $\gamma(1) = \sum_{i=1}^k y_i$, and such that the induced paths $\gamma_1, \dots, \gamma_k : [0,1] \to C_{Fuchs}$ such that $\gamma_i(0) = \widehat{x_i}$ and $\gamma_i(1) = y_i$ are piecewise geodesic and disjoint. Then we take a subdivision $\{\gamma^m\}$ of γ such that each piece of the induced subdivision $\{\gamma_i^m\}$ of each γ_i is an embedded geodesic arc. By 2.3.11 we can move branch points on σ along γ^1 to obtain a structure σ_1 and iteratively we can move branch points on σ_{m-1} along γ^m to obtain a structure σ_m. After a finite number of steps we obtain a structure whose branching divisor develops to $\sum_{i=1}^k y_i$, as desired.

Now let again $\sum_{i=1}^k y_i \in \mathcal{S}S_{\mathfrak{c}}$ be a generic point and let $\sigma_j \in \mathcal{X}$ be the structures such that $\mathcal{D}_{\mathfrak{c}}(\sigma_j) = \sum_{i=1}^k y_i$. Let R be the injectivity radius of C_{Fuchs}, $d^* = \min\{d_{C_{Fuchs}}(y_i, y_j)\}$ and $0 < \varepsilon < \min\{R, d^*\}$. Let V_ε be the neighbourhood of $\sum_{i=1}^k y_i$ and $U_{\varepsilon,j}$ be the neighbourhood of σ_j defined as above via the balls $B_\varepsilon(y_i)$. By the above discussion $\mathcal{D}_{\mathfrak{c}} : U_{\varepsilon,j} \to V_\varepsilon$ is a diffeomorphism, hence it is enough to prove that $U_{\varepsilon,j} \cap U_{\varepsilon,l} = \varnothing$ if $j \neq l$. If this is not the case, then for some $j \neq l$ there is a structure σ in the intersection; it is possible to move its branch points along two different collections of geodesic embedded twin pairs μ and ν to obtain σ_j and σ_l respectively. When projecting all these arcs to the Fuchsian model, we see (at least) one piecewise geodesic loop based at some y_i which is made of two geodesic segments each of length less than ε. Since the ε-balls on C_{Fuchs} are uniquely geodesic, this loop must be essential. But then the geodesic representative in its free homotopy class is a simple closed geodesic shorter than the systole of C_{Fuchs}, which is absurd. $\square$

Since generalised diagonals have real codimension at least 2, the inclusion of the simple part of a symmetric product induces an epimorphism on fundamental groups;

since the fundamental group of a symmetric product equals the first homology group of the space, the target space in the above statement is not simply connected.

Corollary 2.3.13. *The restriction to any connected component $\mathcal{X}$ of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is a branched covering map branching over the generalized diagonals of $S_{\mathfrak{c}}$.*

As far as the branching locus is concerned, it can be decomposed according to the partition of the branching divisor $div(\sigma)$ and the one of the developed divisor $\mathcal{D}_{\mathfrak{c}}(\sigma)$; the simply developed structures correspond to the choice of the maximal partition $(1,\dots,1)$ for both divisors. The same techniques developed above can be used to prove that the divisor map is a covering map when suitably restricted to strata of the branching locus with a given couple of partitions. Notice that the map has indeed non trivial branching behaviour around non principal strata.

When the combinatorics is such that it prevents the presence of branching, it is easy to use the above theorem to deduce information about the topology of the corresponding cells of the real decomposition.

Corollary 2.3.14. *If a combinatorics $\mathfrak{c}$ is such that all branched components contain exactly one simple branch point, then the universal cover of every component of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is diffeomorphic to a product of k-copies of $\mathbb{H}^2$. In particular every component of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is aspherical.*

Proof. A structure with combinatorics $\mathfrak{c}$ lies necessarily in the principal stratum. Moreover, as already remarked, for such a combinatorics we have that $\mathcal{S}S_{\mathfrak{c}} = S_{\mathfrak{c}} = \prod_{C\in\mathfrak{c}^0} C_{Fuchs}$ and $\mathcal{S}\mathcal{M}^{\mathfrak{c}}_{k,\rho} = \mathcal{M}^{\mathfrak{c}}_{k,\rho}$. Therefore by 2.3.12 we obtain that each component of $\mathcal{M}^{\mathfrak{c}}_{k,\rho}$ is a cover of a product of complete unbranched hyperbolic surfaces, whose universal cover is $\mathbb{H}^2$. □

The same conclusion holds for structures whose developed divisor is a single branch point of order k, since in this case the divisor map takes value in the pure diagonal of the symmetric product, which is a copy of the Fuchsian model itself. However as soon as there is a branched component with more than one branch point, the Fuchsian model contains non trivial symmetric products, which are highly non aspherical.

Remark 2.3.15. The cells of the real decomposition are not simply connected in general, as shown by the following construction. Let S be a genus $g \geq 2$ surface and γ be a simple closed geodesic on the uniformizing structure σ_ρ. For any $x \in \gamma$ let β_x be a bubbleable geodesic arc orthogonal to γ at x and short enough to be contained in a collar neighbourhood of γ. Then let $\sigma_x = Bub(\sigma_\rho, \beta_x)$. This defines a map $B_\gamma : \gamma \to \mathcal{M}^{\mathfrak{c}}_{2,\rho}$, where $\mathfrak{c}$ is the combinatorics of a standard bubbling. The divisor map $\mathcal{D}_{\mathfrak{c}} : \mathcal{M}^{\mathfrak{c}}_{2,\rho} \to S_{\mathfrak{c}} = Sym^2(\sigma_\rho)$ induces a map $(\mathcal{D}_{\mathfrak{c}})_* : \pi_1(\mathcal{M}^{\mathfrak{c}}_{2,\rho}) \to \pi_1(S_{\mathfrak{c}}) \cong H_1(S,\mathbb{Z})$. Then we have that $(\mathcal{D}_{\mathfrak{c}})_*(B_\gamma) = [\gamma] \in H_1(S,\mathbb{Z})$ by construction. Therefore any non separating geodesic gives rise to a non trivial element in $\pi_1(\mathcal{M}^{\mathfrak{c}}_{2,\rho})$.

2.4 Structures with $k = 2$ branch points

In the simple case in which we have just $k = 2$ branch points we can obtain a very neat description of the combinatorics which are allowed in the geometric decomposition of S. Throughout this section S is a closed, connected and oriented surface endowed with a BPS σ with quasi-Fuchsian holonomy ρ and $k = 2$ branch points. We begin by assuming that the branch points are not on the real curve, to apply the index formulae; in 2.4.2 we will add some considerations about structures with real branch points.

In 2.2.2 we observed that in general an unbranched negative component which is not a disk is automatically an incompressible annulus. Under the hypothesis of having two branch points we obtain a precise statement about branched negative incompressible components.

Lemma 2.4.1. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. Let $C \subset \sigma^-$ be a branched negative incompressible component containing k_C branch points. Then*

1. *either $k_C = 1$, C is an annulus with $\partial C = l \cup l'$ such that $I(l) = 0, I(l) = 1$*
2. *or $k_C = 2$, C is a pant or a once-holed torus and $\forall\ l \subset \partial C$ we have $I(l) = 0$*

Proof. Since C is incompressible we can applying the index formula and we get

$$-\chi(C) = -eu\left(\rho_{|_{\pi_1(C)}}\right) = \chi(C) + k_C - \sum_{l \subset \partial C} I(l) \Rightarrow 2\chi(C) + k_C = \sum I(\gamma) \geq 0$$

and here we look for integer solutions with the constraints that $\chi(C) \leq 0$ (being incompressible, C is not a disk) and $k_C \leq 2$. We see that the only possibilities are the following

1. $k_C = 0, \chi(C) = 0$, so that C is an unbranched annulus (which we discard, since C is assumed to be branched)
2. $k_C = 1, \chi(C) = 0$, so that C is an annulus; we get $\sum I(l) = 1$, which means that one boundary component has index 0 and the other has index 1
3. $k_C = 2, \chi(C) = 0$, so that C is again an annulus and $\sum I(l) = 2$; in particular there is a boundary with positive index and the adjacent component should be branched, but C already contains all the branching (so we do not have this possibility)
4. $k_C = 2, \chi(C) = -1$, and we have $\sum I(l) = 0$, which implies that all boundaries have zero index.

□

Notice that performing a bubbling over a grafting of a hyperbolic surface in a suitable way we can obtain structure realising each of the possibilities allowed by the lemma. On the other hand a description of compressible negative branched components follows from the main theorem below (2.4.5).

To do a similar study for positive branched components we need some preliminary results. A straightforward consequence of 2.2.8 is that disks are always unbranched when we have only $k = 2$ branch points; in particular a real component bounding a geometric disk has index 1. We want to prove an analogous statement for essential real components.

Lemma 2.4.2. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. If a component $C \subset \sigma^{\pm}$ is not a disk and contains a single simple branch point, then the inclusion $i : C \hookrightarrow S$ can not be nullhomotopic (i.e. $i_*(\pi_1(C)) \subset \pi_1(S)$ is not the trivial subgroup).*

Proof. By contradiction assume $i_*(\pi_1(C)) \subset \pi_1(S)$ is trivial. In particular C must have genus 0 and its boundary consists of $m \geq 2$ (it is not a disk) non essential boundary components $l_1, \dots, l_m$ with index $I(l_i) \geq 1$. Since $i_*(\pi_1(C))$ is trivial in $\pi_1(S)$, the flat bundle associated to ρ is trivial on C, hence the Euler class vanishes. Applying the index formula we obtained

$$0 = \pm eu\left(\rho_{|_{\pi_1(C)}}\right) = \chi(C) + k_C - \sum_{i=1}^{m} I(l_i) \leq 2 - m + k_C - m \leq k_C - 2$$

which contrasts with the fact that $k_C = 1$. □

Proposition 2.4.3. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. If $l \subset \sigma^{\mathbb{R}}$ is any real component, then $I(l) \leq 1$.*

Proof. Suppose by contradiction we have a real curve $l_0 \subset S_{\mathbb{R}}$ of index $I(l_0) \geq 2$. We distinguish two cases according to the fact that the real curve l_0 is trivial or not.

In the case l_0 is homotopically trivial, it bounds exactly one subsurface D homeomorphic to a disk one one side and another subsurface S' which is not a disk on the other side. This subsurface D can either be a geometric disk, or it can consist of more than just one single geometric component. In the first case it is unbranched by 2.2.8 hence l_0 has index 1; in the second case the geometric component C of D which has l_0 in its boundary is a non disk component, hence it must be branched; since S' must be branched as well, C contains exactly one branch point, but then 2.4.2 applies and we get a contradiction with the fact that C is contained in a disk, hence its inclusion is homotopically trivial.

Suppose now l is essential. Let us call $C^{\pm}$ the adjacent geometric components. Then $C^{\pm}$ are branched; more precisely $k_{C^{\pm}} = 1$, they are not disks since l_0 is essential and all other components are unbranched, since $C^+ \cup C^-$ contain all the branching. The two components $C^{\pm}$ may have $m \geq 0$ more boundaries in common, let us call them $l_1, \dots, l_m$. Moreover each of them can have more boundary components, either essential or not. Let us focus on C^+; its boundary consists of $l_0, l_1, \dots, l_m$ and possibly of some other non essential components $l'_1, \dots, l'_n$ and some essential ones $l''_1, \dots, l''_p$, for some $n, p \geq 0$. Clearly the non essential components $l'_1, \dots, l'_n$ must bound unbranched disks $D'_1, \dots, D'_n$ (hence they have index 1), and the essential components $l''_1, \dots, l''_p$ must bound unbranched components which are not disks (hence they have index 0 and are essential).
We consider the subsurface $E = C^+ \cup D'_1 \dots D'_n$ and we see that it is incompressible: $l''_1, \dots, l''_p$ are essential by definition, $l_1, \dots, l_m$ are non separating curves in S (C^+ and C^- are adjacent along l_0 in any case), hence they are essential as well, as soon as $m \geq 1$. The only case we need to check is when $m = 0$, but we are currently discussing the case in which l_0 is essential.

Then we apply the index formula and get

$$eu\left(\rho_{|_{\pi_1(E)}}\right) = \chi(E) = \chi(C^+) + \sum_{i=1}^{n} \chi(D'_i) = \chi(C) + n$$

On the other hand, since unbranched disks have trivial Euler class, we obtain

$$eu\left(\rho_{|_{\pi_1(E)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) = \chi(C) + k_C - I(l_0) - \sum_{i=1}^{m} I(l_i) - \sum_{j=1}^{n} I(l'_j) - \sum_{h=1}^{p} I(l''_h)$$

$$= \chi(C) + 1 - I(l_0) - \sum_{i=1}^{m} I(l_i) - n$$

By comparing the two expressions we obtain that

$$2n + I(l_0) + \sum_{i=1}^{m} I(l_i) = 1$$

Now we only have that the left hand side is a sum of non negative integers and that $I(l_0) \geq 2$ by hypothesis, therefore in any case we reach an absurd. □

Now we can deal with positive branched components and prove the following result about them. Notice that a loxodromic element in a quasi-Fuchsian group is necessarily non elliptic, since the representation is injective and surface groups are torsion-free.

Lemma 2.4.4. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. Let $C \subset \sigma^+$ be a branched positive component. Then*

1. *if C is incompressible then $k_C = 1$ and there is a unique boundary curve of index 1, loxodromic holonomy and the component beyond it is branched;*
2. *if C is compressible then $k_C = 2$ and there is a unique boundary curve of index 1, trivial holonomy and the component beyond it is an unbranched disk.*

Proof. If C is incompressible then we apply the index formula and get $k_C = \sum_{l \subset \partial C} I(l)$. Moreover every boundary component is essential, hence by 2.4.3 its index is at most 1. Therefore we have exactly k_C components of index 1 (and possibly some components of index 0). Being essential, they do not bound disks, hence the adjacent components are branched. In particular if $k_C = 2$ then there are two boundaries with index 1 and thus some branched component is adjacent to C; but C already contains all the branching, hence $k_C = 1$ and there is a unique real component of index 1. Since it is essential and we are in quasi-Fuchsian holonomy, the holonomy around the curve will be loxodromic. Of course the component beyond it is branched by 2.2.4.
If C is compressible, then let us say there are $m \geq 1$ non essential boundaries $l_1, \dots, l_m$ (which have index 1 by 2.4.3, since non essential curves have strictly positive index) and $n \geq 0$ essential boundaries $l'_1, \dots, l'_n$, n_0 of which have index 1 (and the others have index 0 by 2.4.3). Then we can cap C with these adjacent negative disks and apply the index formula to the resulting incompressible subsurface E

$$\chi(C) + m = \chi(E) = eu\left(\rho_{|_{\pi_1(E)}}\right) = eu\left(\rho_{|_{\pi_1(C)}}\right) = \chi(C) + k_C - m - n_0$$

$$2m + n_0 = k_C$$

Since $m \geq 1$ bu $k_C \leq 2$, this implies that indeed $k_C = 2, m = 1, n_0 = 0$. □

The above study was focused on a single branched component, but now we go global.

Theorem 2.4.5. *Let S be a closed, connected and oriented surface of genus $g \geq 2$, $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation and $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. Let $k^{\pm}$ denote the number of branch points in $\sigma^{\pm}$*

1. *If $k^+ = 2$ then both branch points live in the same positive component; more precisely there exists a unique negative unbranched disk and the branch points live in the positive component which is adjacent to it.*

2. *If $k^- = 2$ then both branch points live in the same negative component; more precisely there exists a negative component of Euler characteristic -1 containing both branch points. Moreover it has at most one non essential boundary component (with trivial holonomy and index 1), while all essential boundaries have loxodromic holonomy and index 0.*

3. *If $k^+ = k^- = 1$ then the two branched components are adjacent along an essential real component with index 1 and loxodromic holonomy; the negative branched component is an incompressible annulus.*

Moreover in each case all the other positive components are unbranched incompressible and all the other negative components are unbranched incompressible annuli and all the other real curves have index 0.

Proof. 1. We have $2\chi(S^-) = k^+ - k^- = 2$, so $\chi(S^-) = 1$, thus there must be a negative disk D. Let C be the positive component adjacent to D. By 2.2.9 C contains 2 (i.e. all) branch points and indeed there are no other negative disks.

2. We have $2\chi(S^-) = k^+ - k^- = -2$, so $\chi(S^-) = -1$. By 2.2.9 negative components are either unbranched incompressible annuli or components with Euler characteristic -1 and 2 branch points; hence there is exactly one of the latter kind. If it is incompressible, then it has the required boundary behaviour by 2.4.1. If it is a pair of pants and it has one non essential boundary component, then the adjacent component is a disk (because it is unbranched), hence the index is 1. If it had two non essential boundaries, then also the third boundary would be non essential, but then all the components adjacent to the three boundaries must be disk and S would be a sphere, so this case is absurd.

3. Let C be the positive branched component. Since it has only one branch point, by 2.4.4 it is incompressible and has a unique boundary component of index 1 and hyperbolic holonomy. The negative component adjacent along it can not be a disk, hence it is branched, with one branch point. By 2.4.1 it is an incompressible annulus and the other boundary component has index 0. Moreover notice that the only negative disks could appear at the boundary of C, but this is forbidden since it is incompressible.

The rest of the statement follows from the initial discussion: the non branched components can not be disks, hence they are incompressible and with zero index boundary by 2.2.1. The negative ones are annuli by 2.2.2. As a consequence all real curves have zero index, except the boundary of the disk in the case $k^+ = 2$ and the curve separating the branch points in the case $k^+ = 1$, which have index 1. □

This gives a description of negative branched components also in the compressible case, which was still missing so far.

Remark 2.4.6. Notice that in the case $k^+ = 1 = k^-$ we can always satisfy the hypothesis of [8, Theorem 7.1], hence we can move branch points without crossing the real curves to obtain a structure which is a bubbling. This is a key fact in the proof of the main theorem below (3.5.6).

We conclude by observing the following funny application of 2.4.5.

Corollary 2.4.7. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. Then the number of branch points contained in σ^+ and the total number of real components always sum to an odd number.*

Proof. If $k^+ = 0, 2$ then there is a negative component of Euler char ± 1. In both cases it has an odd number of boundary components. All the other negative components are incompressible annuli. The total number of real components is therefore odd. if $k^+ = 1$ then the positive branched component is incompressible and there is exactly one index 1 real boundary, beyond which the negative branched component sits. And it is an annulus. All other negative components are annuli too, hence we have an even number of real components. □

2.4.1 Classification of combinatorics

Theorem 2.4.5 gives some restrictive conditions for a combinatorics to be allowed, i.e. to be the combinatorics of a geometric decomposition of a quasi-Fuchsian BPS with two branch points. For instance it implies that the interesting part of the structure, such as branch points and real components with positive index, is quite localised, and that the structure of the negative part is almost the same as in an unbranched structure; more precisely we have the following statement.

Lemma 2.4.8. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be geometrically branched. Then, up to moving branch points in their geometric component, every unbranched negative incompressible component is contained in a grafting annulus.*

Proof. We already know by 2.2.2 and 2.2.3 that such an unbranched negative incompressible component is an annulus A bounded by two real curves l_1, l_2 with loxodromic holonomy and index 0. It is enough to prove that the ends E_1 and E_2 of l_1 and l_2 inside the adjacent positive component(s) do not contain branch points and that their peripheral geodesics γ_1, γ_2 are embedded, so that $E_1 \cup A \cup E_2$ carries the complete unbranched structure of a grafting annulus. If $k^+ = 0$ then this is obvious. In the other cases a component C adjacent to A contains branch points if and only if it has exactly one additional real boundary $l_0 \neq l_1, l_2$ of index 1 (for instance by 2.4.5). But in this case it is enough to move these branch points in C so that they leave the closure of the ends E_1, E_2 (recall that a peripheral geodesic can self-intersect only at a branch point by 2.2.11). In the case $k^+ = 1$ this movement is possible because by 2.2.12 the branch point must live in the closure of the end relative to l_0; moreover ends are disjoint by [8, Lemma 3.13], hence it is enough to move the branch point so that it does not belong to the closure of E_1, E_2 but only to the closure of E_0. In the case $k^+ = 2$ we apply the degeneration dichotomy of [8, Proposition 8.1]: either we can collapse the two branch points to a single branch point of order 6π or a bubble appears. In the first case the argument is similar to the previous one, i.e. we just move this branch point so that it belongs to the closure of E_0 but not of E_1, E_2. In the second case, the bubble which appears must necessarily intersect the real curve l_0, simply because there are no other real curves of index 1 around. But for such a bubbling actually both branch points must belong to the closure of the end E_0; once again, it is thus enough to move them a little bit so that they do not belong to the closure of E_1, E_2. □

Lemma 2.4.9. *Let $\mathfrak{c}$ be an allowed combinatorics. If $\mathfrak{c}$ has $k^+ = 2, 1$ or 0 respectively, then it can be degrafted to a combinatorics $\mathfrak{c}'$ with at most $1, 2$ or 3 real components respectively.*

Proof. Let $\sigma \in \mathcal{M}^{\mathfrak{c}}_{2,\rho}$. By the previous lemma, up to a a movement of branch points which does not cross the real curves, hence does not change the combinatorics, every negative unbranched incompressible component can be degrafted. Therefore we only need to look at real curves bounding non essential or branched negative components. By 2.4.5 when $k^+ = 2$ there is only one negative disk, when $k^+ = 1$ there is only one branched annulus and when $k^+ = 0$ there is only one branched negative component with one or three real boundaries, and all other negative components are incompressible unbranched annuli. □

Motivated by these observations, we now proceed to the explicit classification of combinatorics having at most 3 real components. Moreover we show that each of them can be realised by performing a grafting and bubbling of the uniformizing structure along suitable curves.

We begin with structures with a single component of the real curve. We know by 2.4.7 that either both branch points are in the positive part, or both are in the negative part.

1. In the first case we have a negative unbranched disk and a genus g positive component glued to the disk and containing both branch points; the index of the real curve is 1. Notice this combinatorics is obtained by a standard bubbling on a hyperbolic surface

2. In the second case there is a negative branched component containing both branch points with Euler characteristic -1. Since there is a single real component, it is necessarily a once-holed torus. The real curve has index 0 and there is just another positive unbranched component of genus $g-1$ glued to it. Notice this combinatorics is obtained by taking a bubbling over a non separating grafting of a hyperbolic surface (as in the right side of 2.4).

Now we consider structures with two components of the real curve. By 2.4.7 we know branch points have different signs, hence there is a real component l of index 1 which is shared by a positive branched component with one branch point and a negative branched incompressible annulus with one branch point. We have two cases here

1. If the positive component has another boundary besides l, then it must be the other boundary l' of the annulus. In this case the positive component has genus $g-1$. This combinatorics is obtained by bubbling a non separating grafting of a hyperbolic surface (as in the left side of 2.4, but with the grafting done on a non separating geodesic).

2. If the positive component has no other boundaries, then it has some genus g_1 and there is another positive component of genus $g-g_1$ adjacent to the other boundary l' of the annulus. This combinatorics is obtained by bubbling a separating grafting of a hyperbolic surface (as in the left side of 2.4)..

Finally we come to structures with three real components. In this case by 2.4.7 we have $k^+=0$ or 2. We begin with the case $k^+=2$. As above, in this case there is a negative disk and a positive branched component C adjacent to it.

1. if C has two more boundaries, then it has genus $g-1$ and there is a negative annulus joining the two extra boundaries. This combinatorics is obtained by bubbling a non separating grafting of a hyperbolic surface (outside the grafting annulus).

2. if C has only one more boundary, then it has some genus $g_1>0$, beyond the other boundary there is a negative annulus, and then another positive component of genus $g-g_1$. This combinatorics is obtained by bubbling a separating grafting of a hyperbolic surface (outside the grafting annulus).

Now consider the case $k^+=0$. In this case by 2.4.5 the negative branched component C is either a pair of pants (with at most one non essential boundary) or an incompressible once-holed torus.

1. If C is an incompressible torus, then adjacent to its boundary we find a positive incompressible component of genus $g_1>0$ with another boundary, then a negative incompressible annulus, and finally another positive incompressible component of genus $g-g_1$. This combinatorics can be obtained by grafting a structure with one real component and $k^+=0$ along a geodesic in the positive part (as the one coming from the right side of 2.4).

2. If C is a pair of pants with a non essential boundary, then this curve bounds a positive disk, and thus we have the following cases

   - if the other two boundaries are adjacent to the same positive component C, then it has genus $g-1$. This combinatorics is obtained by bubbling a non separating grafting of a hyperbolic surface (inside the grafting annulus).
   - if the other two boundaries are adjacent to different positive components, then they have genus $g', g''>0$ summing to g. This combinatorics is obtained by bubbling a separating grafting of a hyperbolic surface (inside the grafting annulus).

3. If C is an incompressible pair of pants, then we have the following cases (see Picture 2.6 below):

   - if all boundaries are adjacent to the same positive component C', then it has genus $g-2$. This combinatorics is obtained by grafting a hyperbolic surface along two nearby non separating geodesics and then bubbling along an arc with one endpoint in each negative annulus;
   - if one of the boundaries is adjacent to a positive component C' while the other two are adjacent to another single component C'', then C' has genus $g'>0$ and C'' has genus $g-g'-1$ This combinatorics is obtained by bubbling a separating grafting of a hyperbolic surface along an arc with endpoints in the negative annulus but not itself contained inside it;
   - if each boundary is adjacent to a different component, then the three components have genus $g', g'', g'''>0$ summing to g (this requires $g\geq 3$). This combinatorics is obtained by grafting a hyperbolic surface along two nearby separating geodesics and then bubbling along an arc with one endpoint in each negative annulus.

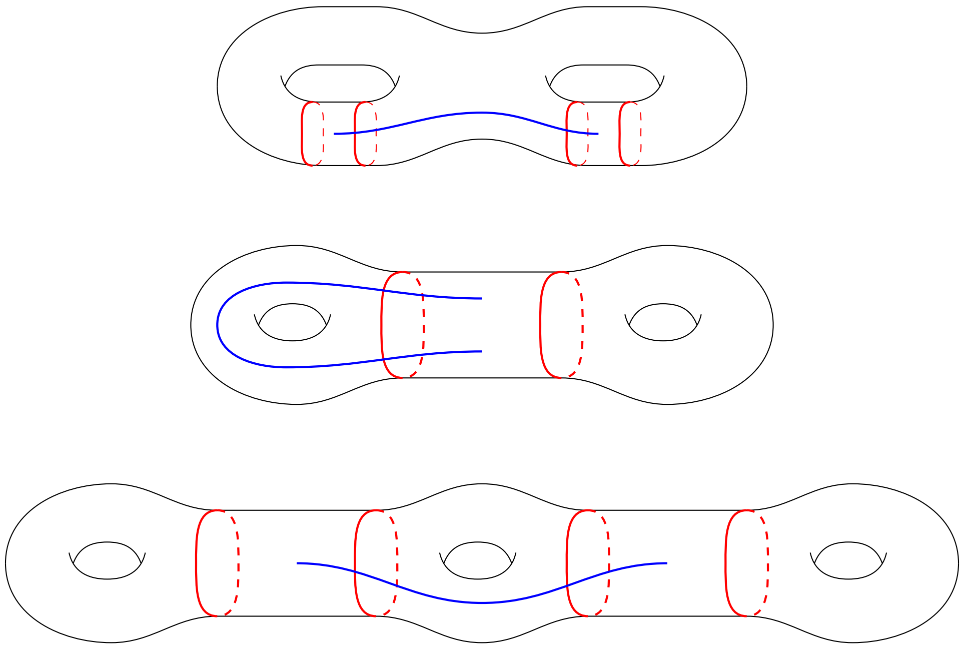

Figure 2.6: Examples of bubbling for the third case.

This concludes the classification of allowed combinatorics with at most three real components. Notice that it is possible to obtain each of them as the combinatorics of a structure which is a suitable bubbling over a suitable grafting (with at most two grafting annuli) of the uniformizing structure. This means that a combinatorial analysis of the geometric decomposition can not provide any obstruction to the fact that a given BPS with quasi-Fuchsian holonomy is a bubbling over some unbranched structure.

2.4.2 Really branched structures

In the previous sections it was always assumed that all branch points were outside the real curve. The index formula 2.1.16 from [8] has been used to obtain an explicit description of which combinatorics can arise for such a decomposition. In this section we want to obtain analogous results for really branched structures as an application of what we know about the geometrically branched ones.

Remark 2.4.10. If σ is a really branched BPS, then the real curve $\sigma^{\mathbb{R}}$ is a finite union of disjoint closed curves on S (i.e. a finite but possibly disconnected 1-dimensional CW complex), which we will call the **real graph**, to stress the fact that it has singular components. A real branch point of order k gives rise to a vertex of order $2(k+1)$ on this graph. Moreover the orientation of the surface induces a cyclic order on the (germs of) edges at a vertex, and edges carry a canonical orientation such that the positive (respectively negative) part sits on the left (respectively right) side of the edge, as in the non really branched case. Therefore we can view the real graph as an oriented ribbon graph.

We are now interested in understanding which are the allowed combinatorics for such a graph and for the resulting geometric decomposition. The strategy will be to combine an analysis of the branched $(\mathrm{PSL}_2\mathbb{R}, \mathbb{RP}^1)$-structure on the real curve with an analysis of "neighbouring structures", i.e. those obtained by a tiny movement of branch points outside the real curve, and maybe also a splitting of a double branch point into two simple ones. To cover all the possibilities occurring when the total branching order is 2, we need to consider the cases in which a double branch point, two simple branch points or one simple branch point live on the real graph. To simplify the discussion we also begin with assuming that the real graph is connected.

Let us start by considering the case of a BPS σ with **one double real branch point**. In this case the singular point p of the real graph has order six. We cyclically label the six germs of edges at p as $(1, 1', 2, 2', 3, 3')$ in such a way that the germ corresponding to 1 is oriented outwards. Given such a local picture, the combinatorics

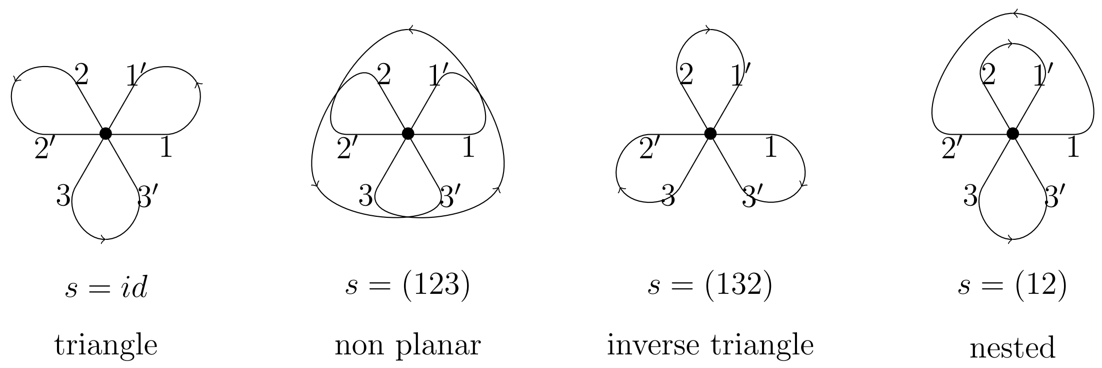

Figure 2.7: The four possibilities for the real graph for structures with a double real branch point.

of the whole real component containing p is determined by the choice of a way to join each germ with label m with a germ with label n', i.e. a choice of a permutation in $\mathfrak{S}_3$. It is easily checked that the three transpositions give rise to the same graph, hence we are left with the following four cases, which we label by the corresponding permutation $s \in \mathfrak{S}_3$ (in cycle notation): the triangle $s = id$, the non planar $s = (123)$, the inverse triangle $s = (132)$ and the nested $s = (12)$, shown in Picture 2.7. Then the whole structure on S is obtained by gluing geometric components along a tubular neighbourhood of this graph, recalling that positive components lie on the left of the edges, and negative components to the right. The one obtained from $s = (123)$ is not planar as a ribbon graph: a tubular neighbourhood of it inside the surface is a torus with two holes. Moving branch points outside the real curve and studying the resulting structure allows to forbid some of the above possibilities.

Lemma 2.4.11. *There can be no real graph with permutation* $s = (132)$.

