THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

Riemann surfaces with symmetry: algorithms and applications

T. P. Northover

Doctor of Philosophy
University of Edinburgh
2010

ABSTRACT

Riemann surfaces frequently possess automorphisms which can be exploited to simplify calculations. However, existing computer software (Maple in particular) is designed for maximum generality and has not yet been extended to make use of available symmetries. In many calculations, the symmetries can be most easily used by choosing a specific basis for $H_1(\Sigma, \mathbb{Z})$ under which the automorphism group acts neatly. This thesis describes a Maple library, designed to be used in conjunction with the existing `algcurves`, which allows such a choice to be made. In addition we create a visual tool to simplify the presentation of Riemann surfaces as sheeted covers of $\mathbb{C}$ and the creation of homology bases suitable for use in the Maple library.

Two applications are considered for these techniques, first Klein's curve and then Bring's. Both of these possess maximal symmetry groups for their genus, and this fact is exploited to obtain neat algebraic homology bases. In the Klein case the basis is novel; Bring's is derived from work in the hyperbolic setting by Riera. In both cases previous hyperbolic work and calculations are related to the algebraic results. Vectors of Riemann constants are calculated for both curves, again exploiting the symmetry.

Finally this thesis moves back to simpler cases, and presents a general algorithm to convert results from general genus 2 curves into results based on a symmetric basis if one exists. This is applied to algebraic and numeric examples where we discover an elliptic surface covered in each case.

ACKNOWLEDGEMENTS

I would first like to thank my supervisor, Professor Harry Braden, for his constant encouragement and suggestions throughout my time in Edinburgh. Without his invaluable advice this thesis could never have been completed.

I found the lunchtime discussions in the department invaluable, providing insight into many diverse aspects of the world. Thankyou!

I would also like to thank my family for their support. My aunt and uncle, Sula and Niall, kept me sane through the four years here. And my parents, Anna and Pete gave me my initial interest in mathematics, and encouragement right through to the end of this project.

DECLARATION

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

The results of Chapter 4 exist as `arXiv:0905.4202v1` and have been submitted for publication in collaboration with my supervisor Harry Braden.

(T. P. Northover)

CONTENTS

1

INTRODUCTION

Riemann surfaces, compact one dimensional complex manifolds, have been studied for more than a century and have a rich theory, bringing together disparate branches of mathematics. Algebra, topology, analysis and others have found application in revealing aspects of these fascinating objects.

The simplest Riemann surfaces are the elliptic curves with genus 1, which provide analytic answers to many seemingly simple yet intriguing problems. From the motion of a pendulum to the solution of a quintic equation, classical problems find their home with the elliptic functions, as for example in [1].

More recently, these surfaces have played a fundamental role in providing explicit solutions to integrable systems. As discussed by Babelon and Talon in [2], problems can often be recast as a Lax system with spectral parameter

$$\frac{\mathrm{d}L(\lambda)}{\mathrm{d}t} = [L(\lambda), M(\lambda)],$$

where L and M are matrix functions of the dynamical variables and $[,]$ is a matrix commutator. In this case Babelon shows that the spectral curve

$$\Sigma : \det(L(\lambda) - \mu\mathbf{1}) = 0,$$

a Riemann surface, is invariant under the motion and the eigenvectors of $L - \mu\mathbf{1}$ provide a line bundle on Σ which will vary with time. One can often show, as for example in [19], that this line bundle $D(t)$ has linear motion in the Jacobian

$$D(t) = \alpha t + \beta,$$

where α is a function of the system and β obviously encodes the initial conditions. Remarkably, this divisor is enough to characterise the eigenvectors of L up to normalization; and so Θ-functions (specifically the ability to give a meromorphic function on Σ with prescribed poles) allow Babelon and Talon to write down the eigenvectors that originally gave $D(t)$.

Finally, these Θ-functions (together with the spectral data encoded in Σ itself) are enough to reconstruct L itself, generally with each entry a rational function of Θ-functions of time. From L one can usually recover the original dynamical variables by algebraic manipulation.

Via these integrable systems, a path has been made to modern particle physics. In

particular explicit monopole solutions in various backgrounds reduce to well-known integrable systems with Lax representations, as in [3]. Discovering analytic and global properties of solutions obtained like this leads to further consideration of the surfaces involved.

The Maple software package already has a module called `algcurves` which provides routines for dealing with Riemann surfaces expressed as plane algebraic curves. Many of these functions provide complete information in their realm, however on the analytic side they lack flexibility. In particular the choice of basis for the first homology group is an implicit variable in the Θ-functions, via the period matrix of the surface. Maple chooses this basis behind the scenes, and while it can be exposed to scrutiny it is nevertheless fixed in all calculations Maple undertakes. The disadvantage of this approach is that in cases where the surface has holomorphic automorphisms there may be preferred bases which reflect the symmetry in the associated analytic objects as described (for example) by Fay in [16]. The algorithm used by Maple to select its basis was designed for use on completely general surfaces, and so the results obtained frequently don't reach their potential simplicity when symmetries are available.