Proof. Assume that σ is a really branched BPS with a real graph Γ of type (132). Notice that the neighbourhood of this graph in S is a sphere with four holes; moreover the external boundary is adjacent to a positive component, whereas the three inner boundaries are adjacent to some negative component(s). Moving the branch point to the positive part, we get a structure σ' with three real components and $k^+ = 2$. By assumptions there are no other real components, hence by the above classification we know that the negative part consists of an unbranched disk and an unbranched annulus. But then on σ we have that one loop of Γ bounds a negative disk, hence the developing map is surjective (onto the limit set of ρ) along it, while on the other hand the concatenation of the three loops gives the external boundary and bounds a positive unbranched non contractible component, thus the developing map along it should be injective, absurd. □

Remark 2.4.12. This graph is not allowed even if we allow for more real components (beyond the singular one). Indeed suppose that the real graph contains some other component. As we have seen, moving the point to the positive component gives a structure with three real components; by 2.4.7 the extra real components are in even number. The same argument of the previous proof shows that a geometric component adjacent to the singular real component can not be a disk, hence it must be an annulus, the other boundary of which is either on another loop of the singular real component or on a regular real component. But then if we move to the negative component, we see that the singular component gives rise to a single real component and we get a structure with a negative component containing two branch points and having two or four boundaries, hence even Euler characteristic, which is absurd by the classification in 2.2.9.

Lemma 2.4.13. *There can be no real graph with permutation* $s = (12)$ *(or any other transposition).*

Proof. Assume that σ is a really branched BPS with a real graph Γ of type (12). Then we consider a movement of branch points which blows up the double point into two simple branch points and moves one to the positive part and one to the negative part. By assumptions there are no other real components, hence we would get a structure with $k^+ = 1$ and an odd number of real components, which is absurd by 2.4.7. □

Remark 2.4.14. The graph associated to transpositions can be realised as a component of the real graph of a BPS, but we need to allow for at least one non singular component beyond the singular one. If we do so, we can realise this structure as a (limit of) bubbling: graft σ_ρ along a geodesic, then bubble along a subarc of one of the two real curves; finally move branch points away from each other along the real curve to get a doubly branched structure.

We now show that graphs of type id and (123) are indeed realised.

Example 2.4.15. (The triangle id) This combinatorics is obtained by (degrafting) the structure called "the triangle" constructed in [8] as an example of a Fuchsian structure on a surface of genus 2 with a real component whose peripheral geodesic in the positive part is not embedded. This construction relies on the choice of a Schottky group of a pair of pants $\Gamma =< \alpha, \beta >$ and of a basepoint $x \in \mathbb{RP}^1$ such that $x, \alpha(x), \beta\alpha(x)$ appear in cyclic order.

Example 2.4.16. (The non-planar (123)) We produce a structure on a genus 2 surface. First of all we "compute" what is the expected combinatorics. Since the neighbourhood of the graph is a twice-holed torus, to get a surface of genus 2 we can only add a disk and a once-holed torus. Moving the branch point to the positive part we get a structure with a single real component and 2 branch points in the positive region; by the classification this is the combinatorics of a standard bubbling, therefore necessarily the disk is negative and the holed torus is positive. A substantial difference with the triangle should be noticed here: none of the loops a, b or c of the graph is a geometric boundary, but only their concatenations abc and acb are. This makes it trickier to follow the developing map along the real graph. Nevertheless we can adapt the construction of the triangle from [8] starting with a Schottky group of a one-holed torus $\Gamma =< \alpha, \beta >$.

We conclude this chapter by saying that a similar analysis can of course be carried out for real structures inside the principal stratum, i.e. structures with two simple branch points on the real curve or with one simple branch point on the real curve and one simple branch point in some geometric component. There will clearly be more possible combinatorics for the real graphs, and we could try to understand how these can vary when moving, collapsing and splitting branch points inside $\mathcal{M}^{\mathbb{R}}_{2,\rho}$. However the resulting picture is not so well behaved as one might expect: for instance any really simply branched structure whose branch points belong to different connected components of the real graph can not be deformed into a really doubly branched structure by a deformation inside $\mathcal{M}^{\mathbb{R}}_{2,\rho}$, simply because moving branch points on the real graph does not change the number of its connected components. As a result it is not even clear if $\mathcal{M}^{\mathbb{R}}_{2,\rho}$ is a connected hypersurface inside $\mathcal{M}_{2,\rho}$.

Chapter 3

Spaces of bubblings

As mentioned in 1.4.3, given an unbranched projective structure with holonomy ρ we can produce many examples of BPSs with the same holonomy and two simple branch points by performing a bubbling along an embedded arc which is injectively developed. This gives rise to a subspace of the principal stratum $\mathcal{M}_{(1,1),\rho}$, which we denote by $\mathcal{BM}_{(1,1),\rho}$, and in this chapter we are interested in its topology: we will prove that in quasi-Fuchsian holonomy it is an open and dense subspace of full measure (see 3.5.6). More generally we will use the following notation.

Definition 3.0.17. Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a representation. If $\mathcal{X} \subset \mathcal{M}_{(1,1),\rho}$ is any subspace, we denote by $\mathcal{BX}$ the subspace of $\mathcal{X}$ consisting of all the structures σ of $\mathcal{X}$ which are obtained by bubbling some unbranched structure $\sigma_0 \in \mathcal{M}_{0,\rho}$.

Notice that there is in general no "underlying unbranched structure", i.e. a structure σ can be a bubbling over different unbranched structures along different arcs; in 3.6 below we will study this phenomenon in detail and exploit it to reduce multi(de)graftings to bubbling/debubbling constructions on quasi-Fuchsian structures.

As a warm-up, in the first section we show that if the holonomy is the trivial representation, then any structure in $\mathcal{BM}_{(1,1),\rho}$ has actually genus 0 and is realised as a bubbling of the Riemann sphere. The other sections are concerned with the case of non elementary holonomy (and in particular with the quasi-Fuchsian case). In this context moving branch points provides coordinates for the moduli space (as established in [8]), hence we will address some problems that arise when bubbling and moving branch points interact.

3.1 Trivial holonomy

In this section we let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be the trivial representation $\rho = id$. Under this assumption it is easy to prove "by hands" (i.e. without developing more technology) that a structure in $\mathcal{M}_{(1,1),id}$ is indeed obtainable as a bubbling. The trick is that such a structure can occur only in genus $g = 0$ and therefore the developing map is just a self-map of $\mathbb{CP}^1$ of degree 2.

Lemma 3.1.1. *A BPS on S has trivial holonomy if and only if it is induced by a finite branched cover $S \to \mathbb{CP}^1$ (i.e. a non constant meromorphic function).*

Proof. As already observed in 1.2.8 a branched cover $f : S \to \mathbb{CP}^1$ endows S with a branched projective structures σ with a branch point of angle $2\pi(k+1)$ at every branch point of order k; the map f defines a developing map for σ, equivariant with respect to the trivial representation. Conversely suppose we have a developing map $dev : \widetilde{S} \to \mathbb{CP}^1$ which is equivariant with respect to the trivial representation; then $dev(\gamma x) = dev(x)$ for any $x \in \widetilde{S}$, hence dev induces a map $f : S \to \mathbb{CP}^1$ which is a branched cover (since dev is) of finite degree (since S is compact). □

Here the developing map is a branched cover defined already on the surface; we will use the branch values as an analogue in this context of the developed divisor on the quasi-Fuchsian model in the case of quasi-Fuchsian holonomy (see 2.3). Hence we can accordingly extend the definition of simply developed structure, by saying that a BPS with trivial holonomy is simply developed if the branched cover defining it maps injectively branch points to branch values, or, equivalently, if the branching is total (i.e. the fibre over a critical value consists of one critical point of multiplicity equal to the degree of the cover).

Lemma 3.1.2. *If S admits a branched covering map $f : S \to \mathbb{CP}^1$ with exactly two branch values, then S is a sphere and the branching is total.*

Proof. Let $x, y \in \mathbb{CP}^1$ be the two branch values. Let x_i and y_j be the points in the fibre above x and y, with orders λ_i and μ_j respectively, for $i = 1 \dots m, j = 1 \dots n$. Notice that the sum of the branching order over all points above the same value equals the degree of the cover f. Then Riemann-Hurwitz implies

$$\chi(S) = deg(f)\chi(\mathbb{CP}^1) - \sum_{i=1}^{m}(\lambda_i - 1) - \sum_{j=1}^{n}(\mu_j - 1) =$$

$$= 2deg(f) - \sum_{i=1}^{m}\lambda_i - \sum_{j=1}^{m}\mu_j + m + n = m + n$$

In particular the characteristic of S does not depend on the degree of the cover and is always $\chi(S) = m + n \geq 2$, hence S is always a sphere. This implies that actually $2 = \chi(S) = m + n$, hence $m = n = 1$. □

We can now prove the following.

Proposition 3.1.3. *Let S be a closed surface and $\sigma \in \mathcal{M}_{(1,1),id}$. Then S is a sphere, σ is simply developed and it is a bubbling over the Riemann sphere.*

Proof. Since the holonomy of σ is trivial, it is defined by a branched cover $f : S \to \mathbb{CP}^1$ by 3.1.1. A Riemman-Hurwitz computation as the one above shows that there are no branched covers $S \to \mathbb{CP}^1$ with a single branch value, hence f has exactly two branch values, which implies that σ is simply developed (and also that $deg(f) = 2$ and the branching is total). By 3.1.2 then S is a sphere. Let $p, q \in S$ be the branch points of f and $\widehat{p}, \widehat{q}$ be the corresponding branch values on $\mathbb{CP}^1$. Let $\gamma : [0, 1] \to \mathbb{CP}^1$ be a smooth arc between $\widehat{p}$ and $\widehat{q}$. Then $\mathbb{CP}^1 \setminus \gamma$ is a disk avoiding branch points, hence we can lift it to an open disk $D \subset S \setminus \{p, q\}$. The boundary of D will be given by a couple of paths γ_1, γ_2 between p and q, each of which is embedded and injectively developed (since γ is a simple arc). Moreover the two paths are indeed disjoint outside branch points since they live in $S \setminus \{p, q\}$ which is an unbranched structure (on an open surface). Therefore $\overline{D}$ is actually a bubble on σ. □

Since the trivial representation preserves the standard round metric on the sphere, we have the following straightforward consequence.

Corollary 3.1.4. *Any spherical metric on S^2 with two cone points of angle 4π is obtained by a bubbling of the standard round sphere.*

3.2 Standard BM-configurations

In this section let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a non elementary representation. It is proved in the Appendix of [8] that moving branch points gives coordinates for the moduli space $\mathcal{M}_{k,\rho}$ (for the definition of this deformation see 1.4.23). Hence it is natural to try to analyse the topology of the subspace of bubblings $\mathcal{BM}_{k,\rho}$ by moving branch points. To do this we need to understand what happens when we try to move branch points along an embedded twin pair based at one of the vertices of a bubble.

Definition 3.2.1. Let $\sigma \in \mathcal{BM}_{(1,1),\rho}$. A **BM-configuration** (Bubbling-Movement configuration) on σ is the datum of a bubble B together with an embedded twin pair μ based at a vertex p of B (i.e. at a branch point of σ). We denote the configuration by (B, μ, p).

We introduce now the nicest type of BM-configuration, which will allow us to perform local deformations of the structure preserving the bubble.

Definition 3.2.2. A BM-configuration (B, μ, p) on $\sigma \in \mathcal{BM}_{(1,1),\rho}$ is said to be a **standard BM-configuration** if either all the arcs are disjoint and disjointly developed outside the obvious intersections (i.e. $\partial B \cap \mu = \{p\}$ and $dev(\partial B) \cap dev(\mu) = \{dev(p)\}$) or the embedded twin pair is entirely contained in the boundary of the bubble (i.e. $\mu \subset \partial B$).

Notice that, given a standard BM-configuration (B, μ, x) of the second type, a very tiny isotopy of the bubble (which is allowed by 1.4.13) reduces (B, μ, x) to a standard BM-configuration of the first type. Namely in any projective coordinate we can push the developed image of the arc of bubbling slightly to the left or right of itself; when referring to a standard BM-configuration we will really always think of the first case. We have the following characterisation.

Lemma 3.2.3. *Let $\sigma \in \mathcal{BM}_{(1,1),\rho}$ and let (B, μ, p) be a BM-configuration on it. Let β be the induced bubbleable arc on $\sigma_0 = Debub(\sigma, B)$. Then (B, μ, p) is a standard BM-configuration if and only if μ induces an arc μ' on σ_0 such that the concatenation of β and μ' is a bubbleable arc on σ_0.*

Proof. When debubbling σ with respect to B we naturally end up with an unbranched structure σ_0 endowed with a bubbleable arc β such that $Bub(\sigma_0, \beta) = \sigma$. One of the two arcs contained in the embedded twin pair, let us say μ_2, starts outside the bubble, hence its germ survives in σ_0, and we can try to analytically continue it

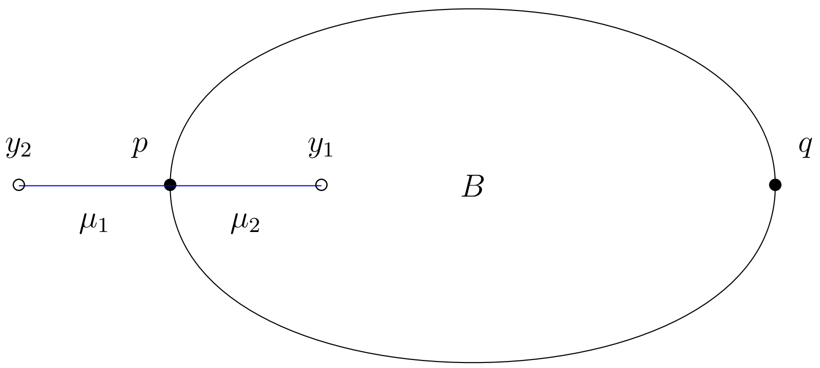

Figure 3.1: A standard BM-configuration

to a path μ_0 which has the same developed image of μ. If the BM-configuration is standard then μ_2 never meets the bubble, thus μ_0 is a simple arc on σ_0, which does not meet β away from p; in other words the concatenation of β and μ_0 is a simple arc on σ_0. Moreover the developed image of this arc is given by the concatenation of the developed image of ∂B and μ, which are disjoint. Thus this arc is bubbleable on σ_0. Conversely, if this arc is bubbleable, then when we perform the bubbling we can reconstruct the embedded twin pair μ by looking for the twin of μ_0 inside the bubble. Since the whole $\beta\mu_0$ is bubbleable and we are bubbling only along the subarc β, we see that the developed image of the remaining part does not cross the developed image of β. This means exactly that the twin starting inside the bubble will not leave it. Therefore the induced BM-configuration is standard. □

The interest in standard BM-configurations is motivated by the following lemma.

Lemma 3.2.4. *Let $\sigma_0 \in \mathcal{BM}_{(1,1),\rho}$ and (B, μ, p) be a standard BM-configuration based at p; let σ_t be the BPS obtained by moving branch points on σ_0 along μ up to time t, where $t \in [0,1]$ is a parameter along the developed image of μ. Then $\sigma_t \in \mathcal{BM}_{(1,1),\rho}$ for all $t \in [0,1]$.*

Proof. This directly follows from the characterisation in 3.2.3 together with [8, Lemma 2.9]. Indeed, with the above notations we have that $\sigma_t = Move(\sigma_0, \mu^t) = Bub(\sigma_0, \beta\mu'^t) =$, where μ^t and μ'^t are the subarcs of μ and μ' respectively from time 0 to time t. Concretely, moving a branch point along a standard BM-configuration gives rise to a 1-parameter family of BPS which can actually be obtained by bubbling a fixed unbranched structure along increasing subarcs of a fixed bubbleable arc. □

We are now ready to prove the following result.

Theorem 3.2.5. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a non elementary representation. Then $\mathcal{BM}_{(1,1),\rho}$ is open in $\mathcal{M}_{(1,1),\rho}$ (hence in $\mathcal{M}_{2,\rho}$).*

Proof. By 3.2.4 it is enough to show that given $\sigma_0 \in \mathcal{BM}_{(1,1),\rho}$ there is a small neighbourhood U of it such that any structure in U is obtained by moving branch points along a standard BM-configuration on σ_0. This easily follows from the fact that moving branch points gives a full neighbourhood of σ_0 in the moduli space, because local movements of branch points can always be performed along embedded twin pairs which are in standard BM-configuration with a given bubble on σ_0: namely a local neighbourhood of the developed image of a branch point is path connected, and still is if we remove the developed image $\widehat{\beta}$ of the boundary of the bubble, so that we can move to any point in that neighbourhood by means of an arc which either avoids $\widehat{\beta}$ or is entirely contained in it (depending on the fact that the point is outside $\widehat{\beta}$ or on it). □

Notice however that a priori more complicated BM-configurations might arise, which can not be used to move branch points preserving the bubble; namely if the embedded twin pair intersects the boundary of the bubble (or if this holds for their developed images), then moving branch points results in the break of the bubble: the aspiring bubbleable arc is either not embedded or not injectively developed. In this case it is not clear if it is possible to find another bubble. This heuristic argument

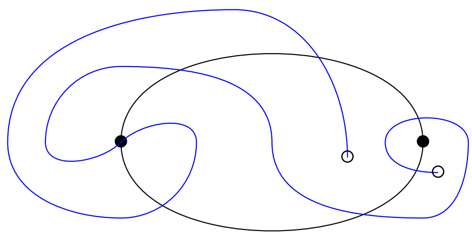

Figure 3.2: A non standard BM-configuration

can be made more precise by the following observation: moving branch points on a standard BM-configuration preserves the isotopy class (relative to endpoints) of the bubble; in particular it does not change the underlying unbranched structure. On the other hand it is not difficult to produce examples of movements of branch points which do not preserve the underlying unbranched structure.

Example 3.2.6. Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a Fuchsian representation, β be a bubbleable arc on the hyperbolic surface $\sigma_\rho = \mathbb{H}^2/\rho$ and $\sigma = Bub(\sigma_\rho, \beta)$. Notice that σ is simply developed (see 2.3.7 for the definition). On the other hand it is possible to move branch points on σ along suitable embedded twin pairs μ and ν with both endpoints inside the bubble in such a way that the resulting structure does not have this property (see Picture 3.3). This prevents the structure $Move(\sigma, \mu)$ from being a bubbling over σ_ρ, for instance by 2.3.9. Of course the BM-configuration on σ is not standard.

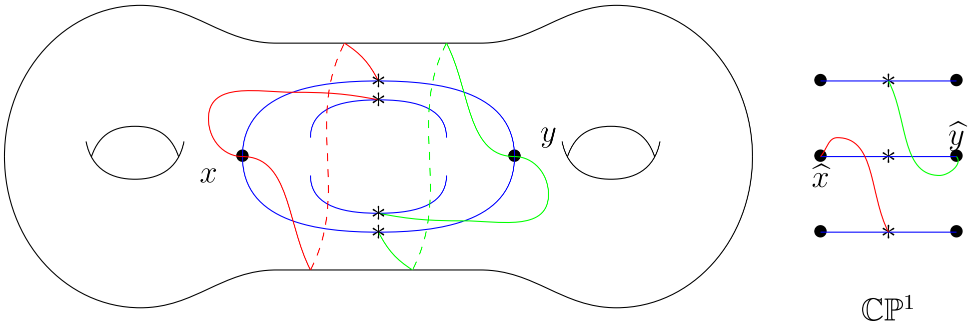

Figure 3.3: Picture for Example 3.2.6.

3.3 Taming developed images and avatars

One of the main technical issues about $\mathbb{CP}^1$-structures is that the developing map is dramatically non injective (already in the case of unbranched structures), hence it is quite difficult to control the relative behaviour of the developed images of some configuration of objects on the surface, even when the configuration is well behaved on the surface. In this section we introduce the collection of avatars of a given object on the surface, and prove some technical lemmas to handle it. The first definitions can be given in full generality, but we will soon focus on unbranched quasi-Fuchsian structures.

Definition 3.3.1. Let σ be a BPS on S and let $U \subset S$ be any subset. We say that U is **tame** if for some lift $\widetilde{U}$ of U we have that a developing map for σ is injective when restricted to $\cup_{\gamma \in \pi_1(S)} \gamma.\widetilde{U}$.

Notice that a tame simple arc is in particular bubbleable, and that a tame simple closed curve is in particular graftable as soon as it has loxodromic holonomy. The following lemma provides an easy criterion to prove tameness in this case of quasi-Fuchsian holonomy.

Corollary 3.3.2. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian . Then any subset of the uniformizing structure σ_ρ is tame. More generally if $\sigma \in \mathcal{M}_{0,\rho}$ and $C \subset \sigma^\pm$ is a geometric component, then any subset of the convex core of C is tame.*

Proof. The first statement follows directly from the fact that the developing map of σ_ρ is globally injective. For the second one we observe that a subset of the convex core avoids all grafting annuli. Therefore the collection of its developed images is the same that we would see if we removed all grafting annuli, i.e. the same that we see on the uniformizing structure. □

The following easily follows from the previous results.

Corollary 3.3.3. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian and $\sigma \in \mathcal{M}_{2,\rho}$. If σ is a bubbling of some $\sigma_0 \in \mathcal{M}_{0,\rho}$ along a tame arc or along an arc contained in the convex core of a geometric component, then σ is simply developed.*

Being able to control the collection of developed images of a given object on the surface will not be enough in the following. For example, even if we start with a very well behaved structure σ_0 (e.g. a hyperbolic surface), when we perform a bubbling or a grafting we introduce in our structure σ_0 a whole region R whose full developed image is the whole model space $\mathbb{CP}^1$; as a result inside R we "see" a lot of developed images of any given object $O \subset \sigma_0$. The following definition aims at making this more precise.

Definition 3.3.4. Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a representation, $\sigma \in \mathcal{M}_{k,\rho}$ and $U \subset \sigma$ be any subset. An **avatar** of U is any subset $V \subset \sigma$ such that there exist a lift $\widetilde{U}$ of U and a lift $\widetilde{V}$ of V such that $dev(\widetilde{U}) = dev(\widetilde{V})$.

Of course, by equivariance of the developing map, if V is an avatar of U then for any lift of U there exist a lift of V with the same developed image. Let us discuss some examples.

Example 3.3.5. If a structure has an injective developing map, then having the same developed image means being the same set, so that there are no non-trivial avatars. This happens for the uniformizing structure σ_ρ for a quasi-Fuchsian representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$.

Example 3.3.6. This example wants to show that the concept of avatar extends the one of non simply developed structures in quasi-Fuchsian holonomy. Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation and $\sigma \in \mathcal{M}_{k,\rho}$ be a geometrically branched BPS. Let x, y be two branch points belonging to the same component $C \subset \sigma^\pm$; if $D_C(x) = D_C(y)$ then x and y are avatars: indeed there exist lifts $\widetilde{x}, \widetilde{y}$ whose developed image differ by some $g = \rho(\gamma) \in \rho(\pi_1(C))$ by definition of $D_C : C \to C_{Fuchs}$; but then $\widetilde{x}, \gamma\widetilde{y}$ are lifts witnessing that x and y are avatars. The converse does not hold in general, simply because branch points might live in different components, but of course if $k = 2$ then x, y are avatars if and only if σ is not simply developed: if the points are avatars then they have the same sign, thus they belong to the same component by 2.4.5.

Example 3.3.7. Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation and $\sigma \in \mathcal{M}_{0,\rho}$. Let $C \subset \sigma^+$ be a positive geometric component (the same analysis works for a negative one) and $U \subseteq C$. Let us fix a lift $\widetilde{C}$ of C to the universal cover, a lift $\widetilde{U}$ of U contained in it and let $\widehat{U} = dev(\widetilde{U})$. Then we have that $\gamma\widetilde{U}$ sits in $\widetilde{C}$ if and only if $\gamma \in \pi_1(C)$ and that the whole collection $\pi_1(S)\widetilde{U}$ develops to the upper-half plane. Now recall that geometric components of σ carry complete unbranched structures, hence the restriction of the developing map to any lift of C is a diffeomorphism with the upper-half plane (in particular it is surjective). Therefore the whole collection $\{\rho(\gamma)\widehat{U} \mid \gamma \in \pi_1(S)\}$ can be pulled back to $\widetilde{C}$: elements with $\gamma \in \pi_1(C)$ will give the collection of lifts of U to $\widetilde{C}$, whereas those coming from $\gamma \notin \pi_1(C)$ will give us a collection of sets not corresponding to lifts of U to $\widetilde{C}$ but which still develop in the same way as some lifts of U in other lifts of C; projecting these ones to C gives a lot of non-trivial avatars of U in C. As we will see in detail below, the avatars of U in C can be labelled by the cosets of $\pi_1(C)$ in $\pi_1(S)$. Of course we can do the same construction for any component of the geometric decomposition of the structure to produce avatars of U in components different from C, but with the same sign (otherwise the collection turns out to be empty by definition). Namely the collection of avatars of U in some component C' is given by

$$\pi_{|_{C'}}\left(dev^{-1}_{|_{C'}}\left(dev\left(\pi_1(S)\widetilde{U}\right)\right)\right)$$

where π denotes the universal cover projection and $\widetilde{U}$ is any lift of U to the universal cover. The idea is that any geometric component will *see* avatars of the chosen subset U, and actually infinitely many of them, so the situation can in general be quite complicated.

However in quasi-Fuchsian holonomy we have a well-defined notion of size for subsets avoiding the real curve, which allows us to control the collection of avatars of a small set, as the following result shows. Let us denote by $sys(\rho)$ the systole of the uniformizing structure σ_ρ or, equivalently, the minimum of the translation lengths of the elements in $\rho(\pi_1(S))$.

Lemma 3.3.8. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian and $\sigma \in \mathcal{M}_{0,\rho}$. Let $U \subset \sigma$ be a connected set with $diam(U) < sys(\rho)$ and which is π_1-trivial (i.e. $i_*(\pi_1(U)) \subset \pi_1(S)$ is the trivial subgroup). Then U sits inside a geometric component, it is tame and its avatars are disjoint.*

Proof. Recall that when the holonomy is quasi-Fuchsian there is a well defined hyperbolic metric on the complement of the real curve, which blows up in a neighbourhood of it; hence we can define a generalised path metric on the whole surface. Any connected subset of σ which intersects the real curve must have infinite diameter with respect to this metric, because any path intersecting the real curve has infinite length. Therefore U can not intersect the real curve, hence it is contained in some geometric component.

Since U is π_1-trivial, it lifts homeomorphically to the universal cover. To prove tameness, assume that there are two lifts $\widetilde{U}_1$ and $\widetilde{U}_2$ which overlap once developed, i.e. $\exists\ x_i \in \widetilde{U}_i$ such that $dev(x_1) = dev(x_2)$. Let $\gamma \in \pi_1(S)$ be the unique deck transformation such that $\gamma\widetilde{U}_1 = \widetilde{U}_2$. Then we have the following absurd chain of inequalities

$$sys(\rho) \leq d(\rho(\gamma)dev(x_1), dev(x_1)) = d(\rho(\gamma)dev(x_1), dev(x_2)) =$$

$$= d(dev(\gamma x_1), dev(x_2)) \leq diam(dev(\widetilde{U}_2)) = diam(U) < sys(\rho)$$

where d denotes the hyperbolic distance on $\mathbb{CP}^1 \setminus \mathbb{RP}^1$ and the last equality follows from the fact that the developing map is an isometry for each geometric component.

Finally let us prove that the avatars in each geometric component are disjoint. Let C be a geometric component, and choose a lift $\widetilde{C}$ of it and a lift $\widetilde{U}$ of U. By the above discussion, we have to check that the collection

$$\pi_{|_C}\left(dev^{-1}_{|_C}\left(dev\left(\pi_1(S)\widetilde{U}\right)\right)\right)$$

is a disjoint collection. By tameness we know that the collection $dev\left(\pi_1(S)\widetilde{U}\right)$ is disjoint, and the same is true for $dev^{-1}_{|_C}\left(dev\left(\pi_1(S)\widetilde{U}\right)\right)$, since the restriction of the developing map to each geometric component is a diffeomorphism. So we only need to prove that the projection π does not overlap things too much. Let us introduce the following notation: if $\gamma \in \pi_1(S)$ then

$$\gamma * \widetilde{U} := dev_{|_C}^{-1}\left(dev\left(\gamma\widetilde{U}\right)\right)$$

With this notation what we want to prove now is that if there exist $\gamma_1, \gamma_2 \in \pi_1(S)$ such that $\pi\left(\gamma_1 * \widetilde{U}\right) \cap \pi\left(\gamma_2 * \widetilde{U}\right) \neq \varnothing$ then actually $\pi\left(\gamma_1 * \widetilde{U}\right) = \pi\left(\gamma_2 * \widetilde{U}\right)$. So let $x_i \in \gamma_i * \widetilde{U}$ such that $\pi(x_1) = \pi(x_2)$. Then $\exists\ \gamma \in \pi_1(C)$ such that $\gamma x_1 = x_2$. If we develop these points we see that

$$dev(x_2) = dev(\gamma x_1) = \rho(\gamma)dev(x_1)$$

and that $dev(x_2) \in dev(\gamma_2 * \widetilde{U}) = \rho(\gamma_2)dev(\widetilde{U})$ and $\rho(\gamma)dev(x_1) \in \rho(\gamma)dev(\gamma_1 * \widetilde{U}) = \rho(\gamma\gamma_1)dev(\widetilde{U})$. Since we already know that U is tame, we can conclude that $\rho(\gamma\gamma_1) = \rho(\gamma_2)$, hence that $\gamma\gamma_1 = \gamma_2$ since quasi-Fuchsian representations are faithful. But then we have that

$$\gamma_2 * \widetilde{U} = (\gamma\gamma_1) * \widetilde{U} = dev_{|_C}^{-1}\left(dev\left(\gamma\gamma_1\widetilde{U}\right)\right) = dev_{|_C}^{-1}\left(\rho(\gamma)dev\left(\gamma_1\widetilde{U}\right)\right)$$

The last term is indeed equal to $\gamma dev_{|_C}^{-1}\left(dev\left(\gamma_1\widetilde{U}\right)\right)$, because $\gamma \in \pi_1(C)$. So we have proved that $\gamma_2 * \widetilde{U} = \gamma\left(\gamma_1 * \widetilde{U}\right)$ for $\gamma \in \pi_1(C)$, which of course implies that $\pi\left(\gamma_1 * \widetilde{U}\right) = \pi\left(\gamma_2 * \widetilde{U}\right)$ as desired. □

Notice that the proof above shows that in the collection $dev_{|_C}^{-1}(dev(\pi_1(S)\widetilde{U}))$ either two elements differ by an automorphism of the universal cover $\pi : \widetilde{C} \to C$ and project to the same set on C, or they project to disjoint sets on C. In other words the avatars of U in C can be labelled by the cosets of $\pi_1(C)$ in $\pi_1(S)$; the index of $\pi_1(C)$ in $\pi_1(S)$ is 1 in the case of the uniformizing structure (where there are no non-trivial avatars, as already observed), and infinite otherwise, because in all the other cases any geometric component is a non compact (incompressible) subsurface and free subgroups of surface groups have infinite index. We conclude with the following technical lemma which says that it is always possible to nicely isotope a bubbleable arc in order to minimise its intersections with a sufficiently small neighbourhood of its endpoints.

Lemma 3.3.9. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian and $\sigma \in \mathcal{M}_{0,\rho}$. Let $\beta \subset \sigma$ be a bubbleable arc with endpoints x, y, with $x \notin \sigma^{\mathbb{R}}$. Let $U \subset \sigma$ be a connected π_1-trivial neighbourhood of x with $diam(U) < sys(\rho)$ and not containing any avatar of y. Then there is an injectively developed isotopy (relative to x and y) from β to another bubbleable arc β', such that β' does not intersect any non-trivial avatar of U and $\beta' \cap U$ is connected (i.e. β' does not come back to U after the first time it leaves it).*

Proof. First of all notice that if U does not contain avatars of y, then in particular y is not an avatar of x. Moreover no avatar of U contains avatars of y; in particular no avatar of U contains y. We also know by 3.3.8 that U is geometric (i.e it avoids the real curve), tame and its avatars are disjoint. Since U is geometric, for $\varepsilon > 0$ small enough the ε-neighbourhood $\mathcal{N}_\varepsilon(U)$ of U satisfies the same properties.

Let $\{U_i\}_{i\in I}$ be the collection of avatars of U crossed by β. Going along β from x to y we see that, apart from the initial segment starting at x inside U, every time β enters one of the U_i's it crosses it and leaves it (this is exactly because no avatar contains the second endpoint y). Therefore we can isotope all the arcs given by $\beta \cap U_i$ to arcs living in $\mathcal{N}_\varepsilon(U_i) \setminus U_i$, for each $i \in I$, without touching the first segment starting at x; since the chosen neighbourhood is tame this can be done in such a way that the isotopy is injectively developed. Since all the $\mathcal{N}_\varepsilon(U_i)$ are disjoint, this gives an isotopy on σ from β to an arc β' which intersects the whole collection of avatars only in the initial segment starting at x in U. It is still a bubbleable arc because it coincides with β (which is bubbleable) outside the $\mathcal{N}_\varepsilon(U_i)$'s, and the deformations inside these sets do not produce any new intersection because $\mathcal{N}_\varepsilon(U)$ is tame. □

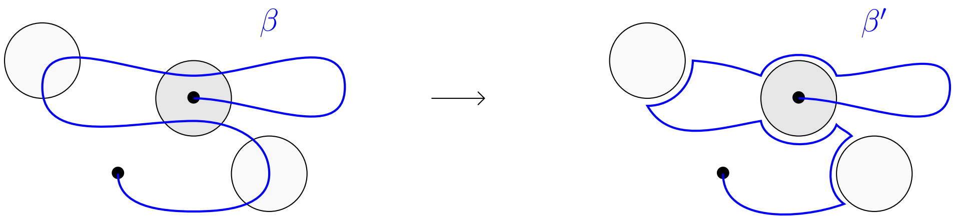

Figure 3.4: Pushing an arc outside the avatars of a neighbourhood of one endpoint.

To get a visual intuition of what can go wrong if we are not able to verify the hypothesis that avatars of U are disjoint and do not contain avatars of y, just consider what happens in this picture if the second endpoint y belongs to one of the avatars: there is no guarantee that the deformation that we want to perform is a bubbleable isotopy.

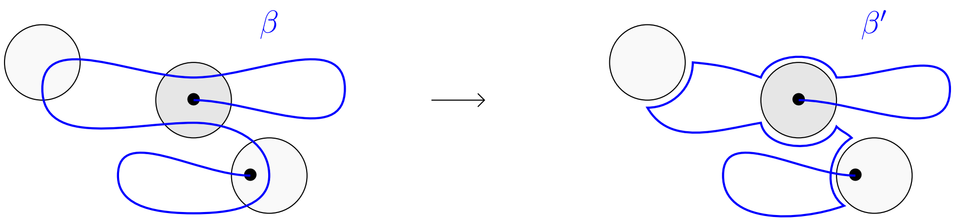

Figure 3.5: Avoiding avatars may result in self-intersections.

3.4 Visible BM-configurations

This section is about a class of BM-configurations with the property that, roughly speaking, the embedded twin pair survives after debubbling the structure; these should be thought as a strict generalisation of standard BM-configurations, which can still be dealt with by exploiting the underlying unbranched structure, where deformations are more easily defined and controlled. Once more, we give the main definitions in the general case, then specialise to the quasi-Fuchsian case.

Definition 3.4.1. Let $\sigma_0 \in \mathcal{M}_{0,\rho}$, $\beta \subset \sigma_0$ a bubbleable arc and $\sigma = Bub(\sigma_0, \beta) \in \mathcal{M}_{2,\rho}$ with distinguished bubble B coming from β. Let p be a branch point of σ and μ an embedded twin pair based at p with developed image $\widehat{\mu}$. Notice that the germ of μ is well-defined on σ_0. We say that the BM-configuration (B, μ, p) is a **visible BM-configuration** if we can take the analytic continuation of this germ on σ_0 to obtain a properly embedded path μ_0 on σ_0 which develops to $\widehat{\mu}$.

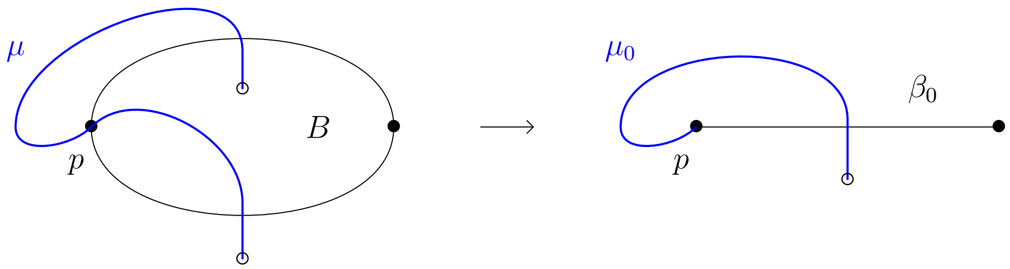

Figure 3.6: A visible BM-configuration.

Example 3.4.2. A standard BM-configuration is visible: as already observed in 3.2.3, the boundary of the bubble and the embedded twin pair of a standard BM-configuration induce a pair of adjacent embedded arcs on the debubbled structure, whose concatenation is actually a bubbleable arc itself.

Example 3.4.3. If σ is a standard bubbling over a hyperbolic surface and we pick an embedded twin pair μ which intersects the real curve, then the resulting configuration is not visible: the debubbled structure is the uniformizing hyperbolic structure, which has no real curve, so there can be no path on it developing as needed; the analytic continuation of the germ of μ is a geodesic which wraps around the surface without converging to a compact embedded arc.

The next result shows that (under suitable conditions) visible BM-configurations can be deformed to standard BM-configurations in a controlled way.

Proposition 3.4.4. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian, $\sigma_0 \in \mathcal{M}_{0,\rho}$, $\beta \subset \sigma_0$ a bubbleable arc. Let x, y be the branch points of $\sigma = Bub(\sigma_0, \beta)$ and B the bubble coming from β. Assume σ is simply developed and $x \notin \sigma^{\mathbb{R}}$. Let $K = \inf_{\gamma \in \pi_1(S)} d(dev(x), \rho(\gamma)dev(y)) \in]0, +\infty]$ and let μ be an embedded twin pair based at x such that (B, μ, x) is a visible BM-configuration and with length $l(\mu) < \min\{sys(\rho), K\}$. Then there is another bubble $B' \subset \sigma$ such that $Debub(\sigma, B') = \sigma_0$ and (B', μ, x) is a standard BM-configuration.*

Proof. Since the BM-configuration is visible, after debubbling σ we can define an arc μ_0 on σ_0 starting at x and developing as μ. By hypothesis this arc is shorter than $sys(\rho)$ and K; in particular it can be put inside a connected contractible neighbourhood U of x with $diam(U) < sys(\rho)$ and which does not contain any avatar of y. By 3.3.9 there is an injectively developed isotopy from β to a bubbleable arc β' which avoids all non trivial avatars of U and intersects U just once at the starting segment at x. Since this isotopy is injectively developed, bubbling σ_0 along β' gives a structure isomorphic to σ by 1.4.11. Moreover the fact that $\mu \subset U$ and that β' avoids all non trivial avatars of U and does not come back to it after the first time it leaves it implies that the concatenation of μ and β' is a bubbleable arc; this is equivalent to saying that the resulting BM-configuration (B', μ, x) is standard by the characterisation in 3.2.3. $\square$

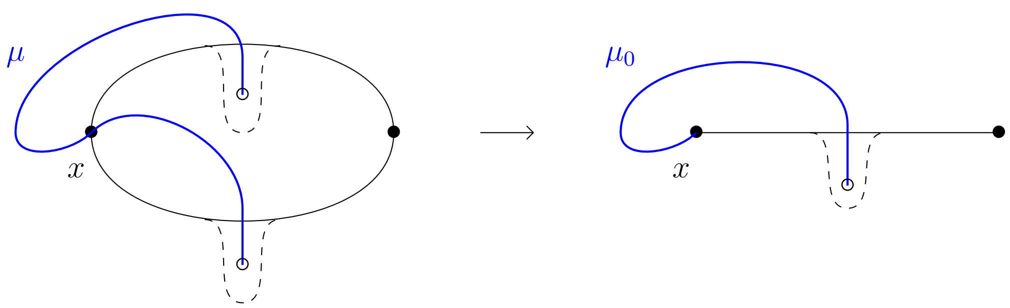

Figure 3.7: Deforming a visible BM-configuration into a standard one.

Remark 3.4.5. The above result means that moving branch points on a given bubbling by a very small displacement (with respect to the representation) does preserve the [bubbleable isotopy class of the] given bubble. In particular the underlying unbranched structure can be left unchanged throughout the movement. Notice that the hypothesis on the constant K is indeed necessary: namely we can bubble a hyperbolic surface and then move one branch point to obtain a doubly developed structure, which can not be a bubbling of the hyperbolic surface, as already mentioned in 2.3.9.

The condition of being visible is a bit obscure, if compared to that of being standard, in the sense that we have to debubble the structure to check visibility, and we do not have a simple characterisation as the one in 3.2.3 for standard BM-configurations; but visibility is always at least locally available at geometric branch points, as shown by the following result.

Lemma 3.4.6. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be quasi-Fuchsian, $\sigma_0 \in \mathcal{M}_{0,\rho}$, $\beta \subset \sigma_0$ a bubbleable arc such that $\sigma = Bub(\sigma_0, \beta) \in \mathcal{M}_{2,\rho}$ has a branch point x not on the real curve. Let μ be an embedded twin pair based at x of length $l(\mu) < sys(\rho)$. Then the resulting BM-configuration is visible.*

Proof. Let us fix a developed image $\widehat{x}$ of x and $\widehat{\mu}$ for μ. Since $l(\widehat{\mu}) < sys(\rho)$, it is contained in a fundamental domain for ρ, and a fortiori in a fundamental domain for

$\rho_{|H}$, for any subgroup $H \leq \pi_1(S)$. Therefore it projects injectively to every quotient $\mathbb{H}^2/\rho(H)$. Now consider the debubbled structure σ_0, and let C be the geometric component of σ_0 containing x; since σ_0 is unbranched, C is incompressible and the developing map induces an isometry $D_C : C \to C_{Fuchs} = \mathbb{H}^2/\rho(\pi_1(C))$, where $\widehat{\mu}$ projects injectively to an embedded arc. Pulling that arc back by D_C gives the desired arc on σ_0, which proves the visibility of the BM-configuration. □

Remark 3.4.7. We want to remark that it is not possible to apply these ideas to a movement of a branch point which sits on the real curve. Indeed, here geometricity is used to produce neighbourhoods of the relevant objects which have disjoint avatars. On the other hand if a point belongs to the real curve, then any of its neighbourhoods will contain infinitely many avatars of both branch points, and actually of whatever object we want to consider. This follows from the fact that if Γ is a Fuchsian group, then the collection of fixed points of its hyperbolic elements is dense in the limit set $\Lambda_\Gamma = \mathbb{RP}^1$. Therefore if $x \in \sigma^{\mathbb{R}}$ and $U \subset \sigma$ is an open neighbourhood of x, then there exists some $\gamma \in \pi_1(S)$ such that the attracting fix point of $\rho(\gamma)$ lies in $dev(U) \cap \mathbb{RP}^1$; the $\rho(\gamma)$-orbit of the developed image of any subset $V \subset \sigma$ will converge to that point, hence definitely intersect $dev(U)$, which of course implies that we see infinitely many avatars of V inside U.

3.5 Bubbles everywhere

In this section we prove that bubblings our quite ubiquitous in $\mathcal{M}_{2,\rho}$. The strategy is to show first that if a geometric component of the real stratification (i.e. a connected component of a space with fixed combinatorics $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$) contains a bubbling, then the subspace of bubblings is actually dense in it. Then we prove that every cell is adjacent to one which contains a bubbling, and that this is enough to obtain that actually every cell contains a bubbling. What is left out of this approach are the subspace of really branched structures $\mathcal{M}_{2,\rho}^{\mathbb{R}}$ and the subspace of non simply developed structures $\mathcal{M}_{2,\rho} \setminus \mathcal{SM}_{2,\rho}$, which have real codimension 1 and 2 respectively; of course both of them have non trivial intersection with the space of bubblings, i.e. this result can be improved. Recall from 3.3.6 that a structure with two branch points is simply developed if and only if the branch points are not avatars, and that it is geometrically branched if all branch points are outside the real curve, by definition.

Theorem 3.5.1. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian. Let $\mathcal{X} = \{\sigma \in \mathcal{M}_{2,\rho} \mid \sigma$ is simply developed and geometrically branched$\}$ and $\mathcal{Y}$ a connected component of $\mathcal{X}$ containing a bubbling. Then every structure in $\mathcal{Y}$ is a bubbling.*

Proof. As usual let us denote by $\mathcal{BY}$ the subspace of $\mathcal{Y}$ made of bubblings. By hypothesis is it non empty; moreover it is open in $\mathcal{Y}$ because $\mathcal{BY} = \mathcal{Y} \cap \mathcal{BM}_{2,\rho}$ and the second one is open in $\mathcal{M}_{2,\rho}$ by 3.2.5. We will prove that $\mathcal{BY}$ is also closed in $\mathcal{Y}$ and conclude by connectedness of $\mathcal{Y}$.

Let $\sigma_\infty \in \mathcal{Y} \cap \overline{\mathcal{BY}}$. By hypothesis the branch points x_∞ and y_∞ of σ_∞ are outside the real curve of σ_∞ and not avatar of each other. Fix any developed image $\widehat{x}_\infty$ of x_∞ and $\widehat{y}_\infty$ of y_∞; then define $K_\infty = \inf_{\gamma \in \pi_1(S)} d(\widehat{x}_\infty, \rho(\gamma)\widehat{y}_\infty)$. The distance here is the one induced by the hyperbolic metrics on the domain of discontinuity of ρ; K_∞ is strictly positive since the branch points of σ_∞ are not avatars, but can be $+\infty$ in the case they have opposite sign. Then let $A = \min\{sys(\rho), \frac{1}{3}K_\infty\}$.

Choose $L < A$ and consider the neighbourhood $\mathcal{N}_L(\sigma_\infty)$ of σ_∞ in $\mathcal{Y}$ obtained by moving branch points by a distance $L < A$ (this is well defined since σ_∞ is geometrically branched; also notice that neighbourhoods in $\mathcal{Y}$ are the same thing as neighbourhood in $\mathcal{M}_{2,\rho}$ since $\mathcal{Y}$ is an open submanifold). Since σ_∞ is in the closure of $\mathcal{BY}$, $\mathcal{N}_L(\sigma_\infty)$ will contain a bubbling $\sigma \in \mathcal{BY}$. Let ζ be an embedded twin pair based at x_∞ and ξ be an embedded twin pair based at y_∞ such that $Move(\sigma_\infty, \zeta, \xi) = \sigma$. By construction they can be chosen to have length smaller than L. Let $\widehat{\zeta}$ and $\widehat{\xi}$ be the developed images based at $\widehat{x}_\infty, \widehat{y}_\infty$, and let $\sigma_0 \in \mathcal{M}_{0,\rho}$ and $\beta \subset \sigma_0$ be such that $\sigma = Bub(\sigma_0, \beta)$. Also let x, y be the branch points of σ corresponding to x_∞, y_∞ respectively.

We are now going to show that σ_∞ is actually a bubbling over the same σ_0. First of all notice that by 3.4.6 both BM-configurations are visible, since both embedded twin pairs are shorter than the systole of the representation. Moreover, by definition of A, the two movements are independent from each other; more precisely they do not interfere with each other in the sense that each twin pair avoids all avatars of the other twin pair. We begin by focusing at x; let us denote by $\widehat{x}, \widehat{y}$ the developed images of x, y which are seen at the endpoints of $\widehat{\zeta}$ and $\widehat{\xi}$. We have that for any $\gamma \in \pi_1(S)$

$$K_\infty \leq d(\widehat{x}_\infty, \rho(\gamma)\widehat{y}_\infty) \leq d(\widehat{x}_\infty, \widehat{x}) + d(\widehat{x}, \rho(\gamma)\widehat{y}) + d(\rho(\gamma)\widehat{y}, \rho(\gamma)\widehat{y}_\infty) = 2L + d(\widehat{x}, \rho(\gamma)\widehat{y})$$

so that

$$d(\widehat{x}, \rho(\gamma)\widehat{y}) \geq K_\infty - 2L$$

and we get

$$\inf_{\gamma \in \pi_1(S)} (d(\widehat{x}, \rho(\gamma)\widehat{y})) \geq K_\infty - 2L > 3L - 2L = L = l(\widehat{\zeta})$$

by definition of A. Therefore we can apply 3.4.4 and replace β by a new (but isotopic) bubbleable arc which is in standard BM-configuration on σ with respect to ζ. We now let $\sigma' = Move(\sigma, \zeta)$, which is still a bubbling over σ_0 by 3.2.4.

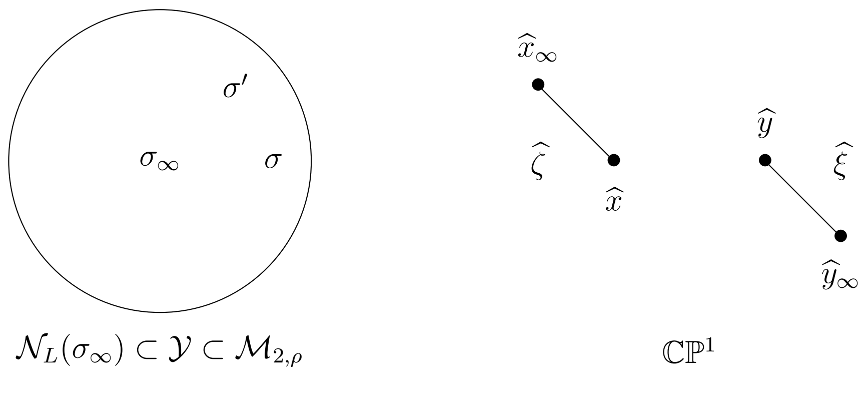

Figure 3.8: The neighbourhood $\mathcal{N}_L(\sigma_\infty)$ and the movement of points in $\mathbb{CP}^1$.

We now want to play the same strategy again at y to get back to σ_∞; to do so, we just have to check that the movement is small enough with respect to the distance between the two branch points of σ', which now develop to $\widehat{x}_\infty$ and $\widehat{y}$. But this is easily checked: if $\gamma \in \pi_1(S)$ then

$$K_\infty \leq d(\widehat{x}_\infty, \rho(\gamma)\widehat{y}_\infty) \leq d(\widehat{x}_\infty, \rho(\gamma)\widehat{y}) + d(\rho(\gamma)\widehat{y}, \rho(\gamma)\widehat{y}_\infty) = L + d(\widehat{x}_\infty, \rho(\gamma)\widehat{y})$$

so that

$$d(\widehat{x}_\infty, \rho(\gamma)\widehat{y}) \geq K_\infty - L$$

and we get

$$\inf_{\gamma \in \pi_1(S)} (d(\widehat{x}_\infty, \rho(\gamma)\widehat{y})) \geq K_\infty - L > 3L - L > L = l(\widehat{\xi})$$

So we can apply 3.4.4 again and replace the bubbleable arc with an isotopic one which is in standard BM-configuration and safely move branch points along ξ. This movement results in our structure σ_∞ and does not break the bubble by 3.2.4. In other words this proves that $\sigma_\infty \in \mathcal{BY}$ (and indeed the underlying unbranched structure is the same as that of σ and σ'), so that $\mathcal{BY}$ is closed. By connectedness we get $\mathcal{Y} = \mathcal{BY}$, i.e. $\mathcal{Y}$ is entirely made of bubblings. □

We now apply this result to geometric components of the real decomposition of $\mathcal{M}_{2,\rho}$, i.e. connected components of the space $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$ for some combinatorics $\mathfrak{c}$. Notice that structures in a geometric component are geometrically branched and recall that, given a combinatorics, we denote by $k^\pm$ the number of positive and negative branch points. As far as the sign of branch points is concerned, for $k = 2$ we have three types of combinatorics, namely $k^+ = 0, 1, 2$. At this point, one of them is easily dealt with.

Corollary 3.5.2. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian. Let $\mathcal{C}$ be a geometric component of the real decomposition of $\mathcal{M}_{2,\rho}$ whose combinatorics has $k^+ = 1$. Then $\mathcal{C}$ is entirely made of bubblings.*

Proof. First observe that $\mathcal{C}$ does not contain doubly developed structures, since branch points have different sign, hence $\mathcal{C} = \mathcal{SC}$. Let $\sigma \in \mathcal{C}$. By 2.4.5 it satisfies the hypothesis of [8, Theorem 7.1]. Therefore it is possible to move branch points inside their own geometric components so that a bubble appears, which proves that $\mathcal{C}$ contains a bubbling. Now apply 3.5.1 with $\mathcal{Y} = \mathcal{C} = \mathcal{SC}$. □

We now have to care about geometric components of the real decomposition of $\mathcal{M}_{2,\rho}$ whose combinatorics has $k^+ = 0, 2$. Let us fix such a combinatorics $\mathfrak{c}$ and a connected component $\mathcal{X}$ of $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$, and prove that $\mathcal{X}$ actually contains a bubbling. The results in [8, Proposition 8.1, Lemma 10.5-6] imply that in some cases points can be moved inside their own geometric components so that a bubble appears, but it is not clear how to verify a priori when this occurs. On the other hand we know by the discussion in 2.4.1 that all combinatorics can be realised by bubbling suitable unbranched structures, which means that fixing a combinatorics does not provide any obstruction to the existence of bubbles; but here we are not just looking for a bubbling in $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$: indeed we look for a bubbling in the chosen component $\mathcal{X}$, and we have already remarked in 2.1.22 that the spaces $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$ are quite far from being connected. The strategy will be to look for bubblings in the components adjacent to $\mathcal{X}$ and drag them from there back into $\mathcal{X}$. In trying to do so, two problems occur: on one side if we naively take a bubbling in some component adjacent to $\mathcal{X}$ and move branch points on it beyond the real curve, then it is quite difficult to control that we are actually moving to the chosen component $\mathcal{X}$; on the other hand if we start with a structure $\sigma \in \mathcal{X}$ and move branch points on it across the real curve to get to a bubbling, then it is quite difficult to check that when we move branch points to get back to σ we do not break the bubble. Some lemmas are in order to guarantee that we can handle these issues. Recall that a path α on a quasi-Fuchsian BPS σ is said to be geodesic if $\alpha \cap \sigma^{\pm}$ is a collection of geodesics.

Lemma 3.5.3. *Let $\mathfrak{c}$ be a combinatorics with $k^+ = 0$ or 2, $\mathcal{X}$ a connected component of $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$ and $\sigma \in \mathcal{X}$. Then there exists a combinatorics $\mathfrak{a}$ with $k^+ = 1$, a connected component $\mathcal{Y}$ of $\mathcal{M}_{2,\rho}^{\mathfrak{a}}$ and a geodesic embedded twin pair μ on σ such that $Move(\sigma, \mu) \in \mathcal{Y}$. In particular $\mathcal{X}$ and $\mathcal{Y}$ are adjacent in $\mathcal{M}_{2,\rho}$.*

Proof. This is just a reformulation of the results in [8, §9], which say that it is always possible to move a branch point along a geodesic embedded twin pair reaching the real curve. □

Even if we will not need it, notice that in general there will be many different components adjacent to $\mathcal{X}$ to which we can geodesically move, possibly with many different combinatorics. We remark that in the process of moving a branch point towards the real curve with the techniques of [8, §9] a bubble might appear before actually crossing the real curve; this would be fine for us, since our ultimate goal now is to prove that $\mathcal{X}$ contains a bubbling; therefore we will forget about this detail in the following.

Of course it is interesting to move to a component whose combinatorics has $k^+ = 1$ because we already know by 3.5.2 that such a component $\mathcal{Y}$ is entirely made of bubblings, so we have a lot of freedom in the choice of the bubble. We can obtain a satisfying level of freedom also in the choice of the embedded twin pairs to be used in the movements of the branch point as indicated by the following lemma.

Lemma 3.5.4. *Let $\sigma \in \mathcal{M}_{2,\rho}$ be a geometrically branched BPS and $\mu = \{\mu_1, \mu_2\}$ an embedded twin pair on σ. Suppose that μ_i crosses $\sigma^{\mathbb{R}}$ at only one point r_i. Then there exists a geodesic embedded twin pair ν on σ such that $Move(\sigma, \mu) = Move(\sigma, \nu)$.*

Proof. Let p be the base point of the embedded twin pair μ and y_i be the endpoint of μ_i. By hypothesis the subarcs $\mu_i^1 \subset \mu_i$ from p to r_i are entirely contained in a geometric component C. We let ν_i^1 be the unique geodesic in C from p to r_i which is isotopic to μ_i^1 relatively to $\{p, r_i\}$. Then we can do the same in the adjacent components to obtain geodesic arcs ν_i^2 isotopic to the subarcs $\mu_i^2 \subset \mu_i$ from r_i to y_i. The concatenation of these paths gives rise to a couple of geodesic paths ν_i from p to y_i which are isotopic to μ_i relatively to $\{p, r_i, y_i\}$. Each geometric subarc ν_i is the geodesic representative of the embedded and injectively developed arc μ_i, hence it is embedded and injectively developed; moreover the two geometric subarcs of ν_i live in two adjacent component, hence their developed images are disjoint and ν is thus actually a geodesic embedded twin pair. The isotopy from μ to ν can be chosen to be an isotopy of embedded twin pairs, so that 1.4.26 applies and gives us that $Move(\sigma, \mu) = Move(\sigma, \nu)$. ☐

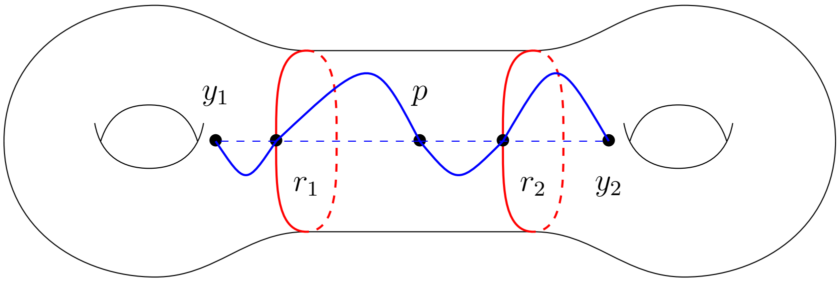

Figure 3.9: Straightening the embedded twin pair.

Notice that this result does not hold for paths which cross more components: subarcs of ν contained in geometric components of the same sign could overlap once developed, even if μ is an embedded twin pair.

We are now ready to prove that all the pieces of the real decomposition of $\mathcal{M}_{2,\rho}$ contain a bubbling.

Proposition 3.5.5. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a quasi-Fuchsian representation. Let $\mathcal{X}$ be a connected component of the real decomposition of $\mathcal{M}_{2,\rho}$. Then $\mathcal{X}$ contains at least one structure which is a bubbling over an unbranched structure in $\mathcal{M}_{0,\rho}$.*

Proof. Let $\mathfrak{c}$ be the combinatorics of $\mathcal{X}$. If it has $k^+ = 1$ then this follows directly from 3.5.2. So let us assume that $k^+ = 2$, the case $k^+ = 0$ being the same up to switching the signs of the branch points. We start by selecting an adjacent component in which we know how to find bubbles. To do this we use 3.5.3: choose some $\sigma_1 \in \mathcal{X}$ and move branch points along an embedded twin pair μ to get to a structure $\sigma_2 = Move(\sigma_1, \mu)$ in some adjacent component $\mathcal{Y} \subset \mathcal{M}_{2,\rho}^{\mathfrak{a}}$ for some combinatorics $\mathfrak{a}$ with exactly one positive and one negative branch point. By 2.4.5 we know the combinatorial properties of the geometric decomposition of σ_2: all real curves are essential, one has index 1 and the others have index 0. Let us call l the unique real curve of index 1; the branch points $p^\pm$ live in the two geometric components $C^\pm$ adjacent to l. By construction we have an induced embedded twin pair ν at p^- on σ_2 such that $Move(\sigma_2, \nu) = \sigma_1$. Here we know by [8, Theorem 7.1] that we can move both branch points inside their own components to get to a structure $\sigma_3 \in \mathcal{Y}$ such that the peripheral geodesics of l go through the branch points $q^\pm$ of σ_3 with angles $\{\pi, 3\pi\}$ and also such that it has a geodesic bubble B (such a bubble can indeed be chosen in many ways, which will be exploited below). Of course we have an induced couple of embedded twin pairs $\zeta^\pm \subset \sigma_3$ based at $q^\pm$ such that $Move(\sigma_3, \zeta^+, \zeta^-) = \sigma_1$, and we would like to operate this movement of branch points on σ_3 without breaking the bubble B; unfortunately there is no reason why $(B, \zeta^\pm, q^\pm)$ should be a standard BM-configuration.

However for our purposes we do not actually need to move branch points to go back to σ_1: it is enough to move to a structure in the same component $\mathcal{X}$ without breaking the bubble B. Therefore we can forget about the embedded twin pair ζ^+,

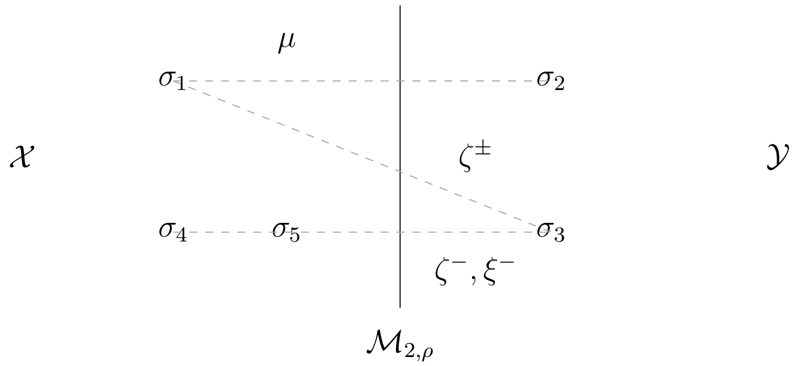

Figure 3.10: The structures $\sigma_1, \sigma_2, \sigma_3, \sigma_4$ and σ_5 involved in the proof.

since we only need to move q^- to go back to that component. Since ζ^- crosses the real curve just once, by 3.5.4 we can replace it with a geodesic embedded twin pair ξ^- which is such that $\sigma_4 = Move(\sigma_3, \xi^-) = Move(\sigma_3, \zeta^-) \in \mathcal{X}$. As mentioned above, the bubble B on σ_3 can be chosen in a quite free way, and our aim now is to prove that it is always possible to choose the bubble so that the BM-configuration (B, ξ^-_{cut}, q^-) is standard, for some suitable truncation ξ^-_{cut} of the embedded twin pair ξ^-; of course we still have that $\sigma_5 = Move(\sigma_3, \xi^-_{cut}) \in \mathcal{X}$.

First of all we recall from [8, §7] that the real curve l carries a natural action of the infinite cyclic group generated by $\rho(l)$ and a natural $\rho(l)$-invariant decomposition $l = \{0\} \cup l^+ \cup \{\infty\} \cup l^-$, corresponding to the decomposition of the limit set of ρ given by the fixed points of $\rho(l)$; according to [8, Proposition 7.8] for any $u \in l^+$ we can find a geodesic bubble B_u intersecting l exactly at u and $\rho(l)^{-1}u$. Suppose we pick one of these geodesic bubbles B_u and look at the situation on C^-, neglecting for a moment what happens beyond the real curve l. Since the embedded twin pair ξ^- and the bubble B_u are both geodesic, when one of the paths of ξ^- enters the bubble it can never leave it, and must reach the real curve l. One of them, let us say ξ^-_1 starts inside B_u (up to an arbitrarily small displacement of u), hence hits l at some point v_1. If the BM-configuration (B_u, ξ^-, q^-) is not already standard, it means that the twin ξ^-_2 starting outside B_u goes somewhere around the surface and then comes back to intersect B_u at some point x, and finally hits the real curve l at some point v_2, distinct from v_1, because ξ^- is an embedded twin pair. Now, let us show that v_2 must live in l^+. To do this, we choose u so that the bubble B_u is orthogonal at q^- to the peripheral geodesic of l. Since ξ^-_2 is a geodesic from q^- to l, once it enters the end relative to l it constantly increases its distance from the peripheral geodesic; in particular, when it intersects the bubble at x it forms an angle smaller than $\frac{\pi}{2}$ with the boundary of B_u. Since u is in l^+, this forces $v_2 \in l^+$ as well. But then it is now possible to choose a different u' in such a way that the arc $\alpha \subset l$ from u' to $\rho(l)^{-1}u'$ containing 0 and ∞ (i.e. the part of l contained in $B_{u'}$) does not contain v_2. This choice guarantees that v_2 is outside the bubble $B_{u'}$, hence that ξ^-_2 does not intersect $B_{u'}$ before crossing the real curve l. We have no tools to control what happens beyond l, but we can truncate ξ^- to a sub-embedded twin pair ξ^-_{cut} which ends beyond l and which is in standard BM-configuration with respect to the bubble $B_{u'}$. By 3.2.4 $\sigma_5 = Move(\sigma_3, \xi^-_{cut})$ is still a bubbling. But we can clearly

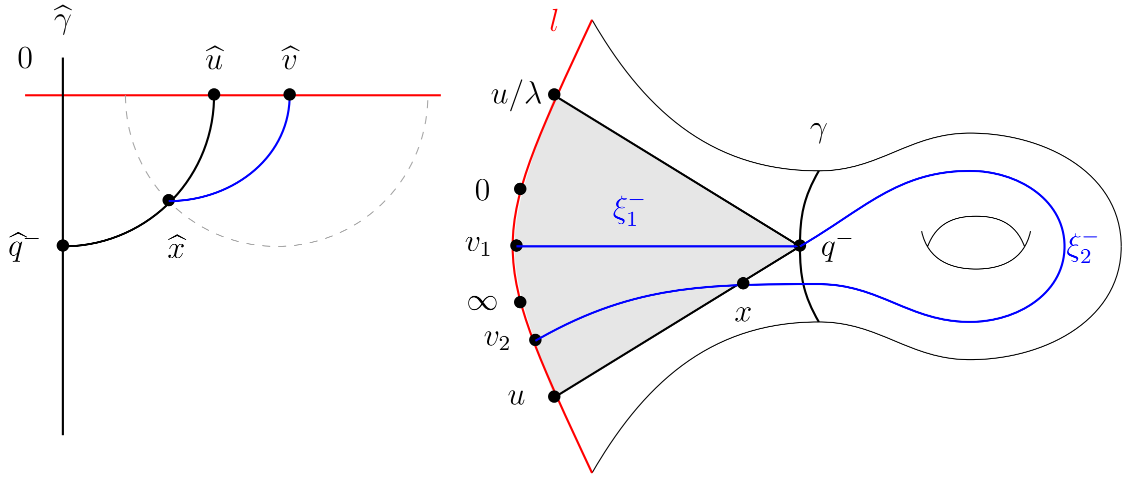

Figure 3.11: The configuration in $\mathbb{CP}^1$ and $C^- \subset \sigma_3$ when B_u is the bubble orthogonal to the peripheral geodesic.

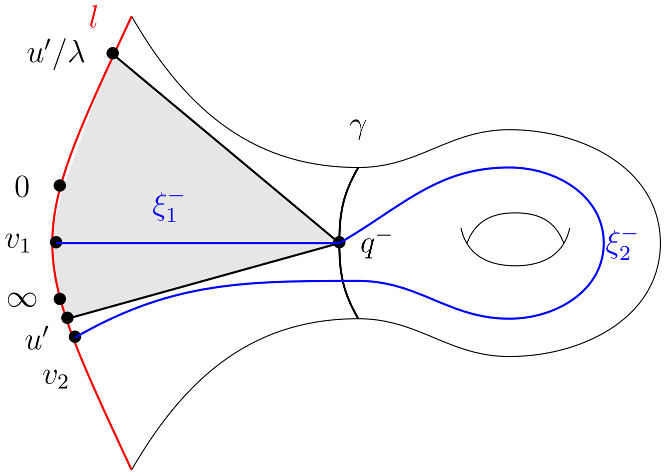

Figure 3.12: A bubble in standard BM-configuration.

keep moving branch points on σ_5 along what is left of ξ^- to reach the structure $\sigma_4 = Move(\sigma_3, \xi^-)$, which, as we already know, lives in the same component $\mathcal{X}$ containing σ_1. Since this movement does not cross the real curve, the structure σ_5 lives in $\mathcal{X}$ too, which proves that $\mathcal{X}$ contains a bubbling. $\square$

We can finally state the main result. Recall that a structure σ is geometrically branched if its branch points are outside the real curve $\sigma^{\mathbb{R}}$, and that it is simply developed if for every branched component $C \subset \sigma^{\pm}$ the canonical map to the Fuchsian model $D_C : C \to C_{Fuchs}$ is injective on the collection of branch points.

Theorem 3.5.6. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian. Then any simply developed structure with at most one real branch point is a bubbling. In particular the space of bubblings is a connected, open and dense submanifold of $\mathcal{M}_{2,\rho}$ of full measure.*

Proof. At first, let σ be a geometrically branched and simply developed structure. Since its branch points are outside the real curve it belongs to some space $\mathcal{M}_{2,\rho}^{\mathfrak{c}}$ for some combinatorics . Let $\mathcal{X} \subset \mathcal{M}_{2,\rho}^{\mathfrak{c}}$ be the connected component containing σ. By 3.5.5 we know that $\mathcal{X}$ contains a bubbling. Moreover σ avoids the subspace of $\mathcal{X}$ made of non simply developed structures, i.e. $\sigma \in \mathcal{SX}$, which is still connected. Then by 3.5.1 σ is a bubbling, since any structure in $\mathcal{SX}$ is. In the case σ has one real branch point, we can perform a movement of that branch point to go from σ to some structure σ' in some geometric piece of the real decomposition with $k^+ = 1$. Then the previous arguments apply verbatim, because the isotopy in 3.5.4 fixes the points of intersection between the embedded twin pair and the real curve, so that we are able to pick a bubble on σ' and move back to σ as in 3.5.5.

To prove the last claim, recall that by [8] $\mathcal{M}_{2,\rho}$ is a connected manifold of real dimension 4 and observe that the subspace of structures left outside by this approach is the union of the subspace of non simply developed structures and the one of structures with both branch points on the real curve; the first one is a submanifold of real dimension 2, whereas the second one is expressed in local charts by the image of the limit set of ρ in the symmetric product $Sym^2(\mathbb{CP}^1)$: in the Fuchsian case it is a real analytic subspace, and in the general quasi-Fuchsian case it is a closed subset of measure zero (with respect to the Lebesgue measure class for the manifold structure of $\mathcal{M}_{2,\rho}$).

□

As a consequence we get a generically positive answer in our setting to the question asked by Gallo-Kapovich-Marden as Problem 12.1.2 in [13], i.e. if any two BPS with the same holonomy are related by a sequence of grafting, degrafting, bubbling and debubbling. More precisely our theorem shows that, if $\{\sigma, \tau\}$ is a generic couple of BPS with at most two branch points and a fixed quasi-Fuchsian holonomy, then we can apply one debubbling to each of them (if needed), to reduce to a couple of unbranched structures with the same holonomy $\{\sigma_0, \tau_0\}$. By Goldman's theorem in [16] we can then apply m degraftings on σ_0 to obtain the uniformizing structure σ_ρ and then n graftings on σ_ρ to obtain τ_0, for suitable $m, n \in \mathbb{N}$.

$$\sigma \xrightarrow{\text{1 debub}} \sigma_0 \xrightarrow{m \text{ degraft}} \sigma_\rho \xrightarrow{n \text{ graft}} \tau_0 \xrightarrow{\text{1 bub}} \tau$$

Actually it is possible to do even better, since we can remove the need for degraftings; by the proof of [9, Theorem 11], there exists a simple closed geodesic γ on σ_ρ such that $\sigma_\gamma = Gr(\sigma_\rho, \gamma)$ can be obtained by m' graftings on σ_0 and τ_0 can be obtained by n' graftings on σ_γ, for suitable $m', n' \in \mathbb{N}$. Finally, according to [8,

$$\sigma \xrightarrow{\text{1 debub}} \sigma_0 \xrightarrow{m' \text{ graft}} \sigma_\gamma \xrightarrow{n' \text{ graft}} \tau_0 \xrightarrow{\text{1 bub}} \tau$$

Theorem 5.1] every simple grafting can be realised by a sequence of one bubbling and one debubbling; this implies the following, which shows that it is possible to move around the moduli space only via bubblings and debubblings.

Corollary 3.5.7. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian. There is an open dense subspace $\mathcal{B} \subset \mathcal{M}_{2,\rho}$ such that if $\sigma, \tau \in \mathcal{M}_{0,\rho} \cup \mathcal{B}$ then τ is obtained from σ by a finite sequence of bubblings and debubblings.*

$$\sigma \xrightarrow{\text{1 debub}} \sigma_0 \xrightarrow[m' \text{ debub}]{m' \text{ bub}} \sigma_\gamma \xrightarrow[n' \text{ debub}]{n' \text{ bub}} \tau_0 \xrightarrow{\text{1 bub}} \tau$$

Notice that the length of this sequence depends on the choice of the bubbles on σ, τ, i.e. of the unbranched structures σ_0, τ_0; as we have already suggested a BPS with two branch points can in general be realised as a bubbling over different unbranched structures along different arcs. This will be discussed in detail in the next section, where we will also improve 3.5.7 by providing an explicit uniform bound on the length of the above sequence.

3.6 Multi(de)grafting via bubbles

The idea of the previous sections has been to deform a bubble within its isotopy class, i.e. to work in the debubbled structure; now we consider the problem of changing the isotopy class, i.e. of changing the underlying unbranched structure. In this section we show how to remove a grafting region on a projective structure by bubbling and debubbling it, extending [8, Theorem 5.1] with dealt with the case of a simple grafting annulus. More precisely, given a bubbleable arc which crosses a grafting region we show how to find (in the bubbled structure) a different bubble whose boundary has strictly fewer intersections with the real curve.

To simplify the exposition we adopt here the convention that normal letters denote objects on the surface and letters with a hat denote a developed image of the corresponding object. For the same reason we will state and prove results for Fuchsian representations; everything extends to the quasi-Fuchsian case by replacing the hyperbolic plane by the positive component of the domain of discontinuity of ρ.

Let us fix a Fuchsian representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ and a projective structure σ with holonomy ρ. By [16] σ is obtained by a multigrafting on the uniformizing structure $\sigma_\rho = \mathbb{H}^2/\rho(\pi_1(S))$, hence it decomposes into a hyperbolic core of finite volume (coming directly from σ_ρ) plus a certain number of grafting regions. We will denote by $A_\gamma = A_\gamma^1 \cup \cdots \cup A_\gamma^M$ the grafting region obtained by grafting M times some simple closed geodesic γ of σ_ρ. Recall that in our terminology a grafting annulus is made of a negative annulus and also of a couple of ends in the adjacent positive component(s), see Picture 1.2. Notice that the structure on the interior of each grafting annulus is uniformizable, in the sense that the developing image is injective on the interior of the universal cover of the annulus.

Given a bubbleable arc β which crosses A_γ transversely from side to side, we introduce some auxiliary objects which will be useful in the main construction. For simplicity let us begin with the case in which the grafting is simple, i.e. $M = 1$, and come back later to the general case. Let us denote by γ_L, γ_R the two copies of γ appearing as the boundary of A_γ.

Definition 3.6.1. Let β be an oriented bubbleable arc properly embedded in A_γ (i.e. $\partial\beta = \beta \cap \partial A_\gamma$). We call I (in) and O (out) the two points of $\partial\beta$ at which β respectively enters in the annulus and leaves it. Notice that there is a unique point on ∂A_γ which is different from I but is developed to the same point $\widehat{I}$. We will refer to it as the twin of I, and similarly for O.

Definition 3.6.2. Let β be an oriented bubbleable arc properly embedded in A_γ. We can define a preferred orientation for γ so that in the developed image $\widehat{O}$ sits after $\widehat{I}$ along $\widehat{\gamma}$ (since β is bubbleable, $\widehat{I} \neq \widehat{O}$, thus this is well defined). We refer to it as the orientation of γ induced by β.

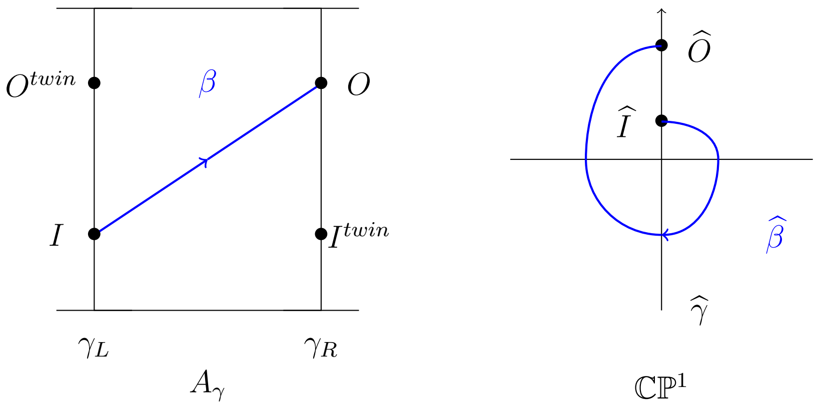

Figure 3.13: An arc inducing In and Out points and an orientation.

Definition 3.6.3. Let β, β' be oriented bubbleable arcs properly embedded in A_γ. Then we say that β' is coherent with β if $\widehat{O'}$ sits after $\widehat{I'}$ along $\widehat{\gamma}$ with respect to the orientation induced by β. Otherwise we say that it is incoherent.

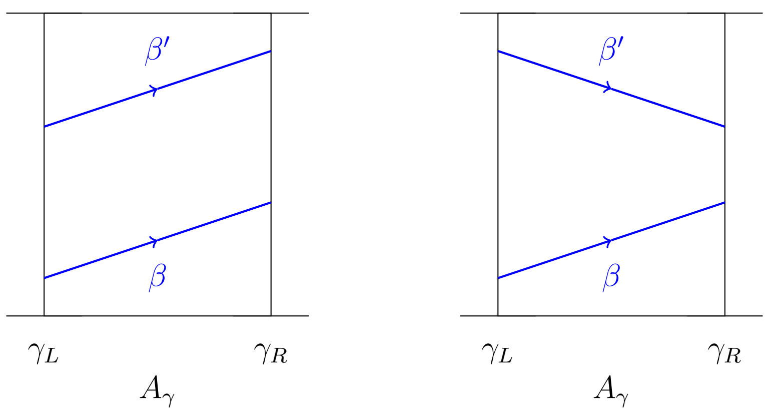

Figure 3.14: Coherent and incoherent arcs.

Now let β be an oriented bubbleable arc which transversely crosses some grafting annulus A_γ, i.e. every time it enters A_γ it crosses it and leaves it on the other side. Then $\beta \cap A_\gamma = \beta_1 \cup \cdots \cup \beta_N$ is a disjoint union of oriented bubbleable arcs properly embedded in A_γ, which we will call crossings; the labelling of $\beta_1, \dots, \beta_N$ is such that they appear in this order along β. For each crossing β_k we can define the entry and exit point I_k and O_k, and the induced orientation of γ as above. Since β is embedded and bubbleable all these points are distinct, and the same holds for their developed images. We agree to fix the orientation of γ determined by the first crossing β_1, but of course more generally we can also decide if two given crossings are coherent or not with respect to each other.

Let us introduce another useful way to order the crossings $\beta_1, \dots, \beta_N$, according to the way they appear when travelling along γ with respect to the orientation of γ induced by β_1: set $\alpha_1 = \beta_1$, then let α_{k+1} be the crossing we meet after α_k along γ with respect to the chosen orientation. We get an ordering of the crossings as $\alpha_1, \dots, \alpha_N$ which is actually a $\mathbb{Z}_N$-order (i.e. $\alpha_{N+1} = \alpha_1$); moreover there exists a unique permutation $\sigma \in \mathfrak{S}_N$ such that $\alpha_k = \beta_{\sigma(k)}$ and $\sigma(1) = 1$. We keep track of the coherence between crossings by defining the following **coherence parameters**

$$\varepsilon_k = \begin{cases} 1 & \text{if } \alpha_k \text{ coherent with } \alpha_1 = \beta_1 \\ -1 & \text{if } \alpha_k \text{ incoherent with } \alpha_1 = \beta_1 \end{cases}$$

$$\varepsilon_{k,l} = \begin{cases} 1 & \text{if } \alpha_k, \alpha_l \text{ coherent with each other} \\ -1 & \text{if } \alpha_k, \alpha_l \text{ incoherent with each other} \end{cases}$$

Let us roughly describe the idea behind the main construction of this section. Given a bubbleable arc which transversely crosses a grafting annulus, we would like to perform the bubbling along it and then find another bubble which avoids the real curve. The naive approach is to start from a branch point and follow the given bubble until we meet the region corresponding to the grafting annulus at the points coming from I_1; here one path can follow the curve coming from the boundary of the grafting annulus until the twin of O_1, and the other one can follow its analytic extension inside the bubble to cross the bubble from side to side. Notice that in doing this it also crosses the grafting annulus from side to side; in particular it reaches O_1. Then they keep travelling along the boundaries of the grafting annulus in the direction induced by β_1, until they meet α_2. One of them will meet that crossing before the other and will follow the analytic extension of γ inside the bubble, while the other one will follow the boundary of the annulus; the coherence parameters ε_k and $\varepsilon_{k,l}$ determine the order in which points are met, and the direction in which the paths will go. Anyway they will reach points on the same side of the annulus, but on opposite sides of the bubble, hence they can keep walking along the original bubble. This works because at every crossing there is an analytic extensions of γ inside the bubble which crosses it from side to side. However in general this naive procedure does not result in a couple of disjoint embedded arcs: already in the case of a single crossing ($N = 1$) the analytic extension of γ inside the bubble is used twice, hence we do not get a new bubble.

To fix this we consider a small collar neighbourhood $A_\gamma^\#$ of A_γ; this can be obtained by slightly pushing the boundary curves of A_γ into the hyperbolic core of the adjacent components (i.e. away from A_γ). More precisely it can be taken to be the region bounded by the couple of simple closed curves $\gamma_{\pm 1} = \{x \in S^+ \setminus A_\gamma \mid d(x, \gamma) = \varepsilon\}$, for some small $\varepsilon > 0$, which develop to the two boundaries of the region $\mathcal{N}_\varepsilon(\widehat{\gamma}) = \{\widehat{x} \in \mathbb{H}^2 \mid d(\widehat{x}, \widehat{\gamma}) \leq \varepsilon\}$. Notice that the developing image is no longer injective in the interior of $A_\gamma^\#$.

We have that $\beta \cap A_\gamma^\# = \beta_1^\# \cup \cdots \cup \beta_N^\#$ is a disjoint union of oriented bubbleable arcs properly embedded in $A_\gamma^\#$ and such that $\beta_k \subset \beta_k^\#$, which we still call crossing. Moreover each crossing $\beta_k^\#$ will intersect $\partial A_\gamma^\#$ in two points; let us label them by I_k^{-1} and O_k^{+1} in such a way that $I_k^{-1}, I_k, O_k, O_k^{+1}$ appear in this order along β; then label the curves $\gamma_{\pm 1}$ so that $I_1^{-1} \in \gamma_{-1}$ and $O_1^{+1} \in \gamma_{+1}$. Notice that for the other crossings it may happen that $I_k^{-1} \in \gamma_{\pm 1}$ and $O_k^{+1} \in \gamma_{\mp 1}$, according to the fact that $\beta_k^\#$ enters in $A_\gamma^\#$ on the same side as $\beta_1^\#$ leaves it or not; however this is not going to be a relevant in our construction.

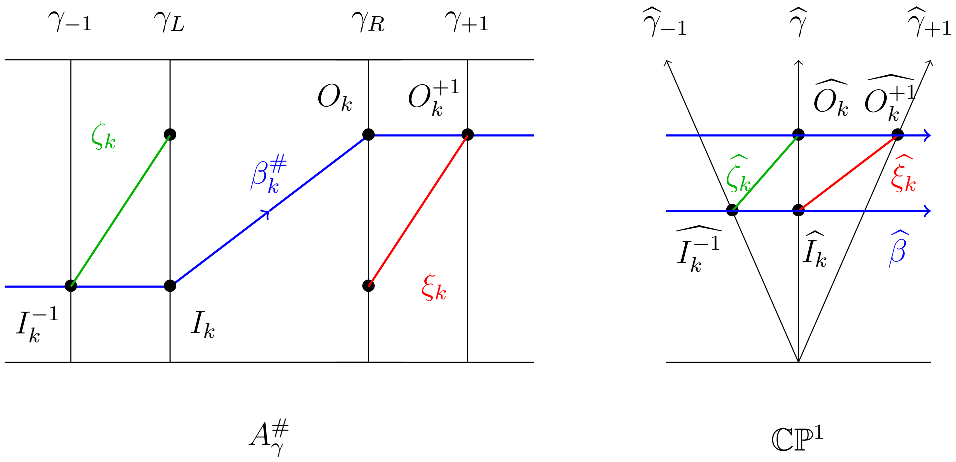

Figure 3.15: The extended annulus $A_\gamma^\#$ and the auxiliary objects.

Now for any $k = 1, \ldots, N$ we consider in the developed image in $\mathbb{H}^2$ the geodesic segment $\widehat{\zeta}_k$ from $\widehat{I_k^{-1}}$ to $\widehat{O_k}$ and the geodesic segment $\widehat{\xi}_k$ from $\widehat{I}_k$ to $\widehat{O_k^{+1}}$. This defines for us an arc ζ_k in $A_\gamma^\# \setminus A_\gamma$ starting from I_k^{-1} and ending at the twin of O_k, and an arc ξ_k in $A_\gamma^\# \setminus A_\gamma$ starting from the twin of I_k and ending at O_k^{+1}. Since β is embedded and bubbleable, all these arcs are disjoint; notice that the behaviour of ζ_k and ξ_k in $A_\gamma^\#$ essentially mimics that of β_k (e.g. they wrap around A_γ the same number of times), with the only difference that they are entirely contained in the positive region, while β_k crosses the real curve twice inside A_γ. To simplify the exposition we also find it convenient to introduce an action of $\mathbb{Z}_2 = \{\pm 1\}$ on all the auxiliary objects we have defined: we let 1 act as the identity, while -1 acts by exchanging an "entry object" with the corresponding "exit object", i.e.

$$-1.I_k = O_k \qquad -1.I_k^{-1} = O_k^{+1} \qquad -1.\zeta_k = \xi_k$$

Moreover notice that all arcs involved are oriented; for any path μ, let μ^{-1} denote the same path with the opposite orientation. We now have all the ingredients required to prove the following result. Recall that we are restricting to the easier case of simple grafting and that we will address later the general case.

Proposition 3.6.4. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be Fuchsian and let $\sigma \in \mathcal{M}_{2,\rho}$ be a bubbling of some $\sigma_0 \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta \subset \sigma_0$ such that β transversely crosses some grafting annulus A_γ of σ_0 but avoids all the other grafting annuli of σ_0. Then σ is also a bubbling of some other $\sigma_0' \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta' \subset \sigma_0'$ which avoids all the real curves of σ_0'.*

Proof. We will prove this by directly finding a new bubble with the required properties on $Bub(\sigma_0, \beta)$. Pick an orientation of β; then we have all the auxiliary objects defined above, in particular fix the orientation of γ induced by the first crossing $\alpha_1 = \beta_1$. We will define a new bubble roughly in the following way: each time the bubble coming from β enters $A_\gamma^\#$ in correspondence of some crossing α_k we will describe how to leave $A_\gamma^\#$ in correspondence of the crossing $\alpha_{k\pm1}$ by suitably following some of the auxiliary arcs (the sign depends on some coherence parameters); then we keep following β until we reach another crossing, if any, and we iterate.

Let us now define a procedure to handle the k-th crossing in the developed image (see Picture 3.16). Suppose $\widehat{\beta}$ enters $\mathcal{N}_\varepsilon(\widehat{\gamma})$ in correspondence of $\widehat{\alpha}_k = \widehat{\beta}_{\sigma(k)}$ at $\omega \widehat{I}_{\sigma(k)}^{-1}$ for some $\omega \in \{\pm1\}$. We begin by following $\omega \widehat{\zeta}_{\sigma(k)}^{\omega}$, so that we get to $\omega \widehat{O}_{\sigma(k)}$. We now distinguish two cases according to the relative position of the endpoints of the two crossings $\widehat{\alpha}_k$ and $\widehat{\alpha}_{k+\omega\varepsilon_k}$

1. if $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{I}_{\sigma(k+\omega\varepsilon_k)}$ sits after $\omega\widehat{O}_{\sigma(k)}$ along $\widehat{\gamma}^{\omega\varepsilon_k}$, then we follow $\widehat{\gamma}^{\omega\varepsilon_k}$ until we reach it; we meet $\widehat{\alpha}_{k+\omega\varepsilon_k}$ at that point $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{I}_{\sigma(k+\omega\varepsilon_k)}$ and then we can follow the arc $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{\xi}_{\sigma(k+\omega\varepsilon_k)}^{\omega\varepsilon_{k,k+\omega\varepsilon_k}}$
2. otherwise $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{I}_{\sigma(k+\omega\varepsilon_k)}$ sits before $\omega\widehat{O}_{\sigma(k)}$ along $\widehat{\gamma}^{\omega\varepsilon_k}$, then the fact that β is embedded implies that $-\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{I}_{\sigma(k+\omega\varepsilon_k)}$ is after $\omega\widehat{O}_{\sigma(k)}$; in this case we can move a little off $\widehat{\gamma}$ along $\widehat{\beta}^{\omega}$ to meet the arc $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{\xi}_{\sigma(k+\omega\varepsilon_k)}^{\omega\varepsilon_{k,k+\omega\varepsilon_k}}$

In both cases we follow the arc $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{\xi}_{\sigma(k+\omega\varepsilon_k)}^{\omega\varepsilon_{k,k+\omega\varepsilon_k}}$ and reach $\omega\varepsilon_{k,k+\omega\varepsilon_k}\widehat{O}_{\sigma(k+\omega\varepsilon_k)}^{+1}$. Then we are ready to leave $\mathcal{N}_\varepsilon(\widehat{\gamma})$ along $\beta^{\omega\varepsilon_{k,k+\omega\varepsilon_k}}$. We use this rule to define a path $\widehat{\beta}'$ in $\mathbb{CP}^1$, starting from the first endpoint of $\widehat{\beta}$.

We should explicitly remark that it is possible that β goes around some topology of the surface between two crossings β_k and β_{k+1}; in this case its developed image does not come back to the region $\mathcal{N}_\varepsilon(\widehat{\gamma})$, but to a different region $g\mathcal{N}_\varepsilon(\widehat{\gamma})$ for some Möbius transformation which depends on the topology around which β travels between β_k and β_{k+1}. However translating $\mathcal{N}_\varepsilon(\widehat{\gamma})$ with the holonomy of the structure does not produce overlaps; this follows from the fact that the developed images of the geodesic γ for the underlying uniformizing structure σ_ρ are disjoint and the fact that $\varepsilon > 0$ can be chosen to be arbitrarily small. On the other hand, if the path does not go around topology (so that $\widehat{\beta}$ keeps intersecting the same region $\mathcal{N}_\varepsilon(\widehat{\gamma})$), then it is enough to notice that the above procedure is completely reversible, in the sense that at any point the knowledge of what arc we have used at the most recent step is enough to know what arc to use to perform the next one, and viceversa. This implies that the path $\widehat{\beta}'$ which is constructed by the above rules does not pass more than once through any of its points.

Finally let us consider what happens to the parts of $\widehat{\beta}'$ which are outside the region $\mathcal{N}_\varepsilon(\widehat{\gamma})$ and its translates. By construction they come from portions of β which

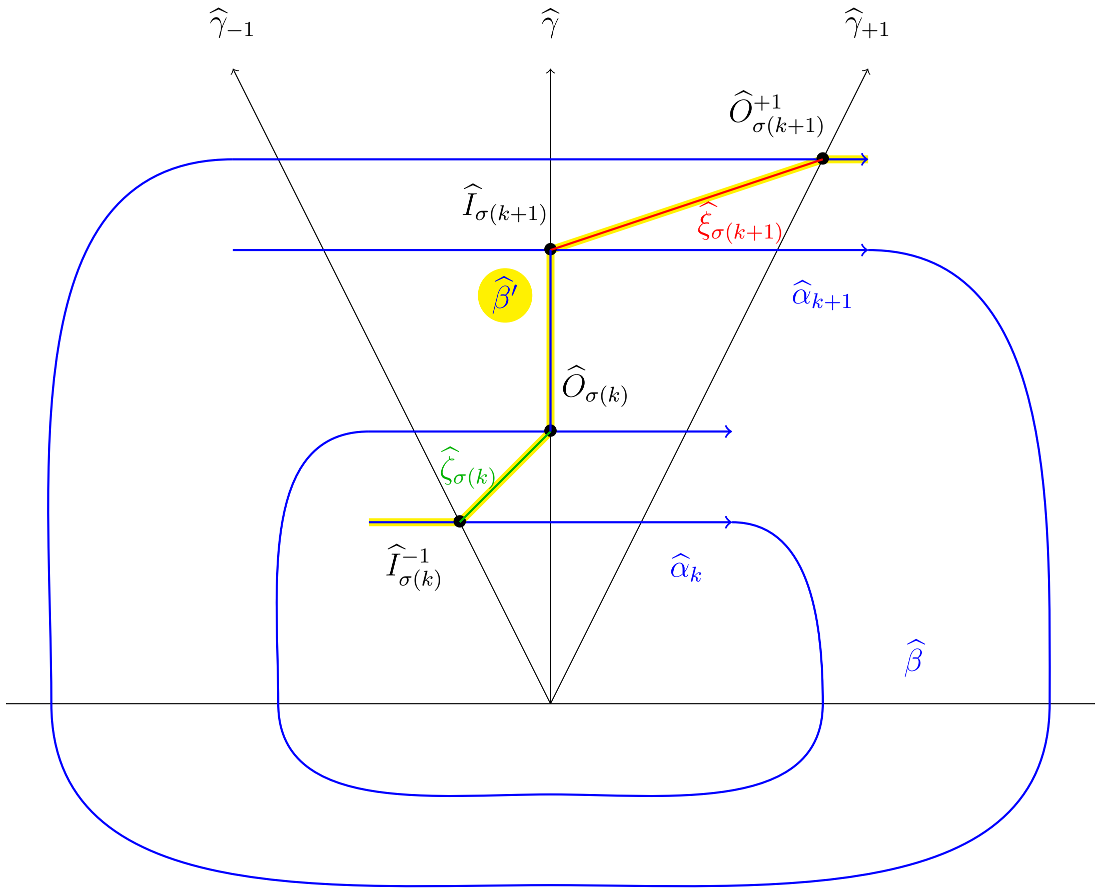

Figure 3.16: The path $\widehat{\beta}'$ in $\mathbb{CP}^1$: the k-th crossing in the case $\omega = \varepsilon_k = \varepsilon_{k,k+1} = 1$ and $\widehat{I}_{\sigma(k+1)}$ sits after $\widehat{O}_{\sigma(k)}$ along $\widehat{\gamma}$.

are outside the grafting annulus A_γ; moreover by hypothesis β does not intersect other grafting annuli. Therefore the developed images of these arcs are the same they would be in the underlying uniformizing structure σ_ρ, in particular they are all disjoint. This proves that the path $\widehat{\beta}'$ is embedded in $\mathbb{CP}^1$. Moreover since the number of marked points (ωI_k^{-1} and ωI_k) is finite, it definitively reaches the point $\widehat{O}_N^{+1}$. After that point we keep following $\widehat{\beta}$ till the end, i.e. its second endpoint. To sum up, $\widehat{\beta}'$ is an embedded path with the same endpoints as $\widehat{\beta}$ but entirely contained in $\mathbb{H}^2$.

We can now follow this path on $Bub(\sigma_0, \beta)$ to identify a new bubble (see Picture 3.17). We start at the branch point of $Bub(\sigma_0, \beta)$ which is the first with respect to the chosen orientation of β and follow the two twin paths developing to $\widehat{\beta}$ which give the boundary of the natural bubble of $Bub(\sigma_0, \beta)$. Then we check that at each crossing α_k there is a couple of embedded arcs developing to subarcs of $\widehat{\beta}'$ which fellow travel from the entry point $\omega I^{-1}_{\sigma(k)}$ to the exit point $\omega \varepsilon_{k,k+\omega\varepsilon_k} O^{+1}_{\sigma(k+\omega\varepsilon_k)}$. This follows from the fact that the auxiliary arcs ζ_k and ξ_k intersect $\beta_k^{\#}$ only at the points I_k^{-1} and O_k^{+1}, hence the copies of ζ_k and ξ_k inside the bubble coming from β cross it from side to side and at the same time they also cross the grafting annulus. As observed

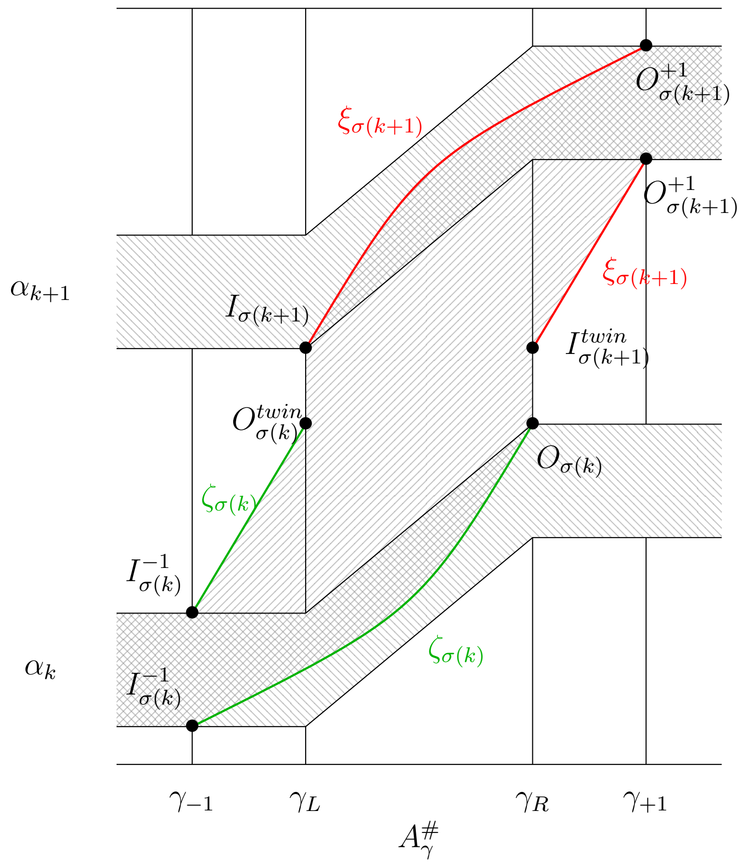

Figure 3.17: The new bubble on the surface: the k-th crossing in the case $\omega = \varepsilon_k = \varepsilon_{k,k+1} = 1$ and $\widehat{I}_{\sigma(k+1)}$ sits after $\widehat{O}_{\sigma(k)}$ along $\widehat{\gamma}$. (The two bubbles are shaded at different angles).

before, the procedure does not use the same auxiliary object twice; this guarantees that coming back to the grafting annulus does not result in new intersections, so that these paths developing to $\widehat{\beta}'$ are actually the boundary of a new bubble. Debubbling with respect to this new bubble gives the desired unbranched structure σ_0' with a bubbleable arc β' such that $Bub(\sigma_0, \beta) = Bub(\sigma_0', \beta')$. Notice that by construction β' does not intersect any real component of σ_0', because $\widehat{\beta}'$ sits entirely in $\mathbb{H}^2$. □

Depending on the intersection pattern between β and A_γ, we have different possibilities for what σ_0' looks like. We are in particular interested in the easiest case, which is the one in which β crosses A_γ just once: the structure σ_0' of the previous result is exactly the one obtained by degrafting σ_0 with respect to A_γ, as established by the following result, which provides a converse to [8, Theorem 5.1].

Corollary 3.6.5 (Degrafting Lemma). *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be Fuchsian. Let $\sigma_0 \in \mathcal{M}_{0,\rho}$, $A_\gamma \subset \sigma_0$ a grafting annulus and $\beta \subset \sigma_0$ a bubbleable arc which transversely crosses A_γ just once. Then there exists $\sigma_0' \in \mathcal{M}_{0,\rho}$ and a bubbleable arc $\beta' \subset \sigma_0'$ such that $\sigma_0 = Gr(\sigma_0', \gamma)$ and $Bub(\sigma_0, \beta) = Bub(\sigma_0', \beta')$.*

Proof. In the previous notations, we have that $\alpha_2 = \alpha_1$. Therefore the new bubble produced by the above procedure does a full turn around A_γ and encompasses the whole real curve contained in it before leaving it. Debubbling with respect to this bubble produces a structure which has no real curves in the homotopy class of γ. By Goldman classification (see [16, Theorem C]), it must be the structure obtained by degrafting σ_0 with respect to A_γ. □

We now address the more general case in which β might cross a grafting region coming from a multigrafting, hence we resume the notation $A_\gamma = A_\gamma^1 \cup \cdots \cup A_\gamma^M$ for the grafting region obtained by grafting M times the simple closed geodesic γ of σ_ρ; recall that A_γ is obtained by taking M copies of $(\mathbb{CP}^1 \setminus \widehat{\gamma})/\rho(\gamma)$ and gluing them in a chain along their geodesic boundaries, so that we see $M+1$ parallel copies of the geodesic γ.

What we want to do is to subdivide A_γ in disjoint annular regions in such a way that we are able to follow the procedure described above for the case of a simple grafting inside each of them. The natural subdivision given by the grafting annuli A_γ^h does not work: the procedure described above makes use of auxiliary curves parallel to γ obtained by slightly enlarging the grafting annulus; if we did the same here we would see a lot of overlaps. To solve this problem we consider more auxiliary curves on each side of the grafting geodesic, as many as the number M of grafting annuli which compose the grafting region $A_\gamma = A_\gamma^1 \cup \cdots \cup A_\gamma^M$. For instance we can consider the curves $\gamma_{\pm h} = \{x \in \sigma_\rho \mid d(x,\gamma) = h\varepsilon\}$ for $h = 1, \dots, M$ and an arbitrarily small $\varepsilon > 0$. They clearly develop to the boundaries of the regions $\mathcal{N}_{h\varepsilon}(\widehat{\gamma}) = \{\widehat{x} \in \mathbb{H}^2 \mid d(\widehat{x}, \widehat{\gamma}) \leq h\varepsilon\}$. Recall that the grafting annuli A_γ^h and A_γ^{h+1} meet along a copy of the grafting geodesic γ, hence around each of these copies we have well defined copies of the curves γ_j for $j = -M, \dots, M$, which we denote in the same way by a little abuse of notation; of course γ_0 is exactly γ (see Picture 3.18).

Given an oriented bubbleable properly embedded arc β which transversely cross A_γ from side to side, we can consider the crossings given by its intersections with the grafting annuli A_γ^h. Let us label the grafting annuli and the auxiliary curves γ_j so that the first annulus met by β is A_γ^1 and the first auxiliary curve is γ_{-M}. We obtain a doubly indexed family of crossings: β_k^h will be the k-th time (with respect to the orientation of β) that β crosses the annulus A_γ^h. We explicitly remark some preliminary facts. First of all the transversality assumption implies that once β enters in A_γ it has to leave on the other side, so that in each annulus A_γ^h we see the same number of crossings, which we call N. Secondly since β is bubbleable and all the grafting annuli have the same developed image, we get that the crossings $\beta_k^1, \dots, \beta_k^M$ have the same coherence and hence induce the same orientation of γ. Therefore we can consistently orient everything using β_1^1. As before this allows us to order the crossings according to the cyclic order in which they appear along this orientation; once again we obtain a doubly indexed family of crossings $\alpha_k^h = \beta_{\sigma(k)}^h$ for some permutation $\sigma \in \mathfrak{S}_N$ such that $\sigma(1) = 1$. Notice that the permutation σ is the same for all the annuli $A_\gamma^1, \dots, A_\gamma^M$ because β is embedded, and that the exit point for β_k^h coincides with the entry point of β_k^{h+1}.

Exactly as before we need to define some auxiliary points and arcs. Recall that around each parallel copy of γ we have a whole package of curves which we have labelled $\gamma_{-M}, \dots, \gamma_M$. Let us denote by $A_\gamma^\#$ the annular region containing A_γ and bounded by $\gamma_{\pm M}$, and by $A_\gamma^{h\#}$ the annular region contained in $A_\gamma^\#$, bounded by

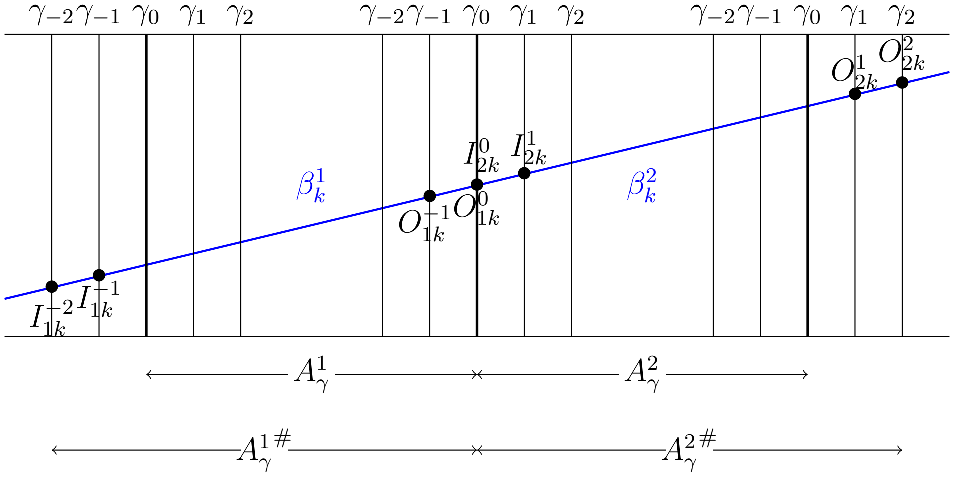

Figure 3.18: The extended region $A^\#_\gamma$ and the auxiliary curves γ_j in the case $M = 2$.

$\gamma_{-M+2h-1\pm 1}$ and containing exactly two real curves, for $h = 1, \ldots, M$; roughly speaking these regions are obtained by slightly moving A^h_γ by a certain amount of ε depending on the index h. Notice that the annuli $A^{h\#}_\gamma$ have disjoint interior and meet pairwise along some γ_j: more precisely $A^{h\#}_\gamma$ meets $A^{h+1\#}_\gamma$ along γ_{-M+2h}. Let us define the crossing $\beta^{h\#}_k = \beta \cap A^{h\#}_\gamma$ and label the intersections of $\beta^{h\#}_k$ with $\gamma_{-M+2h-2}, \gamma_{-M+2h-1}$ and γ_{-M+2h} by $I^{-M+2h-2}_{hk}, I^{-M+2h-1}_{hk}, O^{-M+2h-1}_{hk}, O^{-M+2h}_{hk}$ in such a way that they appear in this order along β. Notice that $O^{-M+2h}_{hk} = I^{-M+2(h+1)-2}_{h+1,k}$ and that a point whose apex is j belongs to an auxiliary curve labelled $\pm j$, according to the fact that crossing enters the grafting region on the same side as β_1 or not.

Finally let us define $\widehat{\zeta}_{hk}$ to be the geodesic from $\widehat{I}^{-M+2h-2}_{hk}$ to $\widehat{O}^{-M+2h-1}_{hk}$ and $\widehat{\xi}_{hk}$ to be the one from $\widehat{I}^{-M+2h-1}_{hk}$ to $\widehat{O}^{-M+2h}_{hk}$, in complete analogy to the case of a simple grating. Then we apply the same procedure described in that case modifying a crossing $\beta^{h\#}_k$ inside the annulus $A^{h\#}_\gamma$. Notice that $A^{h\#}_\gamma$ is almost as good as a genuine grafting annulus, in the sense that the open annular subregion between two copies of $\gamma_{-M+2h-1}$ is injectively developed.

Proposition 3.6.6. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be Fuchsian and let $\sigma \in \mathcal{M}_{2,\rho}$ be a bubbling of some $\sigma_0 \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta \subset \sigma_0$ such that β transversely crosses some grafting region $A_\gamma = A^1_\gamma \cup \cdots \cup A^M_\gamma$ of σ_0 but avoids all the other grafting regions of σ_0. Then σ is also a bubbling of some other $\sigma'_0 \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta' \subset \sigma'_0$ which avoids all the real curves of σ'_0.*

Proof. The strategy is the same as in the case of a simple grafting (i.e. $M = 1$, see 3.6.4), with the only difference that the procedure which resolves the crossing α^h_k must take place inside the annular region $A^{h\#}_\gamma$. These regions are precisely defined so that what happens inside one of them is completely independent from what happens inside the adjacent ones. □

Now that the ideas and the main construction have been explained in detail, let us consider the general case of an arc which crosses many grafting regions. Notice that under the assumption of the next theorem the endpoints of β are outside any grafting region.

Theorem 3.6.7. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be Fuchsian and let $\sigma \in \mathcal{M}_{2,\rho}$ be a bubbling of some $\sigma_0 \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta \subset \sigma_0$. Assume that every time β intersects some grafting region of σ_0 it actually crosses it transversely. Then σ is also a bubbling of some $\sigma_0' \in \mathcal{M}_{0,\rho}$ along a bubbleable arc $\beta' \subset \sigma_0'$ which avoids all the real curves of σ_0'.*

Proof. The strategy is to use the same technique used in 3.6.4 and 3.6.6 in any grafting annulus or region met by β. Notice that now between two crossing of a grafting region A_γ it is possible that β meets some other grafting region A_δ, for a different homotopy class δ. If we tried to resolve the intersections between β and A_γ, it would be impossible to control the behaviour of the developed images of the subarcs coming from $\beta \cap A_\delta$, i.e. to prove that the above procedure produces an injectively path in $\mathbb{CP}^1$. A way to avoid this kind of problems, is to apply the procedure of 3.6.6 *simultaneously* to all grafting regions met by β, without trying to handle different grafting regions one by one. To check that everything works as desired, it is enough to observe that any two different grafting regions A_γ and A_δ are disjoint and that also the ρ-orbits of their developed images are disjoint; this follows from the fact that this holds for any couple of simple closed geodesic on the underlying uniformizing structure. This construction realises $Bub(\sigma_0, \beta)$ as a bubbling of another structure σ_0' along an arc β' as before; by definition it avoids the real curves, exactly because we have replaced the portion crossing the grafting annuli with small geodesic arcs entirely contained in $\mathbb{H}^2$. □

We have already mentioned that [8, Theorem 5.1] states that any simple grafting can be obtained via a sequence of one bubbling and one debubbling, and we have proved an analogous statement for a simple degrafting in 3.6.5 under the assumption of (quasi-)Fuchsian holonomy. Under the same assumption, we can now obtain the same statement for any multi(de)grafting, by 3.6.7. In particular we can show that it is possible to completely degraft a structure and recover the uniformizing structure σ_ρ by just one bubbling and one debubbling.

Corollary 3.6.8 (Multi(de)grafting Lemma)**.** *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$ be Fuchsian and σ_ρ the associated uniformizing structure. Let $\sigma_0 \in \mathcal{M}_{0,\rho}$ and $\beta \subset \sigma_0$ a bubbleable arc which transversely crosses all the grafting region of σ_0 exactly once. Then there exists a bubbleable arc $\beta_\rho \subset \sigma_\rho$ such that $Bub(\sigma_0, \beta) = Bub(\sigma_\rho, \beta_\rho)$.*

Proof. Let $A_{\gamma_1}, \dots, A_{\gamma_n}$ be the grafting regions of σ_0. By 3.6.7 in $Bub(\sigma_0, \beta)$ we can find another bubble avoiding all real curves. Debubbling with respect to this bubble gives an unbranched structure without real curves, as in 3.6.5; once again by Goldman classification in [16, Theorem C] it must be the uniformizing structure. □

Notice that the roles of σ_0 and σ_ρ are symmetric in the above statement, in the sense that the same proof also proves that any multigrafting on σ_ρ can be obtained via a sequence of just one bubbling and one debubbling. By Goldman's Theorem this means that any unbranched Fuchsian structure can be obtained by the associated hyperbolic surface by a simple sequence of one bubbling and one debubbling. We have already observed in 3.5.7 that our main result, together with results from [8], [9] and [16], implies that generically we can go from a BPS σ with quasi-Fuchsian holonomy and at most 2 branch points to another structure τ with the same holonomy and at most 2 branch points by a finite sequence of bubblings and debubblings. Now 3.6.8 lets us obtain an explicit and uniform bound on the number of steps needed in such a sequence.

Corollary 3.6.9. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be quasi-Fuchsian. There is a connected, open and dense subspace of full measure $\mathcal{B} \subset \mathcal{M}_{2,\rho}$ such that if $\sigma, \tau \in \mathcal{M}_{0,\rho} \cup \mathcal{B}$ then τ is obtained from σ via a sequence of at most three bubblings and three debubblings.*

Proof. If σ and τ are unbranched, then as already observed this follows directly from 3.6.8 above (and Goldman's Theorem): a sequence made of one bubbling, one debubbling, one bubbling and one debubbling is enough. To deal with branched structures we can take $\mathcal{B}$ to be the space of structures obtained by bubbling unbranched structures provided by 3.5.6; then one more bubbling and one more debubbling are enough to reduce to the unbranched case (see the following diagram). □

$$\sigma \xrightarrow{\text{1 debub}} \sigma_0 \xrightarrow[\text{1 debub}]{\text{1 bub}} \sigma_\rho \xrightarrow[\text{1 debub}]{\text{1 bub}} \tau_0 \xrightarrow{\text{1 bub}} \tau$$

Let us remark that this result provides an explicit generically positive answer (in our setting) to the aforementioned question asked by Gallo-Kapovich-Marden in [13, Problem 12.2].

Chapter 4

The complex-analytical point of view

If in the previous chapters we studied spaces of structures with a fixed holonomy representation, in this one we want instead to look at the complex structure underlying a BPSs (see 1.1.4). We will provide constructions of BPSs coming from complex geometry (namely ODEs on a Riemann surface, or branched covers of $\mathbb{CP}^1$) and also consider the effect on the complex structure determined by a deformation inside the holonomy fibre $\mathcal{M}_{k,\rho}$. The main technical result of this chapter is a criterion which relates the geometry of the branching divisor of a BPS with the existence of local deformations of it which preserve both its holonomy and its underlying complex structure (see 4.3.9); we also discuss applications to the existence of compact submanifolds in $\mathcal{M}_{k,\rho}$ (see 4.3.12).

4.1 Structures from ODEs on Riemann surfaces

We begin by recalling a classical construction of a family of BPSs having the same holonomy representation and underlying complex structure. Let us fix a complex structure $X \in \mathcal{T}(S)$ and denote by $E = X \times \mathbb{C}^2$ the holomorphically trivial rank 2 complex vector bundle on X. It is naturally equipped with a holomorphic linear connection $\nabla_0 = \frac{\partial}{\partial z}$. More generally, it is a standard fact that any other holomorphic linear connection on E is of the form $\nabla_A = \nabla_0 - A$ for some $A \in H^0(X, K \otimes \mathfrak{gl}_2\mathbb{C})$, where K denotes the canonical divisor of X and $\mathfrak{gl}_2\mathbb{C}$ denotes the endomorphism bundle $End(E)$ of the trivial bundle E.

Definition 4.1.1. A (holomorphic linear rank 2) ODE on the Riemann surface X is the datum of a holomorphic linear connection $\nabla_A = \nabla_0 - A$ on the trivial bundle $E = X \times \mathbb{C}^2$ for some $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$.

Of course this can be also seen as a $\mathrm{SL}_2\mathbb{C}$-Higgs bundle with trivial underlying vector bundle. Let us recall the following classical fact.

Lemma 4.1.2. *A holomorphic linear connection on a holomorphic vector bundle on a Riemann surface X is flat.*

Proof. If V is a holomorphic vector bundle on X, then the curvature of a holomorphic connection on V would be a holomorphic $End(V)$-valued two forms, i.e. a holomorphic global section of $(K \wedge K) \otimes End(V)$. But there are no holomorphic 2-forms on a Riemann surface, by dimension reasons. $\square$

The flat sections of ∇_0 are the holomorphic sections of E for which $\nabla_0(s) = 0$, i.e. constant functions $c : X \to \mathbb{C}^2$; they give rise to a foliation of E by horizontal surfaces of the form $X_c = X \times \{c\}$ for $c \in \mathbb{C}^2$, which has trivial holonomy. More generally the flat sections of ∇_A give rise to a foliation of E by surfaces transverse to the fibres of E which are defined in local complex charts (U, z) by $\{(z, w) \in U \times \mathbb{C}^2 \mid w = u(z)\}$ where $u : U \to \mathbb{C}^2$ is a solution of

$$\frac{\partial u}{\partial z}(z) = A(z)u(z)$$

We can lift the ODE to the universal cover, so that local solutions can be analytically continued. A solution can therefore be seen as a holomorphic map $u : \widetilde{X} \to \mathbb{C}^2$ which is equivariant with respect to the action of $\pi_1(X)$ on $\widetilde{X}$ by deck transformations and on $\mathbb{C}^2$ by some monodromy representation $\rho_A : \pi_1(X) \to \mathrm{SL}_2\mathbb{C}$ which is determined by A. The following question

Problem: Which representations $\rho : \pi_1(S) \to \mathrm{SL}_2\mathbb{C}$ occur as monodromy representation of some ODE ∇_A on a Riemann surface diffeomorphic to S?

is a formulation of the Riemann-Hilbert problem which is still open. For instance it is not known whether real or discrete representations occur (whereas it is known that unitary representations can not, see 4.1.7 below). One approach to this question is based on the observation that if a representation is the monodromy of some ODE on a surface of genus $g \geq 2$, then the holonomy fibre $\mathcal{M}_{2g-2,\rho}$ inside the deformation space of BPSs on S contains a holomorphic sphere. We are now going to explain this classical construction.

Definition 4.1.3. A fundamental matrix for the ODE ∇_A is a holomorphic map $\Phi : \widetilde{X} \to \mathrm{SL}_2\mathbb{C}$ whose columns are a basis for the vector space of solutions (on the universal cover).

It is classical that such a map always exists and is ρ_A-equivariant, i.e.

$$\forall\, \gamma \in \pi_1(X), \Phi\gamma = \Phi\rho_A(\gamma)^{-1}$$

Moreover if Φ' is another fundamental matrix then $\Phi' = \Phi g$ for some $g \in \mathrm{SL}_2\mathbb{C}$. We are now going to use the tools we have introduced to define a family of BPSs on X. The following is an immediate consequence of the uniqueness of solutions for the Cauchy problem for the ODE.

Lemma 4.1.4. *Let $u : \widetilde{X} \to \mathbb{C}^2$ be a solution. If u vanishes at a point, then $u = 0$.*

As a result we can projectivize the whole picture. Namely let $P = \mathbb{P}(E) = X \times \mathbb{CP}^1$ be the trivial $\mathbb{CP}^1$-bundle on X. For any connection ∇_A as above we get an induced flat Ehresmann conncetion $\mathcal{F}_A$ on P transverse to the fibres of P:

the leaves of the corresponding foliations are given by surfaces locally defined by $\{(z,w) \in U \times \mathbb{CP}^1 \mid w = [u_1(z) : u_2(z)]\}$ where (U,z) is a local complex chart on X and $u : U \to \mathbb{C}^2$ is a solution of $\frac{\partial u}{\partial z}(z) = A(z)u(z)$; notice this is well defined by the previous lemma. In particular for $A = 0$ we recover the foliation of P by horizontal surfaces of the form $X_c = X \times \{c\}$ for $c \in \mathbb{CP}^1$. The monodromy of this foliation is precisely the projectivization of the linear monodromy of ∇_A, which we still denote by $\rho_A : \pi_1(X) \to \mathrm{PSL}_2\mathbb{C}$.

As soon as A is irreducible (i.e. it has no non-trivial invariant subbundles) none of the horizontal surfaces X_c is flat for ∇_A, therefore $c : X \to X_c \subset P$ defines a section of P which is generically transverse to the foliation induced by A, and is therefore endowed with a branched complex projective structure, which we denote by $\sigma_{A,c}$: namely, the parallel transport along $\mathcal{F}_A$ gives a way to map open sets of X_c to a fixed fibre of P. More precisely, as observed in [10], a developing map for $\sigma_{A,c}$ is given by

$$dev_{A,c} : \widetilde{X} \to \mathbb{CP}^1, dev_{A,c}(z) = \Phi(z)^{-1}c$$

where $\Phi : \widetilde{X} \to \mathrm{SL}_2\mathbb{C}$ is a fundamental matrix for the ODE. By construction this map is ρ_A-equivariant, as the following direct computation shows: let $\gamma \in \pi_1(X)$ then we have that

$$dev_{A,c}(\gamma z) = \Phi(\gamma z)^{-1}c = (\Phi(z)\rho_A(\gamma)^{-1})^{-1}c = \rho_A(\gamma)\Phi(z)^{-1}c = \rho_A(\gamma)dev_{A,c}(z)$$

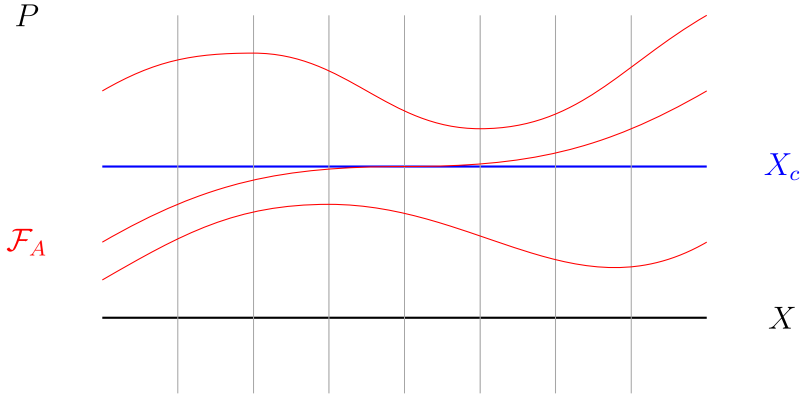

Figure 4.1: The flat projective bundle picture.

Moreover the branching divisor of $\sigma_{A,c}$ corresponds to the tangencies between $\mathcal{F}_A$ and the horizontal section $X \times \{c\} \subseteq P$; its degree can be computed by some index formulae and turns out to be $2g-2$. The following, which is proved in [10, Lemma 6.1-2], gives an complete dictionary between ODE's on Riemann surfaces and holomorphic spheres in the holonomy fibres $\mathcal{M}_{2g-2,\rho}$.

Proposition 4.1.5. *Let $X \in \mathcal{T}(S)$ and let $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$ be irreducible. Then the map $\Sigma_A : \mathbb{CP}^1 \to \mathcal{M}_{2g-2,\rho_A}, c \to \sigma_{A,c}$ is a holomorphic embedding. Conversely, let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be an irreducible representation which can be lifted to $SL_2\mathbb{C}$. Then for any compact holomorphic curve $f : Y \to \mathcal{M}_{2g-2,\rho}$ there exist $X \in \mathcal{T}(S)$, $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$ and a meromorphic map $h : Y \to \mathbb{CP}^1$ such that $f = \Sigma_A \circ h$; in particular $f(Y)$ has genus 0.*

In the following we will refer to $\Sigma_A : \mathbb{CP}^1 \to \mathcal{M}_{2g-2,\rho_A}$ as the rational curve induced by A through the ODE it defines. It is a sphere of structures which, by construction, share the same underlying complex structure X and holonomy representation ρ_A. Contrast this situation to the case of unbranched projective structures, where, by a classical theorem of Poincaré (see [18, Theorem A and 15]), the projections from $\mathcal{P}(S)$ to Teichmüller space and the $\mathrm{PSL}_2\mathbb{C}$-character variety are transverse (i.e. fibres intersect transversely). Here in the intersection between two fibres we find a 1-dimensional family of structures; a local parameter is of course given by local deformations (e.g. movements of branch points). The branching divisor of these structures induced by ODEs enjoy a particular status, as shown by the following.

Lemma 4.1.6. *Let $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$ be irreducible, $c = (c_1, c_2) \in \mathbb{C}^2 \setminus \{(0,0)\}$ and let $\sigma_{A,c} \in \mathcal{M}_{2g-2,\rho_A}$ be the induced BPS. Then the branching divisor $div(\sigma_{A,c})$ is canonical for the underlying complex structure X.*

Proof. As said above, the branch points occur exactly at tangencies between the foliation $\mathcal{F}_A$ and the horizontal section $X_c = X \times \{[c_1 : c_2]\}$ inside the projective bundle $X \times \mathbb{CP}^1$. A direct computation using the local expression for the foliations shows that this occurs exactly at points $z \in X$ at which $c = (c_1, c_2)$ is an eigenvector for $A(z)$. Since $E = X \times \mathbb{C}^2$ is the trivial bundle, we have that $H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C}) = H^0(K)^{\oplus 3}$; more explicitly, A can be written as

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & -a_{11} \end{pmatrix}, \text{ for } \ a_{ij} \in H^0(X, K)$$

Then we see that z is a point at which $c = (c_1, c_2)$ is an eigenvector for $A(z)$ if and only if the following conditions are satisfied

$$\begin{cases} c_1 a_{11}(z) + c_2 a_{12}(z) = c_1 \lambda(z) \\ c_1 a_{21}(z) - c_2 a_{11}(z) = c_2 \lambda(z) \end{cases}$$

where $\lambda \in H^0(X, K)$ is an eigenvalue of A. Since c_1, c_2 are not both zero, let us assume $c_1 \neq 0$, express $\lambda(z) = a_{11}(z) + \frac{c_2}{c_1} a_{12}(z)$ and obtain therefore that

$$c_1^2 a_{21}(z) - 2c_1 c_2 a_{11}(z) - c_2^2 a_{12}(z) = 0$$

In other words the branching divisor of $\sigma_{A,c}$ is exactly the zero divisor of the abelian differential $\Theta_{A,c} = c_1^2 a_{21} - 2c_1 c_2 a_{11} - c_2^2 a_{12}$. □

We will obtain another proof of this fact below in 4.3.11 by more analytic techniques. Let us conclude, for completeness, by recalling the following folklore observation, which tells us that unitary representations can not arise from ODEs; this answers negatively the above Riemann-Hilbert-like question for these representations.

Lemma 4.1.7. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be the monodromy representation of an ODE $\partial u = Au$ on a Riemann surface of genus $g \geq 2$. Then the image of ρ can not be contained in the unitary group* $\mathrm{PSU}(2)$.

Proof. Assume by contradiction that $\rho(\pi_1(S)) \subset \mathrm{PSU}(2)$. Then it preserves the standard unit metric of $\mathbb{CP}^1$ obtained by identifying it with the standard unit sphere $\mathbb{S}^2$, since $\mathrm{PSU}(2)=\mathrm{SO}_3\mathbb{R}=\mathrm{Isom}^+(\mathbb{S}^2)$. By the above considerations we know then that $\mathcal{M}_{2g-2,\rho}$ is not empty; structures in it will carry a natural Riemannian metric with constant sectional curvature 1 and $2g-2$ cone points of angle 4π. A direct computation shows that this is absurd; namely we have

$$\chi(S) - \sum_{i=1}^{2g-2} \frac{2\pi - 4\pi}{2\pi} = 0$$

but by Gauss-Bonnet this computation should be equal to the volume of the metric, which of course should be a strictly positive number. □

In the next sections we look for conditions on a structure $\sigma \in \mathcal{M}_{k,\rho}$ which guarantee the existence of local deformations which preserve both the holonomy and the underlying complex structure, keeping in mind this example of structures induced by a BPS. For example, as we will see, having a canonical divisor is a necessary condition for the existence of these deformations.

4.2 Interlude: Riemann-Roch computations

This section is devoted to proving the existence of holomorphic quadratic differentials with a given behaviour by means of computations with the Riemann-Roch theorem. Let X be a Riemann surface of genus $g \geq 2$. For any divisor E let us denote by $h^0(E)$ the dimension of the space of global holomorphic section of the holomorphic line bundle associated to E. We denote by K the canonical divisor of X. We will freely confuse a divisor with the associated holomorphic line bundle, and use the same letter to denote them, switching from additive to multiplicative notation depending on the point of view.

Lemma 4.2.1. *Let E be a divisor, and let $L = 2K + E$. Then*

1. *If $2-2g < deg(E) \leq 0$ then $h^0(L) = 3g-3+deg(E)$.*
2. *If $deg(E) = 2-2g$ and $-E \neq K$ then $h^0(L) = g-1$.*
3. *If $deg(E) = 2-2g$ and $-E = K$ then $h^0(L) = g$.*

Proof. We will use the classical fact that a line bundle of negative degree has no global sections, and that the only line bundle of degree 0 with non trivial global sections is the trivial one, which has just constant functions as global sections. By Riemann-Roch we get

$$h^0(L) = deg(L) + 1 - g + h^0(K-L) = deg(E) + 3 - 3g + h^0(-K-E)$$

In the first case $deg(-K-E) < 0$ hence $h^0(-K-E) = 0$. In the second case $deg(-K-E) = 0$ but $-K-E \neq 0$, hence $h^0(-K-E) = 0$ too. In the last case $deg(-K-E) = 0$ and $-K-E = 0$, hence $h^0(-K-E) = 1$. □

Notice that for $E = 0$ this recovers the dimension of the space of quadratic differentials $Q(X) = H^0(X, K^2)$; for any divisor E let us also define $Q_E(X) = \{\alpha \in Q(X) \mid (\alpha) + E \geq 0\}$. In particular if $E < 0$ then $Q_E(X)$ consists of holomorphic quadratic differentials vanishing on the points of E with prescribed multiplicity.

Lemma 4.2.2. *Let $D = \sum_{j=1}^{n} \lambda_j p_j$ be a divisor on X of degree $0 \leq deg(D) \leq 2g-2$. Then the following holds:*

1. *if $D \neq -K$, then for any $j = 1, \dots n$ and for any $r = 0, \dots, \lambda_j$ there exists a holomorphic quadratic differential with a zero of order exactly $\lambda_j - r$ at p_j and with a zero of order at least λ_l at p_l for $l \neq j$.*

2. *if $D = -K$, then for any $j = 1, \dots n$ and for any $r = 1, \dots, \lambda_j$ there exists a holomorphic quadratic differential with a zero of order exactly $\lambda_j - r$ at p_j and with a zero of order at least λ_l at p_l for $l \neq j$; moreover a holomorphic quadratic differential with a zero of order exactly $\lambda_j - 1$ at p_j and with a zero of order at least λ_l at p_l for $l \neq j$ actually has a zero of order λ_j at p_j.*

Proof. Let us fix some index j. For $r = 0, \dots, \lambda_j$ we define

$$D_r = rp_j - D = -(\lambda_j - r)p_j - \sum_{l \neq j} \lambda_l p_l$$

and get a chain of subspaces

$$Q_{-D}(X) = Q_{D_0}(X) \subseteq \dots Q_{D_{r-1}}(X) \subseteq Q_{D_r}(X) \dots \subseteq Q_{D_{\lambda_j}}(X) \subseteq Q(X)$$

An element $Q_{D_r}(X) \setminus Q_{D_{r-1}}(X)$ is a holomorphic quadratic differential with a zero of order exactly $\lambda_j - r$ at p_j and with a zero of order at least λ_l at p_l for $l \neq j$. Therefore to conclude it is enough to show that all the inclusions are strict. To do this we compute the dimension of this space, by exploiting the fact that

$$Q_{D_r}(X) = \{\alpha \in Q(X) \mid (\alpha) + D_r \geq 0\} = H^0(X, 2K + D_r)$$

and the dimension $h^0(2K + D_r)$ of the right-hand-side can be compute by the previous lemma: indeed we have $deg(D_r) = r - deg(D)$, so that $2 - 2g \leq deg(D_r) \leq 0$.

Let us begin with the case $deg(D) < 2g - 2$. In this case $2 - 2g < deg(-D) = deg(D_0) \leq D_r \leq 0$. Therefore by 4.2.1 we get

$$h^0(2K + D_r) = 3g - 3 + deg(D_r) = 3g - 3 - deg(D) + r$$

which proves that in the above chain the inclusion $Q_{D_{r-1}}(X) \subseteq Q_{D_r}(X)$ is 1-codimensional.

Let now consider the case $deg(D) = 2g-2$. In this case $deg(D_r) = 2-2g+r$, so that the previous discussion applies as soon as $r > 0$, but we need to check carefully the case $r = 0$. Indeed, if $D \neq -K$, then by 4.2.1

$$h^0(2K + D_0) = g - 1$$

$$h^0(2K + D_1) = 3g - 3 + deg(D_1) = 3g - 3 + 1 - deg(D) + r = g$$

so that also in this case all the inclusions $Q_{D_{r-1}}(X) \subseteq Q_{D_r}(X)$ are 1-codimensional. On the other hand if $D = -K$, then

$$h^0(2K + D_0) = g = h^0(2K + D_1)$$

i.e. in this case the first inclusion $Q_{D_0}(X) \subseteq Q_{D_1}(X)$ is actually the identity. □

4.3 Schiffer variations from movements of branch points

We have introduced in 1.4.23 a family of local deformations of a BPS σ defined around any branch point of σ; when $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ is non elementary these are known (by [8]) to provide local coordinates to the moduli space $\mathcal{M}_{k,\rho}$, as already mentioned in 1.4.22. Recall from 1.1.4 that any BPS has an underlying complex structure. Here we want to understand how this changes through a movement of branch points, i.e. to which extent a local deformation of the projective structure induces a deformation of the underlying complex structure. We are going to do this via an explicit computation of a Beltrami differential. This requires of course a more analytic approach to the surgery defined in 1.4.23.

Let $\sigma \in \mathcal{M}_{k,\rho}$ be a BPS on S, and let p be a branch point of order $ord(p) = \lambda = m - 1 \geq 1$. The underlying complex structure $X \in \mathcal{T}(S)$ is defined by the requirement that for any local projective chart (A, d) the map d is holomorphic as a map $d : A \to d(A) \subset \mathbb{C}$ with respect to the restriction of this complex structure to A. We begin by finding a normal local expression for the projective charts.

Lemma 4.3.1. *If p is a branch point of order $m - 1$ on a BPS σ with underlying complex structure X, then for any projective local chart (A, d) for σ at p there exists a complex local chart (A, φ) for X at p such that $d \circ \varphi^{-1} : \varphi(A) \to d(A)$ is given by $d \circ \varphi^{-1}(z) = z^m + o(z^m)$.*

Proof. Let (A, ψ) be any complex chart on A for X. Then $d \circ \psi^{-1} : \psi(A) \to d(A)$ is a holomorphic map such that $d \circ \psi^{-1}(\psi(p)) = d(p)$. Moreover $\psi(p)$ is a critical point of order $m - 1$ for it, since p is a critical point of order $m - 1$ for d and ψ is a diffeomorphism. Hence, if w is a complex coordinate on $\psi(A)$ we have an expansion of the form $d \circ \psi^{-1}(w) = a(w - \psi(p) + d(p))^m + o((w - \psi(p) + d(p))^m)$ for some $a \in \mathbb{C}^*$. Let us pick a $\lambda \in \mathbb{C}^*$ such that $\lambda^m = a$ and consider the affine transformation $\alpha : \mathbb{C} \to \mathbb{C}, \alpha(w) = \lambda(w - \psi(p) + d(p))$. Replacing ψ by $\varphi = \alpha \circ \psi$ does not change the complex structure, since α is holomorphic. Then (A, φ) is still a chart for X; if we use $z = \alpha(w)$ as a coordinate on $\varphi(A)$, then a direct computation shows that

$$d \circ \varphi^{-1}(z) = d \circ \psi^{-1} \circ \alpha^{-1}(z) = d \circ \psi^{-1}(\lambda^{-1}z + \psi(p) - d(p)) =$$

$$= a(\lambda^{-1}z)^m + o((\lambda^{-1}z)^m) = a\lambda^{-m}z^m + o(z^m) = z^m + o(z^m)$$

as desired. □

Let now fix some objects. Let σ be a BPS defined by some atlas $\mathcal{U}$, and $p \in \sigma$ be a branch point of order $m - 1$. Let (A, d) be a local projective chart at p, and let (A, φ) be a local complex chart for the underlying complex structure, chosen as in 4.3.1. Let us pick an open set B such that $p \in B \subset A$ and such that for any other projective chart $(U, g) \in \mathcal{U} \setminus \{(A, d)\}$ we have $A \cap U \subset A \setminus B$. Let z be a coordinate on $\varphi(A)$ and w on $d(A)$ and let us denote by c the holomorphic map $c = d \circ \varphi^{-1} : \varphi(A) \to d(A), c(z) = z^m + o(z^m)$.

By 1.4.24, we know that any movement of branch point (i.e. local deformation which does not change the structure of the branching divisor of σ, or, equivalently, which keeps it inside the stratum $\mathcal{M}_{\lambda,\rho}$) can always be performed by deforming a local projective chart via a straight-line isotopy. Therefore let us consider a point $q \in d(A)$; then we can find a small neighbourhood $C \subset B$ of p and a smooth bump function $\eta : d(A) \to [0,1]$ which is compactly supported in $d(B)$, $\eta = 1$ on $d(C)$ and also such that we get a well-defined isotopy $H : [0,1] \times d(A) \to d(A), H(t,w) = w + tq\eta(w)$. In particular, the map $H_t : d(A) \to d(A), H_t(w) = H(t,w)$ is a smooth isotopy of $d(A)$ which is projective on $d(A) \setminus d(B)$ and on $d(C)$, as it coincides there with the identity and with the translation $w \mapsto w + tq$ respectively.

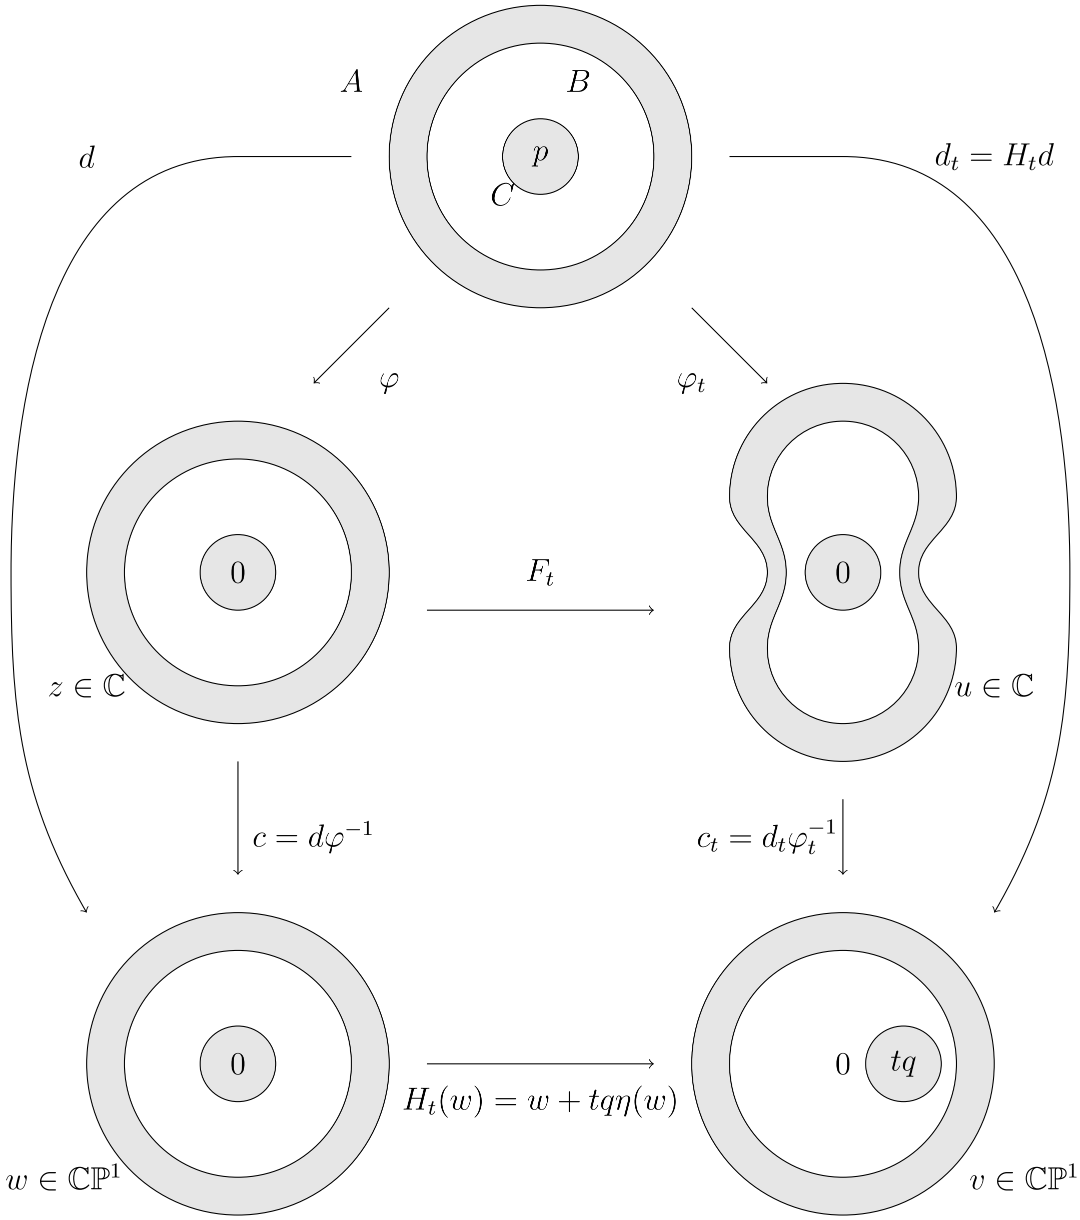

Figure 4.2: Local analysis of the movement of a branch point.

Let us define $d_t = H_t \circ d : A \to d(A)$, which gives an element in the Hurwitz space $\mathcal{H}(d)$. Then replacing the chart (A, d) with the chart (A, d_t) gives a new BPS on S, which we denote by σ_t. Let us denote by X_t the complex structure underlying σ_t. With the same argument of 4.3.1, we can obtain a local complex chart (A, φ_t) for the complex structure X_t such that we get $c_t := d_t \circ \varphi_t^{-1} : \varphi_t(A) \to d_t(A) = d(A)$ is given by $c_t(u) = tq + u^m + o(u^m)$ for a coordinate u on $\varphi_t(A)$. We can choose a diffeomorphism $F_t : \varphi(A) \to \varphi_t(A)$ which lifts $H_t : d(A) \to d_t(A) = d(A)$ (i.e. $H_t \circ c = c_t \circ F_t$); among the possible lifts, there is exactly one which satisfies $F_t \circ \varphi = \varphi_t$, and we choose this one.

Notice that F_t will be holomorphic $\varphi(A) \setminus \varphi(B)$ and on $\varphi(C)$, but not elsewhere; also notice that if $\varphi(A) = \varphi_t(A)$ then the deformation of the complex structure would be trivial, as observed in 1.4.20. Anyway this is never the case: indeed F_t can be regarded as a sort of Schiffer variation of the underlying complex structure (see [30]). To get a feeling of this we can try to get a geometric description of the domain $\varphi_t(A)$ as a subset of the complex plane with coordinate $u \in \mathbb{C}$. Without loss of generality, we can assume that $d(A) = \mathbb{D} = \{|w| < 1\}$; then $d_t(A) = \{|v| < 1\}$. Then the map c_t will take the form $v = c_t(u) = tq + \sum_{k \geq m} e_k u^k$, with $e_m = 1$, and the domain $\varphi_t(A)$ of the u-plane is the one which is mapped by c_t to the unit disk $d_t(A) = \{|v| < 1\}$ in the v-plane. A direct computation shows that

$$1 = |v|^2 = \left(tq + \sum_{k \geq m} e_k u^k\right)\left(t\overline{q} + \sum_{h \geq m} \overline{e}_h \overline{u}^h\right) =$$

$$= t^2 q\overline{q} + t\overline{q} \sum_{k \geq m} e_k u^k + tq \sum_{h \geq m} \overline{e}_h \overline{u}^h + \sum_{l \geq 2m} \sum_{h,k \geq m; h+k=l} \overline{e}_h e_k \overline{u}^h u^k =$$

$$= t|q|^2 + t \sum_{k \geq m} \left(\overline{q} e_k u^k + q\overline{e}_k \overline{u}^k\right) + e_m u^m \overline{e}_m \overline{u}^m + o(u^{2m}) =$$

$$= t^2|q|^2 + 2t \sum_{k=m}^{2m-1} Re(\overline{q} e_k u^k) + |u|^{2m} + o(u^{2m})$$

In other words $\varphi_t(A)$ is a domain bounded by a curve defined by an equation of the form

$$t^2|q|^2 + 2t \sum_{k=m}^{2m-1} Re(\overline{q} e_k u^k) + |u|^{2m} + o(u^{2m}) - 1 = 0$$

Example 4.3.2. Let us consider the easy case of a simple branch point in which $\lambda = 1, m = 2$; let us also assume that the local charts take the form $d(A) = \{|w| = 1\}$, $d_t(A) = \{|v| = 1\}$, $w = c(z) = z^2$ and $v = c_t(u) = tq + u^2$. Then the above computations tell us that $\varphi_t(A)$ is the domain of the u-plane bounded by the curve

$$t^2|q|^2 + 2tRe(\overline{q}u^2) + |u|^4 - 1 = 0$$

Plotting this curve for values such that $tq = 1 - \varepsilon$, for $\varepsilon > 0$ small enough, we can really obtain pictures consistent to the one shown on the right in 1.7. The map $F_t : \varphi(A) \to \varphi_t(A)$ from the unit disk in the z-plane to this bean-like domain is smooth and holomorphic near the boundary and around 0. However it will not be holomorphic also on the remaining annulus, since it should map it to an annular domain inside $\varphi_t(A)$ which has different modulus. We can directly compute, from the definition of the maps involved, that

$$F_t(z)^2 + tq = c_t(F_t(z)) = H_t(c(z)) = H_t(z^2) = z^2 + tq\eta(z^2)$$

i.e. the map $F_t : \varphi(A) \to \varphi_t(A)$ looks like $F_t(z) = \sqrt{z^2 + tq(\eta(z^2) - 1)}$, which coincides with $F_t(z) = z$ around $z = 0$ and with $F_t(z) = \sqrt{z^2 - tq}$ near the boundary of $\varphi(A)$.

We are now going to compute the Beltrami differential of the identity map id_S of the surface S considered as a map between the marked Riemann surfaces X and X_t. For convenience, let us introduce the notation $F(t, z) = F_t(z)$ and $\mu(t, z) = \mu_t(z)$.

Lemma 4.3.3. *In the above notations, the Beltrami differential μ_t of $id_S : X \to X_t$ is zero outside A, and with respect to the coordinate z over $\varphi(A)$ it is given by the following expression*

$$\mu(t, z) = \frac{tq\dfrac{\partial\eta}{\partial\overline{w}}(c(z))\overline{\dfrac{\partial c}{\partial z}(z)}}{\left(1 + tq\dfrac{\partial\eta}{\partial w}(c(z))\right)\dfrac{\partial c}{\partial z}(z)}$$

Proof. The identity map reads as the identity for any choice of charts, with the exception of the choice of charts (A, φ) for X and (A, φ_t) for X_t; by construction, in these charts it reads as the map $F_t : \varphi(A) \to \varphi_t(A)$, hence we reduce to compute the Beltrami differential of F_t. We recall the relation $c_t(F(t, z)) = H(t, c(z)) = c(z) + tq\eta(c(z))$. Taking the derivative with respect to z we obtain

$$\frac{\partial c_t}{\partial u}(F(t, z))\frac{\partial F}{\partial z}(t, z) = \frac{\partial c}{\partial z}(z)\left(1 + tq\frac{\partial\eta}{\partial w}(c(z))\right)$$

and taking the derivative with respect to $\overline{z}$ we obtain

$$\frac{\partial c_t}{\partial u}(F(t, z))\frac{\partial F}{\partial\overline{z}}(t, z) = \overline{\frac{\partial c}{\partial z}(z)}tq\frac{\partial\eta}{\partial\overline{w}}(c(z)))$$

Comparing the two equalities we get the desired expression for $\mu_t(z) = \dfrac{\dfrac{\partial F}{\partial\overline{z}}(t, z)}{\dfrac{\partial F}{\partial z}(t, z)}$. $\square$

Let us now compute the first-order approximation at $t = 0$ for this 1-parameter family of deformations.

Lemma 4.3.4. $\dfrac{\partial\mu}{\partial t}(0, z) = \dfrac{q\dfrac{\partial\eta}{\partial\overline{w}}(c(z))\overline{\dfrac{\partial c}{\partial z}(z)}}{\dfrac{\partial c}{\partial z}(z)} = \dfrac{\partial}{\partial\overline{z}}\left(\dfrac{\partial F}{\partial t}(0, z)\right)$

Proof. We can take a derivative in t in the expression of $\frac{\partial \mu_t}{\partial t}$ from the previous lemma to obtain

$$\frac{\partial \mu_t}{\partial t} = \frac{q\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{\frac{\partial c}{\partial z}} \frac{\partial}{\partial t}\left(\frac{t}{1+tq\frac{\partial \eta}{\partial w}(c(z))}\right) =$$

$$= \frac{q\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{\frac{\partial c}{\partial z}} \cdot \frac{1+tq\frac{\partial \eta}{\partial w}(c(z)) - tq\frac{\partial \eta}{\partial w}(c(z))}{(1+tq\frac{\partial \eta}{\partial w}(c(z)))^2} =$$

$$= \frac{q\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{\frac{\partial c}{\partial z}} \cdot \frac{1}{(1+tq\frac{\partial \eta}{\partial w}(c(z)))^2}$$

and then we evaluate at $t = 0$ to get the first identity. For the second one let us recall the notation $c_t(u) = \sum_{k \geq m} e_k(t)u^k$ and the relation $c(z) + tq\eta(c(z)) = H(t, c(z)) = c_t(F(t, z))$. Taking the derivative with respect to t we obtain

$$q\eta(c(z)) = \sum_{k \geq m} \frac{\partial}{\partial t}\left(e_k(t)F(t,z)^k\right)$$

from which we obtain

$$\frac{\partial F}{\partial t}(t,z) = \frac{q\eta(c(z)) - \sum_{k \geq m} e'_k(t)F(t,z)^k}{\frac{\partial c_t}{\partial u}(F(t,z))}$$

and, since by definition $F(0, z) = z$ and $c_0(F(0, z)) = c(z)$, at $t = 0$ we obtain

$$\frac{\partial F}{\partial t}(0,z) = \frac{q\eta(c(z)) - \sum_{k \geq m} e'_k(0)z^k}{\frac{\partial c}{\partial z}(z)}$$

In this expression the only non-holomorphic term is η, hence we obtain

$$\frac{\partial}{\partial \overline{z}}\left(\frac{\partial F}{\partial t}(0,z)\right) = \frac{\partial}{\partial \overline{z}}\left(\frac{q\eta(c(z)) - \sum_{k \geq m} e'_k(0)z^k}{\frac{\partial c}{\partial z}(z)}\right) =$$

$$= \frac{q}{\frac{\partial c}{\partial z}(z)} \frac{\partial}{\partial \overline{z}}\left(\eta(c(z))\right) = \frac{q\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{\frac{\partial c}{\partial z}(z)}$$

which proves the second identity. ∎

Remark 4.3.5. Notice that the first-order deformation at $t = 0$ has an elementary $\overline{\partial}$-primitive: namely by the above computations we get

$$\frac{\partial \mu}{\partial t}(0, z) = \frac{\partial}{\partial \bar{z}} \left(\frac{q\eta(c(z))}{\frac{\partial c}{\partial z}(z)} \right)$$

This will turn out to be useful in the following computations.

We are now going to compute an expression for the contraction of the first-order approximation $\dot{\mu}_0$ of μ_t at $t = 0$. Let us recall that p has order $\lambda = m - 1$, $c(z) = z^m + o(z^m)$, and let us introduce the notation $\frac{\partial c}{\partial z}(z) = z^{m-1}g(z)$; in particular g is a holomorphic function such that $g(0) = m$.

Proposition 4.3.6. *Let $\alpha \in Q(X) = H^0(X, K_X^2)$ be a holomorphic quadratic differential on X, and let $\alpha = \alpha(z)dz^2$ be its expression in the coordinate z on $\varphi(A)$. Then*

$$< \alpha, \dot{\mu}_0 >= \frac{2\pi i q}{(m-2)!} \frac{\partial^{m-2}}{\partial z^{m-2}} \frac{\alpha(z)}{g(z)}\bigg|_{z=0} = \frac{2\pi i q}{(\lambda - 1)!} \frac{\partial^{\lambda-1}}{\partial z^{\lambda-1}} \frac{\alpha(z)}{g(z)}\bigg|_{z=0}$$

Proof. We begin with

$$< \alpha, \dot{\mu}_0 >= \int_S \alpha \dot{\mu}_0 = \int_A \alpha \dot{\mu}_0 =$$

where we restrict the integral over A since μ_t is compactly supported inside it for any $t \in [0, 1]$. Then we can go in local coordinates in $\varphi(A)$

$$= \int_{\varphi(A)} \alpha(z)\dot{\mu}_0(z)dzd\bar{z} = \int_{\varphi(B \setminus C)} \alpha(z) \frac{q}{\frac{\partial c}{\partial z}(z)} \frac{\partial}{\partial \bar{z}}(\eta(c(z)))dzd\bar{z}$$

where the last equality comes from the above remark 4.3.5 and the fact that η is constant on $c(\varphi(C)) = d(C)$ and $c(\varphi(A \setminus B)) = d(A \setminus B)$. We now observe that

$$d\left(\frac{q\alpha(z)}{\frac{\partial c}{\partial z}(z)} \eta(c(z))dz \right) = -\frac{q\alpha(z)}{\frac{\partial c}{\partial z}(z)} \frac{\partial}{\partial \bar{z}}(\eta(c(z)))dzd\bar{z}$$

because α is holomorphic and $dzdz = 0$. Then we can continue from above and obtain

$$= -\int_{\varphi(B \setminus C)} d\left(\frac{q\alpha(z)}{\frac{\partial c}{\partial z}(z)} \eta(c(z))dz \right) =$$

to which we now apply Stokes Theorem

$$= -\int_{\varphi(\partial B)} \frac{q\alpha(z)}{\frac{\partial c}{\partial z}(z)} \eta(c(z))dz + \int_{\varphi(\partial C)} \frac{q\alpha(z)}{\frac{\partial c}{\partial z}(z)} \eta(c(z))dz = \int_{\varphi(\partial C)} \frac{q\alpha(z)}{z^{m-1}g(z)}dz$$

where the last equality comes from the fact that, by definition, η is 0 in the first integral and 1 in the second, and from the definition of $g(z) = \frac{\partial c}{\partial z}(z)z^{1-m}$. Now observe that everything inside the last integral is holomorphic, therefore we can apply the Cauchy's integral formula to obtain the desired expression

$$= \frac{2\pi i q}{(m-2)!} \frac{\partial^{m-2}}{\partial z^{m-2}} \frac{\alpha(z)}{g(z)}\bigg|_{z=0}$$

□

Example 4.3.7. In the case of a simple branch point, i.e. $m = 2$, we have $c(z) = z^2 + o(z^2)$, hence $g(0) = 2$ and the above formula reduces to

$$< \alpha, \dot{\mu}_0 >= \pi i q \alpha(0)$$

Remark 4.3.8. These expressions hold only for the first-order approximation $\dot{\mu}_0$ of our deformation, therefore they will give information about the projection to Teichmüller space only at first order. To obtain statements about the deformation itself, i.e. to reproduce the same computations with the full expression of $\mu(t, z)$, we would need to solve one of the following equations

$$\frac{tq\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{\left(1 + tq\frac{\partial \eta}{\partial w}(c(z))\right)\frac{\partial c}{\partial z}(z)} = \frac{\partial \omega}{\partial \bar{z}}(t, z) \quad \text{or} \quad \frac{t\frac{\partial \eta}{\partial \overline{w}}(c(z))\overline{\frac{\partial c}{\partial z}(z)}}{1 + tq\frac{\partial \eta}{\partial w}(c(z))} = \frac{\partial \omega}{\partial \bar{z}}(t, z)$$

for some smooth function $\omega : [0, 1] \times \varphi(A) \to \mathbb{C}$ holomorphic on $\varphi(C)$ and such that $\omega(t, z) = 0$ on $\varphi(\partial B)$. What is needed is indeed a closed expression for ω, analogous to the one obtained in 4.3.5 for $\dot{\mu}_0$.

Theorem 4.3.9. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be non elementary, $k \leq 2g - 2$, $\lambda = (\lambda_1, \ldots, \lambda_n)$ be a partition of k and $\pi^\lambda : \mathcal{M}_{\lambda,\rho} \to \mathcal{T}(S)$ denote the restriction of the projection to Teichmüller space to the the λ-stratum of the holonomy fibre. Let $\sigma \in \mathcal{M}_{k,\rho}$ and let $div(\sigma) = \sum_{j=1}^n \lambda_j p_j$ be its branching divisor. If $div(\sigma)$ is not a canonical divisor on the underlying Riemann surface $X = \pi(\sigma)$ then σ is a regular point for π^λ (i.e. π^λ is an immersion at σ).*

Proof. It is classical that a deformation of a complex structure on a surface is trivial if and only if the contraction of its Beltrami differential with any holomorphic quadratic differential is trivial. By the above discussion we know that local deformations of σ inside $\mathcal{M}_{\lambda,\rho}$ are precisely given by movements[1] of branch points. If (A_j, d_j) are local projective coordinates, we know that these deformations are parametrised by the choice of a point $q = (q_1, \ldots, q_n) \in \prod_{j=1}^n d_j(A_j)$. We are now going to show that all these deformations change the underlying complex structure, by showing that they are non trivial at first-order, i.e. we will show that there exists a holomorphic quadratic differential the contraction of which with the first-order approximation of

[1]Recall that a movement of branch points is a local deformation which preserves the structure of the branching divisor, i.e. does not involve any collapsing or splitting of branch points.

our deformations is non trivial. If (A_j, φ_j) are the local complex charts around p_j associated to (A_j, d_j) in the sense of 4.3.1, with coordinate z_j, and if $c_j = d_j \varphi_j^{-1}$ and $g_j(z_j) = \frac{\partial c_j}{\partial z_j}(z_j) z_j^{-\lambda_j}$, then by 4.3.6 we can write the contraction between any holomorphic quadratic differential $\alpha \in Q(X) = H^0(X, K^2)$ and the first-order approximation $\dot{\mu}_0$ of the Beltrami differential of this deformation as

$$< \alpha, \dot{\mu}_0 >= \sum_{j=1}^{n} \frac{2\pi i q_j}{(\lambda_j - 1)!} \frac{\partial^{\lambda_j - 1} \frac{\alpha_j}{g_j}}{\partial z_j^{\lambda_j - 1}}(z_j) \Bigg|_{z_j = 0}$$

Then we observe that

$$\frac{\partial^{\lambda_j - 1} \frac{\alpha_j}{g_j}}{\partial z_j^{\lambda_j - 1}}(z_j) = \frac{1}{g_j(z_j)} \left(\frac{\partial^{\lambda_j - 1} \alpha_j}{\partial z_j^{\lambda_j - 1}}(z_j) - \sum_{l=0}^{\lambda_j - 2} \binom{\lambda_j}{l} \frac{\partial^l \frac{\alpha_j}{g_j}}{\partial z_j^l}(z_j) \frac{\partial^{\lambda_j - 1 - l} g_j}{\partial z_j^{\lambda_j - 1 - l}}(z_j) \right)$$

By the Riemann-Roch computations in 4.2.2, under our hypothesis for any $r = 0, \ldots, \lambda_1$ there exists a holomorphic quadratic differential $\alpha^{[r]} \in Q(X)$ which has a zero of order exactly $\lambda_1 - r$ at p_1 (not higher!) and a zero of order at least λ_j at p_j for $j = 2, \ldots, n$. Let us choose $r = 2$ and compute that

$$\frac{\partial^{\lambda_j - 1} \frac{\alpha_j^{[2]}}{g_j}}{\partial z_j^{\lambda_j - 1}}(0) = \begin{cases} \dfrac{1}{g_1(0)} \dfrac{\partial^{\lambda_1 - 1} \alpha_1^{[2]}}{\partial z_1^{\lambda_1 - 1}}(0) & j = 1 \\ 0 & j = 2, \ldots, n \end{cases}$$

Since by definition $g_j(0) = \lambda_j + 1$, we conclude that

$$< \alpha, \dot{\mu}_0 >= \frac{2\pi i q_1}{(\lambda_1 - 1)!(\lambda_j + 1)} \frac{\partial^{\lambda_1 - 1} \alpha_1^{[2]}}{\partial z_1^{\lambda_1 - 1}}(0) \neq 0$$

where we have $\frac{\partial^{\lambda_1 - 1} \alpha_1^{[2]}}{\partial z_1^{\lambda_1 - 1}}(0) \neq 0$ by the choice of $\alpha^{[2]}$ made above. As a result, any deformations inside the stratum $\mathcal{M}_{\lambda,\rho}$ is non trivial at first-order. □

Corollary 4.3.10. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be non elementary and $k \leq 2g-2$. In the principal stratum $\mathcal{M}_{(1,\ldots,1),\rho}$, any movement of branch points on a non canonically branched structure changes the underlying complex structure. In particular for $k < 2g-2$ any movement of branch points in the principal stratum changes the underlying complex structure.*

Proof. In the principal stratum every local deformation preserves the structure of the divisor, since it is not possible to split simple branch points. For the second statement just observe that a divisor of degree $k < 2g - 2$ is never canonical. □

Let us define the canonical locus of $\mathcal{M}_{k,\rho}$ to be the subspace given by structures whose branching divisor is canonical for the underlying complex structure, i.e. $\mathcal{KM}_{k,\rho} = \{\sigma \in \mathcal{M}_{k,\rho} \mid div(\sigma) = K_{\pi(\sigma)}\}$, where $\pi : \mathcal{M}_{k,\rho} \to \mathcal{T}(S)$ always denotes the projection to Teichmüller space. Of course $\mathcal{KM}_{k,\rho}$ is empty for $k \neq 2g - 2$.

Corollary 4.3.11. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be non elementary, let $X \in \mathcal{T}(S)$ and let F be a positive dimensional submanifold of $\pi^{-1}(X)$. Then $F \cap \mathcal{M}_{(1,\dots,1),\rho} \subseteq \mathcal{KM}_{k,\rho}$.*

Proof. Let $\sigma \in F$, then $T_\sigma F \subseteq \ker d\pi_\sigma$. If $\sigma \in \mathcal{M}_{(1,\dots,1),\rho}$, then $T_\sigma \mathcal{M}_{(1,\dots,1),\rho} = T_\sigma \mathcal{M}_{k,\rho}$ so that actually $T_\sigma F \subseteq \ker d\pi_\sigma^{(1,\dots,1)}$. Then by 4.3.9 we get that $div(\sigma)$ must be canonical with respect to $\pi(\sigma) = X$. □

The issue with non simply branched structures is that $T_\sigma F$ might not be contained in $T_\sigma \mathcal{M}_{\lambda,\rho}$, for instance if F intersect the stratum transversely, so a priori 4.3.9 does not apply in the same straightforward way; nevertheless we can prove the following.

Corollary 4.3.12. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be non elementary and $k < 2g - 2$. Then $\mathcal{M}_{k,\rho}$ does not contain compact complex submanifolds.*

Proof. Let $Z \subset \mathcal{M}_{k,\rho}$ be a compact complex submanifold. Since Teichmüller space is a Stein manifold, the restriction of $\pi : \mathcal{M}_{k,\rho} \to \mathcal{T}(S)$ to Z must be constant, so that $Z \subset \pi^{-1}(X)$ for some $X \in \mathcal{T}(S)$. Pick some $\sigma \in Z$; it will be contained in some stratum $\mathcal{M}_{\lambda,\rho}$. If $T_\sigma Z \subset T_\sigma \mathcal{M}_{\lambda,\rho}$, then we can apply 4.3.9 as in 4.3.11 and conclude that σ is canonically branched. On the other hand if $T_\sigma Z$ is not entirely included in $T_\sigma \mathcal{M}_{\lambda,\rho}$, then it means that it is possible to perform a splitting of branch points on σ, i.e. to find some $\sigma' \in Z$ which is contained in some $\mathcal{M}_{\lambda',\rho}$, for a partition λ' obtained by splitting some entries of λ; iterating this argument, we find some $\sigma^{(m)} \in \mathcal{M}_{\lambda^{(m)},\rho}$ for which $T_{\sigma^{(m)}} Z \subset T_{\sigma^{(m)}} \mathcal{M}_{\lambda^{(m)},\rho}$ and we reduce to the previous argument to conclude that $\sigma^{(m)}$ is canonically branched. In either case we obtain that a structure of Z is canonically branched, which is not allowed by the degree of its branching divisor. □

We explicitly remark that 4.3.11 gives in particular another proof that the holomorphic spheres of BPSs induced by an ODE on X entirely consist of structures which have canonical branching divisor, which was already proved in 4.1.6 by more algebraic techniques. On the other hand comparing 4.3.12 with 4.1.5 confirms the idea that there is a tight relation between canonicity of the branching divisor and existence of compact submanifolds, showing a significantly different behaviour between the $k < 2g - 2$ and the $k = 2g - 2$ case.

4.4 Hyperelliptic structures

We now focus on BPSs which are endowed with an order two automorphism. We begin by reviewing standard material about complex structures which admit an order two biholomorphism.

Definition 4.4.1. A Riemann surface X of genus $g \geq 1$ is hyperelliptic if it satisfies any of the following equivalent definitions

- there exists a meromorphic map $\pi_X : X \to \mathbb{CP}^1$ of degree 2
- there exists a biholomorphism $J : X \to X$ such that $J^2 = id$ and $X/J = \mathbb{CP}^1$

- there exists a biholomorphism $J : X \to X$ with $2g+2$ fixed points
- there exists a polynomial $P \in \mathbb{C}[x]$ of degree D with simple roots such that X is biholomorphic to the Riemann surface obtained by compactifying the affine algebraic curve $\{(x,y) \in \mathbb{C}^2 \mid y^2 = P(x)\}$ inside the total space of the line bundle $\mathcal{O}(d)$ on $\mathbb{CP}^1$, for $d = \frac{D}{2}$ or $d = \frac{D+1}{2}$ depending on the parity of D

The map J is called the hyperelliptic involution of X.

A proof of the only non trivial equivalence of the above conditions can be found in [29, Proposition III.4.11]. In the last description, the hyperelliptic involution reads as $J(x,y) = (x,-y)$, and the projection to $\mathbb{CP}^1$ as $\pi_X(x,y) = x$. Moreover the fixed points of J are exactly the critical points of π_X, i.e. the Weierstrass points of X. This allows for an explicit description of the spaces of differentials on X. The following is proved for instance in [12, III.7.5 Corollary 1-2]

Lemma 4.4.2. *Let $\{(x,y) \in \mathbb{C}^2 \mid y^2 = P(x)\}$ define a hyperelliptic Riemann surface. Then we have that*

1. $\left\{\frac{x^j dz}{y},\ j = 0, \dots, g-1\right\}$ *is a basis of* $H^0(X,K)$
2. $\left\{\frac{x^j dz^2}{y^2}, j = 0, \dots, 2g-2\right\} \cup \left\{\frac{x^j dz^2}{y}, j = 0, \dots, g-3\right\}$ *is a basis of* $H^0(X,K^2)$

In particular we see that J acts as $-id$ on the space of abelian differentials $H^0(X,K)$, whereas the space of quadratic differentials decomposes as a direct sum of two subspaces on which J acts as id and $-id$ respectively; notice that the J-invariant part is a $(2g-1)$-dimensional subspace given exactly by the image of the product map

$$H^0(X,K) \times H^0(X,K) \to H^0(X,K^2), (\alpha,\beta) \mapsto \alpha \otimes \beta$$

Recall that every Riemann surface of genus 2 is hyperelliptic. Also notice that for $g = 2$ the anti-invariant part is the trivial subspace, so that J acts trivially on $H^0(X,K^2)$ in that case.

We are interested in the study of BPSs which admit a projective involution analogous to the hyperelliptic involution of a hyperelliptic Riemann surface. Recall from 1.2 that $Proj(\sigma)$ denotes the group of projective diffeomorphisms of a BPS σ, and from 1.2.3 that projective automorphisms are in particular biholomorphisms for the underlying complex structure. However conversely there is in general no reason why a biholomorphism of the underlying complex structure should be projective for the chosen projective structure.

Definition 4.4.3. A structure $\sigma \in \mathcal{BP}(S)$ is said to be a **hyperelliptic BPS** if the underlying complex structure X is hyperelliptic with hyperelliptic involution J and $J \in Proj(\sigma)$.

Recall from 1.2.7 that J is projective for a structure $\sigma = (dev, \rho)$ if and only if for any lift $\widetilde{J}$ to the universal cover there exists some $g \in \mathrm{PSL}_2\mathbb{C}$ such that $dev \circ \widetilde{J} = g \circ dev$. As a warm-up let us show that every unbranched structure on a torus is hyperelliptic, which follows from this characterisation and the classification of unbranched projective structures in $g = 1$ (see for instance [23]).

Lemma 4.4.4. *Every unbranched projective structure on a torus is hyperelliptic.*

Proof. Unbranched projective structures on a complex torus $X_\tau = \mathbb{C}/span_{\mathbb{Z}}(1,\tau)$ are parametrised by $\mathbb{C}$ and are obtained as affine deformations of the uniformizing structure $\sigma_\tau = \mathbb{C}/span_{\mathbb{Z}}(1,\tau)$ (seen as a projective structure). More precisely the uniformizing structure has developing map $D : \mathbb{C} \to \mathbb{P}^1, z \mapsto z$ with monodromy $\rho : \pi_1(X) = \mathbb{Z} \oplus \mathbb{Z}\tau \to \mathrm{PSL}_2\mathbb{C}, \rho(1)z = z+1$ and $\rho(\tau)z = z + \tau$; the affine structures are parametrised by $c \in \mathbb{C}^*$ and are given by $D_c : \mathbb{C} \to \mathbb{P}^1, z \mapsto e^{cz}$ with monodromy $\rho_c : \pi_1(X) = \mathbb{Z} \oplus \mathbb{Z}\tau \to \mathrm{PSL}_2\mathbb{C}, \rho_c(1)z = e^c z$ and $\rho_c(\tau)z = e^{c\tau}z$. The torus X_τ carries a canonical hyperelliptic involution J given by the map $J : \mathbb{C} \to \mathbb{C}, z \mapsto -z$ on the universal cover. When we look at J in the developed image of the uniformizing structure we see $D(J(z)) = D(-z) = -z = R(D(z))$ where $R \in \mathrm{PSL}_2\mathbb{C}$ is the rotation of π fixing $0, \infty \in \mathbb{CP}^1$. For the other affine structures we get $D(J(z)) = D(-z) = e^{-cz} = D(z)^{-1} = S(D(z))$ where $S \in \mathrm{PSL}_2\mathbb{C}$ is the inversion in the unit circle centred at 0. In both cases we see that the involution reads as a projective transformation in projective coordinates, hence in both cases the structure is hyperelliptic. □

Let us now move on to the study of branched structures. In the following examples we are going to look at the properties of the branching divisor. Recall the following classical characterisation; we use the notation $\mathcal{L}(D) = \{f \in \mathcal{M}(X) \mid (f) + D \geq 0\}$ for the space of meromorphic functions on X with poles bounded by a divisor D and $\ell(D)$ for its dimension.

Lemma 4.4.5. *Let X be a hyperelliptic Riemann surface with canonical map $\pi : X \to \mathbb{CP}^1$ and let $z_0 \in \mathbb{CP}^1$. Then $(g-1)\pi^{-1}(z_0)$ is a canonical divisor.*

Proof. Let z_0 be a branch value for π, so that $2p = \pi^{-1}(z_0)$ is a fix point of the hyperelliptic involution (i.e. a Weierstrass point). Up to an automorphism of $\mathbb{CP}^1$ the canonical map can be chosen to have a double pole at p and be holomorphic elsewhere, so that $\ell(2p) = 2$. The powers of π of course contribute with an extra dimension in the dimension of spaces $\mathcal{L}(2kp)$, so that $\ell(2mp) = m + 1$ for $m = 0, \ldots, g-1$. In particular $\ell((2g-2)p) = g$; but the only divisors with g sections and degree $2g-2$ are the canonical ones, hence $(2g-2)p = (g-1)2p = (g-1)\pi^{-1}(z_0)$ is canonical. □

Let us introduce a first class of examples of hyperelliptic BPS.

Example 4.4.6. Let X be a hyperelliptic Riemann surface of genus $g \geq 2$. The canonical map $\pi : X \to \mathbb{CP}^1$ can be taken as a developing map for a hyperelliptic BPS on X with trivial holonomy and $2g + 2$ simple branch points. Conversely, given an even number n of points $p_1, \ldots, p_n$ on $\mathbb{CP}^1$ we can take a double cover of $\mathbb{CP}^1$ branched at those points to get a Riemann surface of genus $\frac{n-2}{2}$ endowed with a hyperelliptic BPS with trivial holonomy and n simple branch points. Notice that the branching divisors of these structures are invariant under the hyperelliptic involution but not canonical. Also notice that these structures are easily deformed inside $\mathcal{M}_{2g+2,id}$, simply by moving the points p_i on $\mathbb{CP}^1$; generically these deformations do not[2] preserve the underlying complex structure, but they preserve the property that the branching divisor is invariant under the hyperelliptic involution.

The easiest example with non trivial holonomy is the following Fuchsian example.

Example 4.4.7. Let X be a hyperelliptic Riemann surface of genus $g \geq 2$. We can uniformize it to a hyperbolic surface $X = \mathbb{H}^2/\rho(\pi_1(S))$ for some Fuchsian representation $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{R}$. We usually denote by σ_ρ the resulting unbranched projective structure. Since the hyperelliptic involution is an isometry for the hyperbolic structure of σ_ρ, and hyperbolic isometries are in particular projective transformations, we easily get that σ_ρ is an unbranched hyperelliptic projective structure with Fuchsian holonomy. We will see below in 4.4.20 that indeed any unbranched projective structure on a surface of genus $g = 2$ is hyperelliptic. We can also get branched examples as follows: let p be a non Weierstrass point on X; then $J(p) \neq p$ and we can find an embedded geodesic arc β joining p to $J(p)$ such that the action of J on β looks like the action of $\mathbb{Z}_2$ on $[-1, 1] \subset \mathbb{C}$ given by $x \mapsto e^{i\pi}x$; in particular β goes through a Weierstrass point. The hyperelliptic involution clearly extends to a projective automorphism of the bubbling $\sigma = Bub(\sigma_\rho, \beta)$ which acts as a rotation of π on the bubble. Then σ is a hyperelliptic branched structure with Fuchsian holonomy and canonical branching divisor.

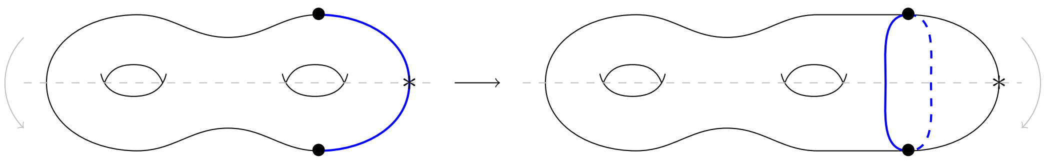

Figure 4.3: Hyperelliptic bubbling of a hyperbolic surface.

We can do a variation of the first example 4.4.6 to obtain more geometric examples, with non trivial holonomy. This is done by taking, as above, a double cover of a sphere, but now the sphere is endowed with a geometric structure which is not the standard unbranched projective structure of $\mathbb{CP}^1$, so we will need the preliminary discussion of OBPSs from Chapter 1.

Example 4.4.8. Let P_m be a right-angled geodesic polygon with m vertices in the 2-dimensional model space of constant curvature M^2_c. We have $P_3 \subset \mathbb{S}^2 = M^2_1$, $P_4 \subset \mathbb{E}^2 = M^2_0$ and for $m \geq 5$ we have $P_m \subset \mathbb{H}^2 = M^2_{-1}$. Then let $2P_m$ be the double of P_m, i.e. the (good) genus 0 orbifold obtained by gluing two copies of P_m along their boundaries. Notice that $2P_m$ carries a natural metric of curvature $1, 0$ or -1 (depending on the value of m according to the previous trichotomy) and m cone points of angle π. We can regard this as an orbibranched complex projective structure in the sense of Definition 1.2.9, with m orbifold points of order $\frac{1}{2}$.

[2]This does not follow from the previous discussion, because the trivial representation is of course elementary, but from the fact that hyperelliptic Riemann surfaces of the same genus are isomorphic if and only if their canonical maps differ by a global automorphism of $\mathbb{CP}^1$, which can not fix more than 2 points.

For $m \geq 6$ even, let us consider the branched cover $\pi : X_m \to 2P_m$ of degree 2 branching exactly over the m cone points of $2P_m$. Then a X_m is a smooth surface of genus $g = \frac{m-2}{2}$, by Riemann-Hurwitz. By 1.2.14 we know that it carries a unique BPS σ_m for which π is a projective map. A direct computation shows that σ_m is actually unbranched, and it is clearly hyperelliptic by construction; this construction actually recovers an unbranched Riemannian metric of constant curvature on X_m.

More generally we can consider a branched cover $\pi' : X'_m \to 2P_m$ of $2P_m$ branching also at some non-orbifold point; such a point will give rise to a couple of branch points of angle 4π exchange by the hyperelliptic involution in the resulting BPS σ'_m on X'_m. A developing map and the corresponding holonomy representation for this structure are obtained by pulling back via π the developing map $dev_m : \widetilde{2P_m}^{orb} \to M^2_c$ and the representation $\rho : \pi_1^{orb}(2P_m) \to Isom(M^2_c)$ defining the OBPS on $2P_m$. The construction always gives rise to hyperelliptic structures whose branching divisor is invariant by the involution, but it is not in general a canonical divisor.

The construction can be further generalised by considering more exotic OBPS on the base sphere. In any case moving the branching locus on the sphere provides non trivial deformations on the underlying complex structure, hence of the BPS.

Another main class of examples comes from ODEs on hyperelliptic Riemann surfaces. Notice that by 4.1.6 we already know these have canonical branching divisor.

Lemma 4.4.9. *Let (X, J) be a hyperelliptic Riemann surface and let $A \in H^0(X, K \otimes \mathfrak{sl}_2\mathbb{C})$ define an ODE. The structures $\sigma_{A,c}$ in the rational curve Σ_A induced by A are hyperelliptic, i.e. $J \in Proj(\sigma_{A,c})$.*

Proof. As said above, a developing map for $\sigma_{A,c}$ is given by

$$dev_{A,c} : \widetilde{X} \to \mathbb{CP}^1, dev_{A,c}(z) = \Phi(z)^{-1}c$$

where $\Phi : \widetilde{X} \to \mathrm{SL}_2\mathbb{C}$ is a fundamental matrix for the ODE. Let us lift J to a biholomorphism $\widetilde{J} : \widetilde{X} \to \widetilde{X}$ and let $du = Au$; then

$$d(u \circ \widetilde{J}) = (du) \circ \widetilde{J} \cdot d\widetilde{J} = -(du) \circ \widetilde{J} = -(A \cdot u) \circ \widetilde{J} = -(A \circ \widetilde{J})(u \circ \widetilde{J})$$

This computation shows that $\Phi \circ \widetilde{J}$ is a fundamental matrix for the ODE defined by $-A \circ J$. However on a hyperelliptic surface the hyperelliptic involution acts as $-id$ on abelian differentials, hence $-A \circ J = A$. As a result Φ and $\Phi \circ \widetilde{J}$ are fundamental matrices for the same ODE, which implies that they differ by some Möbius transformation: $\Phi \circ \widetilde{J} = \Phi g$ for some $g \in \mathrm{SL}_2\mathbb{C}$. Then we can compute that $dev_{A,c} \circ \widetilde{J} : \widetilde{X} \to \mathbb{CP}^1$ is given by

$$(dev_{A,c} \circ \widetilde{J})(z) = (\Phi \circ \widetilde{J})^{-1}(z)c = (\Phi g)^{-1}(z)c = g^{-1}\Phi(x)^{-1}c = g^{-1}dev_{A,c}(z)$$

which of course is equivalent to saying that J is projective, by 1.2.7. □

In the previous examples all branching divisors are invariant under the hyperelliptic involution essentially by construction; some of them fail to be canonical for cardinality reasons, i.e. they are not of the right degree $2g - 2$. The next examples show that it is also possible to construct hyperelliptic BPS in genus $g \geq 2$ whose branching divisor consists of $2g-2$ points, is invariant by the hyperelliptic involution, but is not canonical.

Example 4.4.10. In the setting of example 4.4.8 let us consider P_4, i.e. a regular right-angled square in the Euclidean plane. The double $2P_4$ is a flat orbifold of genus 0 with 4 cone points of angle π. We mark two more points on it and take a double branched cover branching over the 4 cones and over the 2 additional regular points. We obtain a hyperelliptic BPS σ in genus 2 whose branching divisor is of the form $p+q$ with both p and q Weierstrass points, so it is not a canonical divisor by 4.4.5. Notice that σ supports a flat metric with two singular points of angle 4π; its holonomy representation is of course not Fuchsian, not even real.

For a similar example with real holonomy consider the following.

Example 4.4.11. Let $g \in \mathrm{PSL}_2\mathbb{R}$ be a hyperbolic element. The upper and lower half planes $\mathcal{H}^{\pm} \subset \mathbb{C} \subset \mathbb{CP}^1$ are g-invariant, with quotients the Hopf annuli $A_g^{\pm}$. Each of them carries an isometric involution with two fixed points, namely the one which exchanges the two ends. The quotient is a disk with a complete hyperbolic metric with two cone points of angle π and ideal real boundary of index 0. We can glue these two disks (one positive and one negative) to obtain an OBPS on a sphere with 4 cone points of angle π; the holonomy of this structure is real and we can see, by construction, a geometric decomposition which mimics the geometric decomposition of a quasi-Fuchsian BPS. We can mark two addition point x, y on this orbifold and pick a double cover branched over them and over the 4 cone points to obtain a BPS σ in genus 2 with two points of angle 4π located at two of the Weierstrass points; as above, this implies that the branching divisor is not canonical by 4.4.5. Notice that σ carries a natural decomposition in positive, real and negative parts; the combinatorics of this decomposition of course depends on the position of x, y with respect to the decomposition of the orbifold. If both live in the positive (respectively negative) disk, then σ has a positive (respectively negative) component of genus 1 adjacent to a negative (respectively positive) annulus; if x is positive and y is negative, then σ is build from a couple of one-holed tori glued along a common real boundary. None of these combinatorics is compatible with the classification following from 2.4.5, so that σ can not have Fuchsian holonomy.

We will actually see below that this pathology does not occur in quasi-Fuchsian holonomy in genus $g = 2$ (see 4.4.24). A tool for the study of a hyperelliptic structure σ is of course the object obtained as the quotient σ/J. This does not necessarily exist in the category of branched projective structures, since it is not possible in general to projectively uniformize local charts around a fix point of J, as one would do for Riemann surfaces. Anyway it exists as an orbibranched projective structure in the sense of 1.2.9.

4.4.1 Projective automorphisms

In this section we prove a criterion to recognise whether a given biholomorphism f of a Riemann surface X is projective with respect to a given BPS σ on X, which we will use for the study of hyperelliptic structures in genus 2 in the next section.

We have already seen in 4.4.4 that unbranched structures are hyperelliptic in genus 1. We will see that the same holds in genus 2, but not for higher genus. We recall from 1.2 that if $\sigma \in \mathcal{BP}(S)$ and $f \in Diff(S)$, then $f \in Proj(\sigma)$ if it is locally given by restrictions of Möbius transformations, and that it is in particular a biholomorphism for the underlying complex structure $X = \pi(\sigma)$, where as usual we denote by $\pi : \mathcal{BP}(S) \to \mathcal{T}(S)$ the projection to Teichmüller space.

The criterion we will obtain is in term of a classical parametrisation of the fibre $\mathcal{BP}(X) = \pi^{-1}(X)$ in terms of meromorphic differentials on X via the Schwarzian derivative, which we now recall following [11] and [25].

Definition 4.4.12. Let $\Omega \subset \mathbb{C}$ be open and $f : \Omega \to \mathbb{C}$ be a holomorphic function. The Schwarzian derivative of f is defined to be $\mathcal{S}(f) = \left(\frac{f''}{f'}\right)' - \frac{1}{2}\left(\frac{f''}{f'}\right)^2$.

It is clear that if f has no critical points, then $\mathcal{S}(f)$ is holomorphic. On the other hand $\mathcal{S}(f)$ will have poles corresponding to critical points of f: more precisely a direct computation shows that where f has a critical point of order k (e.g. $f(z) = z^{k+1} + o(z^{k+1})$) the Schwarzian derivative of f has a double pole with lowest coefficient in the Laurent expansion given by $-\frac{k(k-2)}{2}$. The basic properties of this operator are the following.

Lemma 4.4.13. *Let $\Omega \subset \mathbb{C}$ be open and $f, g : \Omega \to \mathbb{C}$ be holomorphic functions such that $g(\Omega) \subset \Omega$. Then $\mathcal{S}(f \circ g) = (\mathcal{S}(f) \circ g) \cdot g'^2 + \mathcal{S}(g)$.*

Lemma 4.4.14. *Let $\Omega \subset \mathbb{C}$ be open and $f : \Omega \to \mathbb{C}$ be holomorphic function. Then $\mathcal{S}(f) = 0$ if and only if f is the restriction of a Möbius transformation.*

A direct consequence is that given an unbranched projective structure $\sigma_0 \in \mathcal{P}(X)$ over a Riemann surface $X \in \mathcal{T}(S)$, the Schwarzian derivative of a holomorphic map $f : \widetilde{X} \to \mathbb{CP}^1$ is well defined as a meromorphic quadratic differential $\mathcal{S}_{\sigma_0}(f)$ on X. This differential will have exactly a double pole with residue $-\frac{k(k-2)}{2}$ at a critical point of order k of f, and will be holomorphic elsewhere. In particular if $\sigma \in \mathcal{BP}(S)$ is defined by a developing map $dev : \widetilde{X} \to \mathbb{CP}^1$, then we can define the Schwarzian derivative $\mathcal{S}_{\sigma_0}(\sigma)$ of σ with respect to σ_0. By fixing a section $u : \mathcal{T}(S) \to \mathcal{P}(S)$ of the projection $\pi : \mathcal{BP}(S) \to \mathcal{T}(S)$, we obtain a map

$$\Phi_{u,X} : \mathcal{BP}(X) \to \Gamma(X, \mathcal{M}(K^2))$$

$$\Phi_{u,X}(\sigma) = \mathcal{S}_{u(X)}(\sigma)$$

where $\Gamma(X, \mathcal{M}(K^2))$ denotes the space of global meromorphic sections of the square of the canonical bundle of X, i.e. the space of global meromorphic differentials. For instance one can fix the Fuchsian uniformisation $u : \mathcal{T}(S) \to \mathcal{P}(S) \subset \mathcal{BP}(S), u(X) = \mathbb{H}^2/\Gamma_X$, where Γ_X is the Fuchsian group uniformizing X. This map gives a way to associate to any BPS σ a meromorphic differential on X which measures how much σ is different from the chosen uniformizing structure $u(X)$; of course $\mathcal{S}_{u(X)}(u(X)) = 0$. It is also clear that under this map, unbranched structures give rise to holomorphic quadratic differentials

$$\Phi_{u,X} : \mathcal{P}(X) \to H^0(X, K^2) = Q(X)$$

It is a classical result that, in the unbranched case, this map is actually a bijection, i.e. that any holomorphic quadratic differential is the Schwarzian derivative of some developing map for some unbranched projective structure with underlying complex structure X. This allows for a parametrisation of the space of unbranched structures by the cotangent bundle to Teichmüller space; notice however that since the identification $\mathcal{P}(X) = Q(X)$ relies on the choice of the uniformizing structure u, the intrinsic structure of $\mathcal{P}(S)$ is only that of an affine bundle.

Proving that $\Phi_{u,X} : \mathcal{P}(X) \to Q(X)$ is surjective of course involves the solution of the differential equation $\mathcal{S}_{u(X)}(w) = q$ for a given $q \in Q(X)$; this is obtained via a reduction to a linear differential equation with holomorphic coefficients. However for the general case of branched structures and meromorphic differentials the same technique fails in general and some integrability conditions show up to control the behaviour at the poles; for instance we have already observed that order of the poles is always 2 and that the value of the residue is fixed by the branching order. In [25, Theorem 3] these conditions are studied to show that the space of BPSs on X with a given branching divisor D can be identified via $\Phi_{u,X}$ with an affine algebraic subvariety of the (finite dimensional) vector space $\Gamma(X, \mathcal{M}(K^2))_D$ of meromorphic quadratic differentials on X with simple poles at the points occurring in D.

We are now going to consider the action of the biholomorphism group $Aut(X)$ of X on the space $\mathcal{BP}(X) = \pi^{-1}(X)$ of BPS on X. If $\sigma = (dev, \rho) \in \mathcal{BP}(X)$ and $F \in Aut(X)$, then $F.\sigma$ will be defined by the developing map $dev \circ \widetilde{F}^{-1}$ for some lift of F to the universal cover. By the above discussion, if we fix a section $u : \mathcal{T}(S) \to \mathcal{P}(S)$ then we get an identification $\Phi_{u,X} : \mathcal{BP}(X) \to \Gamma(X, \mathcal{M}(K^2))$ and we can look at this action on the space of meromorphic differentials. Let us denote by $F^*\omega = \omega \circ F^{-1}$ the usual action of the automorphism group on the space of differentials by pullback. Then a direct computation using the properties of the Schwarzian derivative in 4.4.13 shows the following.

Lemma 4.4.15. *Let $F \in Aut(X)$. If $\omega = \mathcal{S}_{u(X)}(\sigma)$ for some $\sigma \in \mathcal{BP}(X)$ then $F.\omega = F^*\omega + \mathcal{S}_{u(X)}(F^{-1})$.*

Proof. By definition $\mathcal{S}_{u(X)}(\sigma) = \mathcal{S}_{u(X)}(dev)$ and $\mathcal{S}_{u(X)}(F.\sigma) = \mathcal{S}_{u(X)}(dev \circ \widetilde{F}^{-1})$. □

In other words $Aut(X)$ acts affinely on $Im(\Phi_{u,X}) \subseteq \Gamma(X, \mathcal{M}(K^2))$, with linear part given by the classical action by pullback. This action can be used to obtain the following criterion.

Proposition 4.4.16. *Let $F \in Aut(X)$, $\sigma \in \mathcal{BP}(X)$ and $\omega = \mathcal{S}_{u(X)}(\sigma)$. Then $F \in Proj(\sigma)$ if and only if $F.\omega = \omega$.*

Proof. Let dev be a developing map for σ; then $\omega = \mathcal{S}_{u(X)}(dev)$. If $F \in Proj(\sigma)$ then by 1.2.7 we have $dev \circ \widetilde{F}^{-1} = g \circ dev$ for some $g \in \mathrm{PSL}_2\mathbb{C}$. Therefore $F.\omega = \mathcal{S}_{u(X)}(dev \circ \widetilde{F}^{-1}) = \mathcal{S}_{u(X)}(g \circ dev) = \mathcal{S}_{u(X)}(dev) = \omega$ by 4.4.13. On the other hand if $\mathcal{S}_{u(X)}(dev \circ \widetilde{F}^{-1}) = F.\omega = \omega = \mathcal{S}_{u(X)}(dev)$ then $dev \circ \widetilde{F}^{-1} = g \circ dev$ for some $g \in \mathrm{PSL}_2\mathbb{C}$ by 4.4.13, which implies that $F \in Proj(\sigma)$ again by 1.2.7. □

It is natural to ask if this affine action can be reduced to a linear action under a suitable choice of the uniformizing section u. This turns out to happen for the

Fuchsian uniformisation $u : \mathcal{T}(S) \to \mathcal{P}(S) \subset \mathcal{BP}(S), u(X) = \mathbb{H}^2/\Gamma_X$, where Γ_X is the Fuchsian group uniformizing X.

Lemma 4.4.17. *Let u be the Fuchsian uniformizing section. Then $Aut(X) = Proj(u(X))$; in other words a biholomorphism of X is projective for the Fuchsian uniformisation of X.*

Proof. This follows from the fact that $Aut(\mathbb{H}^2/\Gamma_X) = N(\Gamma_X)/\Gamma_X$, where $N(\Gamma_X)$ denotes the normaliser of Γ_X inside $\mathrm{Isom}^+\mathbb{H}^2 \subset \mathrm{PSL}_2\mathbb{C}$. □

Corollary 4.4.18. *Let u be the Fuchsian uniformizing section, $F \in Aut(X)$ and $\omega \in Im(\Phi_{u,X}) \subseteq \Gamma(X, \mathcal{M}(K^2))$. Then $F.\omega = F^*\omega$.*

Proof. By 4.4.15, it is enough to check that $\mathcal{S}_{u(X)}(F) = 0$. But this follows directly from 4.4.17. □

On the other hand, there are plenty of couples (σ, F) where σ is an unbranched projective structure and F is a non projective but holomorphic diffeomorphism, as shown by the following result.

Corollary 4.4.19. *Let $F \in Aut(X), F \neq id_X$. If X has genus 2, then also assume F is not the hyperelliptic involution. Let $\sigma \in \mathcal{P}(X)$ such that $F \in Proj(\sigma)$. Then there exists $\sigma' \in \mathcal{BP}(X)$ such that $F \notin Proj(\sigma)$.*

Proof. Since $\sigma \in \mathcal{BP}(X)$, we have that $\mathcal{S}_\sigma(F) = 0$. In particular if u is a uniformizing section passing through σ then the action of F on $Q(X)$ is the linear action $F.\omega = F^*\omega$ by 4.4.15. This action is known to be faithful by [12, §V.2] if and only if F is not the hyperelliptic involution of a genus 2 surface. Therefore under our hypothesis $\exists\ \omega \in Q(X)$ such that $F.\omega = F^*\omega \neq \omega$. By 4.4.16 we have that F is not projective for the projective structure $\sigma' = \Phi_{u,X}^{-1}(\omega) \in \mathcal{P}(X)$. □

For instance recall that by picking the Fuchsian uniformizing section we can satisfy the hypothesis of this statement for any $F \in Aut(X)$. With the same ideas we obtain the following.

Corollary 4.4.20. *Let X be a genus 2 Riemann surface and J its hyperelliptic involution. Then for any $\sigma \in \mathcal{P}(X)$ we have $J \in Proj(\sigma)$.*

Proof. Let us fix the Fuchsian uniformizing section to reduce to a linear action of $Aut(X)$ on $Q(X)$ and then just recall that holomorphic quadratic differentials are J-invariant on a hyperelliptic surface of genus 2 by 4.4.2. □

Notice that the same fails for hyperelliptic surfaces in higher genus. This also fails in genus 2 for branched structures, as we will see in the next section.

4.4.2 The genus $g = 2$ case

Let $g = 2$, so that every complex structure is hyperelliptic. Let $J : X \to X$ be the hyperelliptic involution of the complex structure $X = \pi(\sigma)$ underlying σ. Also recall from 4.4.2 that in genus 2 every holomorphic quadratic differential α can be written as $\alpha = \beta \otimes \gamma$ for a couple of abelian differentials $\beta, \gamma \in H^0(X, K)$. In particular $J^*\alpha = \alpha$ (but this fails for higher genus). Also recall that in genus 2 a canonical divisor is always of the form $p + J(p)$ for some point $p \in X$. We have a sort of converse to 4.3.9.

Lemma 4.4.21. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be non elementary, let $\sigma \in \mathcal{M}_{(1,1),\rho}$ be hyperelliptic and let $div(\sigma) = p + J(p)$ be a canonical divisor. Then there exist a 1-dimensional family of movements of branch points inside $\mathcal{M}_{(1,1),\rho}$ which preserve the complex structure at first-order; in other words $\dim \ker d\pi_\sigma \geq 1$, where $\pi : \mathcal{M}_{2,\rho} \to \mathcal{T}(S)$ is the projection to Teichmüller space.*

Proof. Let us choose a projective chart (A, d) at p and an adapted complex chart (A, φ) according to 4.3.1. Since $J \in Proj(\sigma)$, we have that $(J(A), d \circ J^{-1})$ is a local projective chart at $J(p)$; then a direct computation shows that $(J(A), \varphi \circ J^{-1})$ is a complex chart at $J(p)$ adapted to $d \circ J^{-1}$ in the sense of 4.3.1, and that J reads as the identity map with respect to these local coordinates, both as a biholomorphism and as a projective diffeomorphism. As a result the local deformations of σ are parametrised by the cartesian product $d(A) \times d(A)$; let us pick a generic parameter $(q, r) \in (A) \times d(A)$ and let μ_t be the Beltrami differential of the movement of branch point defined by (tq, tr), for $t \in [0, 1]$.

Now let $\alpha \in Q(X)$. Recall from 4.4.2 that in genus 2 every holomorphic quadratic differential is J-invariant, so that we have

$$\alpha \circ (\varphi \circ J^{-1})^{-1}(0) = \alpha \circ J \circ \varphi^{-1}(0) = \alpha \circ \varphi^{-1}(0)$$

which means that in the local representation in the chosen coordinates α takes the same value at p and at $J(p)$. By 4.3.6 we can compute that, with a little abuse of notation, we have

$$< \dot{\mu}_0, \alpha >= \pi i (q\alpha(p) + r\alpha(J(p))) = \pi i A(q + r)$$

where A is the common value of α at p and $J(p)$ (in the chosen coordinates). Therefore we see that the family of local deformations of σ defined by $q + r = 0$ in $d(A) \times d(A)$ is exactly the 1-dimensional family of deformations for which the first-order contraction with any $\alpha \in Q(X)$ vanishes. □

Remark 4.4.22. The same strategy does not seem to work for a double canonical divisor, i.e. when $\sigma \in \mathcal{M}_{(2),\rho}$, $J(p) = p$ and $div(\sigma) = 2p$. Indeed here we have only one local parameter q for deformations inside the stratum and the contraction gives

$$\frac{2\pi i q}{9}(\alpha'(0)g(0) - \alpha(0)g'(0))$$

which might very well be non zero for suitable choices of α. (However, since $2p$ is canonical, by 4.2.2 we get that $\alpha(p) = 0 \Rightarrow \alpha'(p) = 0$). Anyway this is fine, since the rational curve Σ_A coming from an ODE $du = Au$ is clearly not inside the minimal stratum $\mathcal{M}_{(2),\rho}$, but is somehow transverse to it, since its structures come from sections which are generically transverse to the foliation defined by the (monodromy of the) ODE.

Indeed one can consider a map $\Theta : \mathcal{M}_{2,\rho} \to \mathcal{T}(S)^{(2)}$ from $\mathcal{M}_{2,\rho}$ to the tautological bundle $\mathcal{T}(S)^{(2)}$ over Teichmüller space whose fibre over X is $Sym^2(X)$, i.e. the symmetric product of X, by sending a structure σ to its divisor $div(\sigma)$ seen as an element of the symmetric product of the underlying complex structure. In genus 2 we know that $Sym^2(X)$ is a blow up of the Jacobian of X, with exceptional divisor given by the couples of points representing a canonical class of X. The map Θ is then an isomorphism of Σ_A onto this exceptional divisor (for instance by 4.3.11), and it is classical that in genus 2 a divisor of the form $2p$ is canonical if and only if p is a Weierstrass point, and that there are exactly 6 of them. We see thus that Σ_A intersects $\mathcal{M}_{(2),\rho}$ exactly in (the structures corresponding via Θ to) these 6 points.

Remark 4.4.23. If the deformations of 4.4.21 can be integrated to deformations which actually preserve the complex structure, then by 4.3.9 they keep σ inside the canonical locus. In other words these deformations would preserve both the complex structure and the fact that the divisor is canonical. This is exactly what happens for structures induced by an ODE, by 4.1.6. On the other hand it is easily checked that in the local charts used in the proof of 4.4.21 the hyperelliptic involution reads as the identity map, essentially by construction; therefore it looks like J is suggesting to move along $q - r = 0$ to keep the divisor canonical, and not along $q + r = 0$. Of course this movement changes the underlying complex structure, since, by the above computations, the contraction gives $2\pi i\alpha(p)q$, so to get a non zero contraction it is enough to contract with some $\alpha \in Q(X)$ not vanishing at p, which exists by 4.2.2.

In 4.4.10 and 4.4.11 we have seen examples of hyperelliptic BPS whose branching divisor is invariant under the hyperelliptic involution and of degree $2g - 2$ but not canonical. The second one in particular occurs with real holonomy. However the following shows that this does not happen in quasi-Fuchsian holonomy (at least in genus 2).

Lemma 4.4.24. *Let $g = 2$, $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a non elementary representation whose image is an infinite non cyclic Gromov-hyperbolic group, and let $\sigma \in \mathcal{M}_{(1,1),\rho}$. If σ is hyperelliptic then $div(\sigma)$ is canonical.*

Proof. Let σ be hyperelliptic and let J denote the hyperelliptic involution. Then by 1.2.13 the quotient $O = \sigma/J$ is naturally endowed with an OBPS. By construction the holonomy of this structure is a representation of its orbifold fundamental group $\rho^{orb} : \pi_1^{orb}(O) \to \mathrm{PSL}_2\mathbb{C}$. Since J has order 2, we have an injection $\pi_1(S) \hookrightarrow \pi_1^{orb}(O)$ as an index 2 subgroup. Then we have that $\rho(\pi_1(S)) \hookrightarrow \rho^{orb}(\pi_1^{orb}(O))$; but a group is quasi-isometric to any finite index subgroup, so $\rho^{orb}(\pi_1^{orb}(O))$ is itself an infinite non cyclic Gromov-hyperbolic group. Since σ is hyperelliptic, its branching divisor is necessarily J-invariant. By 4.4.5 if it is not canonical, then it must be the sum of two Weierstrass points. In this case the OBPS on O has 4 points of order $\frac{1}{2}$ (i.e. angle π) coming from the other four Weierstrass points. Let us consider the double cover $T \to O$ branched over these four points. It is a torus, naturally endowed with an unbranched projective structure τ by 1.2.14. Unbranched projective structures are completely classified (see [23] for instance), and are known to be actually affine. In particular the holonomy of τ is a free abelian group of rank 2 (possibly non discrete) which appears as an index 2 subgroup of the hyperbolic group $\rho^{orb}(\pi_1^{orb}(O))$, giving the desired contradiction. □

Corollary 4.4.25. *Let $\rho : \pi_1(S) \to \mathrm{PSL}_2\mathbb{C}$ be a non elementary representation whose image is an infinite non cyclic Gromov-hyperbolic group and let $\sigma \in \mathcal{M}_{(1,1),\rho}$ be hyperelliptic. Then there exist a 1-dimensional family of movements of branch points inside $\mathcal{M}_{(1,1),\rho}$ which preserve the complex structure at first-order; in other words $\dim \ker d\pi_\sigma \geq 1$, where $\pi : \mathcal{M}_{2,\rho} \to \mathcal{T}(S)$ is the projection to Teichmüller space.*

Proof. By 4.4.24 we have that $div(\sigma)$ is canonical, hence we can apply 4.4.21. □

We conclude by observing that it is possible for a structure $\sigma \in \mathcal{BP}(X)$ over X to have a canonical branching divisor even if the hyperelliptic involution is not projective, as follows from the following description. Let us fix a uniformizing section $u : \mathcal{T}(S) \to \mathcal{P}(S)$ of X and let $\omega = \mathcal{S}_{u(X)}(\sigma)$ be the Schwarzian derivative of σ with respect to it. By 4.4.20 we have that $\mathcal{S}_{u(X)}(J.\sigma) = J^*\omega + \mathcal{S}_{u(X)}(J) = J^*\omega$. Therefore σ is hyperelliptic if and only if $J^*\omega = \omega$, by 4.4.16. The J-invariance of quadratic differentials in genus 2 is known by 4.4.2 (and as already observed means that any unbranched projective structure in genus 2 is actually hyperelliptic), but here ω is meromorphic (with double poles on a canonical divisor). By [25, Theorem 3] we know that if ω_0 is any meromorphic differential arising as the Schwarzian derivative of some BPS with two simple branch points, then $\omega - \omega_0$ is given by a meromorphic quadratic differential with at most simple poles on the same points. Through this parametrisation it is actually possible to write down explicit equations for the moduli space of BPS with two simple branch points on a given canonical divisor on X and to explicitly recognise the ones for which the hyperelliptic involution is projective.

Proposition 4.4.26. *Let X be a Riemann surface of genus $g = 2$ with hyperelliptic involution J. Let $D = p + J(p)$ be a canonical divisor on X and $\mathcal{BP}(X)_D = \{\sigma \in \mathcal{BP}(X) \mid div(\sigma) = D\}$. Then $\mathcal{BP}(X)_D$ carries the structure of an affine algebraic variety with two irreducible components; moreover the hyperelliptic structures fill an entire irreducible component.*

Proof. Let us fix a realisation of X as the compactification of the affine curve $C = \{(x, y) \in \mathbb{C}^2 \mid y^2 = h(x)\}$ for some monic polynomial $h(x) = \sum_{k=0}^{6} a_k x^k \in \mathbb{C}[x]$ inside the total space of the line bundle $\mathcal{O}(3) \to \mathbb{CP}^1$. This is the line bundle on $\mathbb{CP}^1$ with transition functions $(x, y) \mapsto (z = x^{-1}, w = x^{-3}y)$. In particular X is obtained from C just by adding two points at infinity $\infty_\pm$ whose coordinates are $(0, \pm 1)$ in the (z, w)-coordinates; in these coordinates the equation of X becomes $w^2 = \sum_{k=0}^{6} a_k z^{6-k}$. Moreover the hyperelliptic involution reads as $J(x, y) = (x, -y)$ and we have a couple of meromorphic functions $\lambda : X \to \mathbb{CP}^1, \lambda(x, y) = x$ and $\mu : X \to \mathbb{CP}^1, \mu(x, y) = y$.

Without loss of generality we can assume that the divisor is at infinity, i.e. $p + J(p) = \infty_\pm$. Following [25] we fix an integrable meromorphic quadratic differential with double poles at D and residue $-\frac{3}{2}$, for instance $-\frac{3}{2}x^4y^{-2}dx^2$ and a basis for the space of meromorphic quadratic differentials with at most simple poles on D and residue 1 at the poles (if any), for instance

$$\left\{\frac{dx^2}{y^2}, \frac{xdx^2}{y^2}, \frac{x^2dx^2}{y^2}, \frac{(y+x^3)dx^2}{y^2}, \frac{(y-x^3)dx^2}{y^2}\right\}$$

Notice that the first three differentials are a basis for the holomorphic quadratic differentials on X (see 4.4.2), whereas $(y \pm x^3)y^{-2}dx^2$ has a simple pole at $\infty_\pm$. Then any meromorphic quadratic differential on X with at most double poles at D and residue $-\frac{3}{2}$ is of the form

$$\omega_{t,s}(x,y) = -\frac{3}{2}\frac{x^4dx^2}{y^2} + t_1\frac{(y+x^3)dx^2}{y^2} + t_2\frac{(y-x^3)dx^2}{y^2} + \sum_{j=0}^{2} s_j\frac{x^jdx^2}{y^2}$$

for some $(t_1, t_2, s_0, s_1, s_2) \in \mathbb{C}^5$. Since x is a J-invariant function but y is anti-J-invariant, the differential $\omega_{t,s}$ is J-invariant if and only if $t_1 + t_2 = 0$. On the other hand it is integrable (i.e. corresponds to a BPS) if and only if its coefficients $(t_1, t_2, s_0, s_1, s_2)$ satisfy certain equations which can be obtained by looking at the Laurent expansion of $\omega_{t,s}$ around the points at infinity $\infty_\pm$ in the (z, w) coordinates, which is the following

$$\omega_{t,s}(z,w) = \frac{dz^2}{z^4}\frac{z^6}{w^2}\left(s_0 + \frac{s_1}{z} + \frac{s_2}{z^2} + \frac{t_1-t_2}{2z^3} - \frac{3}{2z^4} + \frac{(t_1+t_2)w}{2z^3}\right) =$$

$$= \frac{dz^2}{\sum_{k=0}^{6} a_k z^{6-k}}\left(s_0z^2 + s_1z + s_2 + \frac{t_1-t_2}{2z} - \frac{3}{2z^2} + \frac{(t_1+t_2)w}{2z}\right) =$$

$$= \left(-\frac{3}{2z^2} + \frac{3a_5 + t_1 - t_2 \pm (t_1+t_2)}{2z} + s_2 + A + o(1)\right)dz^2$$

$$\text{for } A = -\frac{3(a_4+a_5^2)}{2} - a_5\frac{t_1 - t_2 \pm (t_1+t_2)}{2} \pm \frac{t_1+t_2}{4}a_5 \in \mathbb{C}$$

From [25] we get that $\omega_{t,s}$ is integrable if and only if

$$\begin{cases} t_1^2 + \frac{3}{2}a_5t_1 + \frac{1}{2}a_5t_2 + 2s_2 - 3(a_4+a_5^2) + \frac{9}{4}a_5 = 0 \\ t_2^2 - \frac{3}{2}a_5t_2 - \frac{1}{2}a_5t_1 + 2s_2 - 3(a_4+a_5^2) + \frac{9}{4}a_5 = 0 \end{cases}$$

These equations define $\mathcal{BP}(X)_D$ as an affine algebraic subvariety of $\mathbb{C}^5$, of the form $\mathcal{V} \times \mathbb{C}^2$, for some curve $\mathcal{V} \subset \mathbb{C}^3$ defined by the same set of equations (which actually involve only t_1, t_2 and s_2). These equations can be summed and subtracted to get an equivalent set of defining polynomials for $\mathcal{V}$, which is

$$\begin{cases} t_1^2 + t_2^2 + t_1 - t_2 + 4s_2 + 2B = 0 \\ t_1^2 + t_2^2 + 2t_1 + 2t_2 = 0 \end{cases}$$

for $B = \frac{9}{4}a_5 - 3(a_4 + a_5^2) \in \mathbb{C}$. Up to a change of coordinates $X = t_1 + t_2, Y = t_1 - t_2 + 2, Z = 4s_2 + 2B$ these equations reduce to

$$\begin{cases} \frac{X^2}{4} + \frac{Y^2}{4} + Y + Z = 0 \\ XY = 0 \end{cases}$$

and the J-invariance condition becomes $X = 0$. Up to another change of coordinates $\eta = 2X, \xi = 2Y - 2, \zeta = Z$ the equations finally reduce to

$$\begin{cases} \eta^2 + \xi^2 + \zeta = 1 \\ \eta(\xi - 1) = 0 \end{cases}$$

and the J-invariance condition becomes $\eta = 0$. So we see that $\mathcal{V}$ has two irreducible components, each of which is a parabola, and only one of them is contained inside the plane $\eta = 0$. Only the BPSs corresponding to differentials from this component will have a projective hyperelliptic involution. □

Unfortunately from this point of view it is not clear how to use additional information about the holonomy representation to determine the component a given structure belongs to.

Chapter 5

Appendix: subgroups of $\mathrm{PSL}_2\mathbb{C}$

In this appendix we collect standard background material about the group of Möbius transformations $\mathrm{PSL}_2\mathbb{C}$ and its subgroups, to fix terminology and notations. We recall the fundamental dichotomy for elements of $\mathrm{PSL}_2\mathbb{C}$.

Definition 5.0.27. An element $id \neq g \in \mathrm{PSL}_2\mathbb{C}$ is called

- **parabolic** if it has one fix point in $\mathbb{CP}^1$, i.e. if it lifts to a non diagonalizable matrix in $\mathrm{SL}_2\mathbb{C}$; equivalently $tr^2(g) = 4$

- **loxodromic** if it has two fix points in $\mathbb{CP}^1$, i.e. if it lifts to a diagonalizable matrix in $\mathrm{SL}_2\mathbb{C}$; equivalently $tr^2(g) \neq 4$. We say that it is **real** if $tr^2(g) \in [0, +\infty[$ and in particular that it is **elliptic** if $tr^2(g) \in [0, 4[$, and that it is **hyperbolic** if $tr^2(g) \in]4, +\infty[$.

Notice that tr^2 is well defined, since $\mathrm{PSL}_2\mathbb{C} = \mathrm{SL}_2\mathbb{C}/\pm id$.

We collect now standard facts about subgroups, to fix terminology and notation, following [3].

5.1 Non elementary subgroups

Let $\Gamma < \mathrm{PSL}_2\mathbb{C}$ be a subgroup.

Definition 5.1.1. We define $\Lambda_\Gamma^0 = \{x \in \mathbb{CP}^1 \mid gx = x$ for some non elliptic loxodromic $g \in \Gamma\}$. Its closure is denoted by Λ_Γ and called the limit set of Γ. The complement $\Omega_\Gamma = \mathbb{CP}^1 \setminus \Lambda_\Gamma$ is called the discontinuity domain.

The set Λ_Γ encodes a great amount of information about Γ; it may be empty, and it may coincide with the whole $\mathbb{CP}^1$.

Definition 5.1.2. Γ is **non elementary** if Λ_Γ is not finite. Γ is **Kleinian** if Γ is discrete.

We collect results from [3, §5.3] in the following statements.

Theorem 5.1.3. *If Γ is non elementary, Λ_Γ is the smallest closed non empty Γ-invariant subset of $\mathbb{CP}^1$; moreover it is perfect and uncountable*

Theorem 5.1.4. *If Γ is Kleinian and non elementary, then Λ_Γ is the set of accumulation points of any orbit and Ω_Γ is a maximal domain of $\mathbb{CP}^1$ on which Γ acts properly discontinuously.*

Elementary groups are quite easy to understand, since they actually lie in smaller subgroups of $\mathrm{PSL}_2\mathbb{C}$ (see [3, §5.1] for a detailed discussion). The possibilities are actually the following ones

1. Λ_Γ is empty and up to conjugation $\Gamma < \mathrm{PSU}(2)=\mathrm{SO}(3)$
2. Λ_Γ consists of a single point and up to conjugation $\Gamma < \mathrm{Aff}(\mathbb{C})=\mathbb{C} \rtimes \mathbb{C}^*$; in this case we say that Γ is reducible
3. Λ_Γ consists of two points and up to conjugation $\Gamma < \{z \mapsto az^s \mid a \in \mathbb{C}^*, s^2 = 1\}$; in this case we say that Γ is completely reducible

The terminology is consistent with the one for subgroups of $\mathrm{SL}_2\mathbb{C}$: the action of $\mathrm{PSL}_2\mathbb{C}$ con $\mathbb{CP}^1$ by Möbius transformations is exactly the projectivization of the linear action of $\mathrm{SL}_2\mathbb{C}$ on $\mathbb{C}^2$, and fix points correspond to eigenspaces. In particular abelian subgroups are elementary and non elementary subgroups are irreducible.

5.2 Quasi-Fuchsian subgroups

Traditionally a discrete subgroup Γ of $\mathrm{PSL}_2\mathbb{R}$ is called a Fuchsian group of the first kind if its limit set is the entire $\mathbb{RP}^1$, and of the second kind otherwise (i.e. if it is a proper subset of $\mathbb{RP}^1$). Examples of Fuchsian groups of the first kind are given by groups which uniformize closed Riemann surfaces of genus $g \geq 2$ as quotients of the upper-half plane of $\mathbb{C}$: they admit fundamental domains of finite area. Examples of the second kind are given by elementary (e.g. cyclic) or Schottky (e.g. loxodromic free) subgroups of $\mathrm{PSL}_2\mathbb{R}$. In this work we are not interested in Fuchsian groups of the second kind, thus we give the following definition.

Definition 5.2.1. A **Fuchsian group** is a discrete subgroup Γ of $\mathrm{PSL}_2\mathbb{R}$ whose limit set is the whole $\mathbb{RP}^1$.

Notice that by definition a Fuchsian group preserves the following decomposition of the Riemann sphere

$$\mathbb{CP}^1 = \mathcal{H}^+ \cup \mathbb{RP}^1 \cup \mathcal{H}^-$$

where $\mathcal{H}^+ = \{z \in \mathbb{C} \mid Im(z) > 0\}$ and $\mathcal{H}^- = \{z \in \mathbb{C} \mid Im(z) < 0\}$ and acts by isometries on $\mathcal{H}^\pm$ with respect to the hyperbolic metric $ds^2 = \frac{dx^2+dy^2}{y^2}$ and by real projective transformations on $\mathbb{RP}^1$. This is the geometric property we are interested in, and which is actually shared by a much wider class of subgroups of $\mathrm{PSL}_2\mathbb{C}$.

Definition 5.2.2. A **quasi-Fuchsian** group is a discrete subgroup Γ of $\mathrm{PSL}_2\mathbb{C}$ whose limit set is a Jordan curve.

Once more we could define quasi-Fuchsian groups of the first and second kind (see [5, Section 6]), but we are not interested in groups of the second kind, hence we stick to this definition. Easy examples of quasi-Fuchsian are obtained by conjugating a Fuchsian group by quasi-conformal transformations of $\mathbb{CP}^1$. According to the following result this is the only way to obtain the finitely generated ones.

Theorem 5.2.3. *(Bers, [5, Theorem 4]) A finitely generated quasi-Fuchsian group is a quasi-conformal deformation of a Fuchsian group, i.e. there exist a Fuchsian group $\Gamma_0 \subset \mathrm{PSL}_2\mathbb{R}$ and a quasi-conformal homeomorphism $f : \mathbb{CP}^1 \to \mathbb{CP}^1$ such that $f\Gamma f^{-1} = \Gamma_0$. In particular the limit set is a quasi-circle*[1].

It is immediate to observe that a quasi-Fuchsian group has a domain of discontinuity made of two invariant disks $\Omega_\Gamma^\pm$. This fact characterizes quasi-Fuchsian groups among finitely generated discrete subgroups of $\mathrm{PSL}_2\mathbb{C}$, according to the following results.

Theorem 5.2.4. *(Accola, [1, Lemma 6]) Let Γ be a Kleinian group. If (Ω_1, Ω_2) is a couple of non-empty disjoint Γ-invariant domains which is maximal (with respect to inclusion), then Ω_1 and Ω_2 are simply connected.*

Theorem 5.2.5. *(Maskit, [28, Theorem 2]) A finitely generated discrete subgroup of $\mathrm{PSL}_2\mathbb{C}$ whose domain of discontinuity consists of two invariant components is a quasi-conformal deformation of a Fuchsian group.*

We will not be interested in non finitely generated groups, thus, thanks to the above results, we consider the following as equivalent definitions of a finitely generated quasi-Fuchsian group:

- Γ is topologically conjugated to a Fuchsian group
- Γ is quasi-conformally conjugated to a Fuchsian group
- the limit set Λ_Γ is a Jordan curve
- the limit set Λ_Γ is a quasi-circle
- the domain of discontinuity Ω_Γ consists of two disks

Remark 5.2.6. Given a Fuchsian group Γ, the orientation of $\mathbb{CP}^1$ induces an orientation on $\mathbb{RP}^1 = \Lambda_\Gamma$ such that $\mathcal{H}^+$ lies on the left of $\mathbb{RP}^1$ and $\mathcal{H}^-$ lies on the right. Since quasi-conformal maps are orientation-preserving, we see that the limit set of any quasi-Fuchsian group carries a natural orientation. We can therefore distinguish the two components of the domain of discontinuity as a positive and a negative region; more precisely if Γ is quasi-Fuchsian and arises as a quasi-conformal deformation of some Fuchsian group Γ_0 (i.e. $f\Gamma f^{-1} = \Gamma_0$), then $\Omega_\Gamma^\pm = f^{-1}(\mathcal{H}^\pm)$.

[1]Since not all homeomorphisms of the sphere are quasi-conformal, not all Jordan curves are quasi-circles. A cardioid is an example of a Jordan curve which is not a quasi-circle. However Bers' theorem implies that if a Jordan curve is the limit set of a finitely generated discrete subgroup of $\mathrm{PSL}_2\mathbb{C}$ then it is actually a quasi-circle.

Acknowledgements

I would like to thank my advisor Stefano Francaviglia, for introducing me to the exciting problems which are discussed in this thesis, for sharing with me his 2π-passion for mathematics and for many invaluable conversations about geometry. I am also grateful to the referees, Bertrand Deroin, Luca Migliorini and Gabriele Mondello, for their interest in this work and for providing many useful comments about this manuscript.

Furthermore my gratitude goes to the Department of Mathematics of the University of Bologna, and to all the people I have met there. I especially would like to thank Rita Fioresi, Luca Migliorini and Alberto Parmeggiani for providing lots of advices and stimulating mathematical discussion. A special thanks also goes to my Phd fellows: Gianluca Faraco, Camilla Felisetti, Alessio Savini and Marco Trozzo for the endless afternoons of geometry and coffee we have spent together; Erika Battaglia, Serena Federico and Benedetta Franceschiello for being irreplaceable QdV-musketeers; and of course Giacomo Ferrari for having been at my side throughout these years, literally.

I also would like to express my gratitude to the Département de Mathématiques et Applications of the École Normale Supérieure of Paris for hosting me during my visit, particularly Bertrand Deroin for many fruitful conversations.

Finally I want to thank the people who have provided me with beneficial distractions from my thesis: Danilo Lewanski for being the kind of friend that brings you on a plane just to push you down; Giulia Saguatti for being the pirate I did not expect to meet; the C.I.T. dog shelter for the hard and satisfying work we have done together; Andrea Ferrari, Federica Rocchi and Martina Rocchi for being a timeless and reliable fixed point. And of course my parents, for their constant support and unconditional love.

References

[1] R. D. M. Accola. “Invariant domains for Kleinian groups”. In: *Amer. J. Math.* 88 (1966), pp. 329–336.

[2] S. Baba. “2π-grafting and complex projective structures, I”. In: *Geom. Topol.* 19.6 (2015), pp. 3233–3287.

[3] A. F. Beardon. *The geometry of discrete groups.* Vol. 91. Graduate Texts in Mathematics. Springer-Verlag, New York, 1983, pp. xii+337.

[4] R. Benedetti and C. Petronio. *Lectures on hyperbolic geometry.* Universitext. Springer-Verlag, Berlin, 1992, pp. xiv+330.

[5] L. Bers. “On boundaries of Teichmüller spaces and on Kleinian groups. I”. In: *Ann. of Math. (2)* 91 (1970), pp. 570–600.

[6] L. Bers. “Simultaneous uniformization”. In: *Bull. Amer. Math. Soc.* 66 (1960), pp. 94–97.

[7] M. R. Bridson and A. Haefliger. *Metric spaces of non-positive curvature.* Vol. 319. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xxii+643.

[8] G. Calsamiglia, B. Deroin and S. Francaviglia. “Branched projective structures with Fuchsian holonomy”. In: *Geom. Topol.* 18.1 (2014), pp. 379–446.

[9] G. Calsamiglia, B. Deroin and S. Francaviglia. “The oriented graph of multi-graftings in the Fuchsian case”. In: *Publ. Mat.* 58.1 (2014), pp. 31–46.

[10] G. Calsamiglia, B. Deroin, V. Heu and F. Loray. “The Riemann-Hilbert mapping for $\mathfrak{sl}\_2$ -systems over genus two curves”. In: *ArXiv e-prints* (Feb. 2016). arXiv: 1602.02273 [math.CV].

[11] D. Dumas. “Complex projective structures”. In: *Handbook of Teichmüller theory. Vol. II.* Vol. 13. IRMA Lect. Math. Theor. Phys. Eur. Math. Soc., Zürich, 2009, pp. 455–508.

[12] H. M. Farkas and I. Kra. *Riemann surfaces.* Second. Vol. 71. Graduate Texts in Mathematics. Springer-Verlag, New York, 1992, pp. xvi+363.

[13] D Gallo, M. Kapovich and A. Marden. “The monodromy groups of Schwarzian equations on closed Riemann surfaces”. In: *Ann. of Math. (2)* 151.2 (2000), pp. 625–704.

[14] D. M. Gallo. “Prescribed holonomy for projective structures on compact surfaces”. In: *Bull. Amer. Math. Soc. (N.S.)* 20.1 (1989), pp. 31–34.

[15] W. M. Goldman. "Locally homogeneous geometric manifolds". In: *Proceedings of the International Congress of Mathematicians. Volume II.* Hindustan Book Agency, New Delhi, 2010, pp. 717–744.

[16] W. M. Goldman. "Projective structures with Fuchsian holonomy". In: *J. Differential Geom.* 25.3 (1987), pp. 297–326.

[17] W. M. Goldman. "Topological components of spaces of representations". In: *Invent. Math.* 93.3 (1988), pp. 557–607.

[18] D. A. Hejhal. "Monodromy groups and linearly polymorphic functions". In: *Acta Math.* 135.1 (1975), pp. 1–55.

[19] N. J. Hitchin. "The self-duality equations on a Riemann surface". In: *Proc. London Math. Soc. (3)* 55.1 (1987), pp. 59–126.

[20] A. T. Huckleberry and G. A. Margulis. "Invariant analytic hypersurfaces". In: *Invent. Math.* 71.1 (1983), pp. 235–240.

[21] Y. Imayoshi and M. Taniguchi. *An introduction to Teichmüller spaces.* Translated and revised from the Japanese by the authors. Springer-Verlag, Tokyo, 1992, pp. xiv+279.

[22] M. Kapovich. *Hyperbolic manifolds and discrete groups.* Vol. 183. Progress in Mathematics. Birkhäuser Boston, Inc., Boston, MA, 2001, pp. xxvi+467.

[23] F. Loray and D. Marín Pérez. "Projective structures and projective bundles over compact Riemann surfaces". In: *Astérisque* 323 (2009), pp. 223–252.

[24] R. Mandelbaum. "Branched structures and affine and projective bundles on Riemann surfaces". In: *Trans. Amer. Math. Soc.* 183 (1973), pp. 37–58.

[25] R. Mandelbaum. "Branched structures on Riemann surfaces". In: *Trans. Amer. Math. Soc.* 163 (1972), pp. 261–275.

[26] R. Mandelbaum. "Unstable bundles and branched structures on Riemann surfaces". In: *Math. Ann.* 214 (1975), pp. 49–59.

[27] B. Maskit. "On a class of Kleinian groups". In: *Ann. Acad. Sci. Fenn. Ser. A I No.* 442 (1969), p. 8.

[28] B. Maskit. "On boundaries of Teichmüller spaces and on Kleinian groups. II". In: *Ann. of Math. (2)* 91 (1970), pp. 607–639.

[29] R. Miranda. *Algebraic curves and Riemann surfaces.* Vol. 5. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 1995, pp. xxii+390.

[30] S. Nag. "Schiffer variation of complex structure and coordinates for Teichmüller spaces". In: *Proc. Indian Acad. Sci. Math. Sci.* 94.2-3 (1985), pp. 111–122.

[31] H. Poincaré. "Sur l'uniformisation des fonctions analytiques". In: *Acta Math.* 31.1 (1908), pp. 1–63.

[32] S. P. Tan. "Branched -structures on surfaces with prescribed real holonomy". In: *Math. Ann.* 300.4 (1994), pp. 649–667.

[33] W. P. Thurston. *Three-dimensional geometry and topology. Vol. 1.* Vol. 35. Princeton Mathematical Series. Edited by Silvio Levy. Princeton University Press, Princeton, NJ, 1997, pp. x+311.

[34] M. Troyanov. "Prescribing curvature on compact surfaces with conical singularities". In: *Trans. Amer. Math. Soc.* 324.2 (1991), pp. 793–821.