The first goal of this thesis will be to augment the `algcurves` library with functions that allow an explicit homology basis to be used. The result is the Maple library `extcurves`, built around the new function `isect` which can calculate the intersection number of two paths on a given Riemann surface. The intersection number serves as an analogue of an inner product in the homology group and from this single ability follow a surprising number of functions; a small step tells us how to find the basis-change matrix for these canonical bases, γ_i and η_i. If $\gamma_i = M_{ij}\eta_j$ then

$$\langle \gamma_i, \eta_j \rangle = \langle M_{ik}\eta_k, \eta_j \rangle$$
$$= M_{ik}\langle \eta_k, \eta_j \rangle$$
$$= M_{ik}J_{kj},$$

where J is the matrix used to define the symplectic group, given in block form by

$$J = \begin{pmatrix} 0 & \mathbf{1} \\ -\mathbf{1} & 0 \end{pmatrix}.$$

Hence

$$M_{ij} = -\langle \gamma_i, \eta_k \rangle J_{kj},$$

which is a trivial extension of the ability to calculate intersection numbers. Applying this technique to Maple's own homology choice will allow us to calculate the period matrix itself in an arbitrary basis.

The `extcurves` library, considered alone, goes beyond flexibility and actually makes seemingly easy tasks rather difficult; we now need a good way to manually specify a homology basis. The second part of this thesis tackles this problem. A purely algebraic description is too error-prone and tedious for many purposes, so a new Java program called `CyclePainter` is developed which allows homology paths to be visually drawn and conveniently displays essential information about their analytic continuations so that monodromy properties are respected.

The next step is to apply these techniques we have developed to some examples. First comes Klein's beautiful quartic curve. As is well known, this is a genus 3 surface with a symmetry group of 168 elements – the maximum possible for that genus. There have been frequent calculations of its period matrix in the literature, though rarely making full use of its symmetry group. We continue this tradition and produce a lovely period matrix of the form

$$\tau = \frac{1}{2} \begin{pmatrix} e & 1 & 1 \\ 1 & e & 1 \\ 1 & 1 & e \end{pmatrix},$$

where $e = \frac{1+\mathrm{i}\sqrt{7}}{4}$. This result is published jointly with Harry Braden in the Journal of Physics A: Mathematical and Theoretical [4]. Period matrices with comparable levels of symmetry have been discovered in the past, but not with the additional benefits provided by rational numbers off-diagonal (this added feature permits certain calculations to make use of the fact that this curve covers others, for example the theta functions factorize into lower genus versions, as noted by Martens in [26] for example).

After that, we relate this period matrix and basis to one constructed by Rauch and Lewittes in [29]. This involves a detour into hyperbolic models of Riemann surfaces and makes full use of the group of symmetries. As a final component to the study of Klein's surface, we derive the vector of Riemann constants with respect to our new particularly symmetric basis.

After Klein's curve we move on to another surface with a maximal symmetry

group. This time Bring's curve is studied, a genus 4 curve where Riera and Rodríguez have already produced (in [30] a period matrix with good symmetry properties in the hyperbolic model of the surface. As in the Klein case, we find a novel algebraic representation of the basis producing this matrix and compute the vector of Riemann constants again. In doing this we will have course to create a new relationship between algebraic versions of Bring's curve studied by Dye [13] and Craig [9].

Finally we change direction to deal with simpler surfaces: those with genus 2. In this particular instance an algorithm is derived which automates the selection of a symmetric basis. If such a curve covers an elliptic curve then, as Murabayashi showed in [28], there is a particularly symmetric form of the period matrix. This matrix then allows the genus 2 Θ-functions to factorise and be written completely in terms of elliptic Θ-functions, as Martens noted in [26]. We present a sequence of algorithms that, given a genus two period matrix, finds such covers and searches for the basis transformation needed to manifest the symmetry.

2

EXTCURVES: INTERSECTIONS OF HOMOLOGIES

2.1   Introduction

Source code for the software described in this chapter can be downloaded from `http://gitorious.org/riemanncycles/extcurves`; a snapshot is on the included CD.

The matrix of periods for a Riemann surface is defined for a given canonical homology basis $\gamma_i = \mathfrak{a}_1, \ldots, \mathfrak{a}_g, \mathfrak{b}_1, \ldots, \mathfrak{b}_g$ and basis ω_i of holomorphic differentials as

$$\Pi_{ij} = \int_{\gamma_i} \omega_j = \begin{pmatrix} \mathcal{A} \\ \mathcal{B} \end{pmatrix}.$$

The $\mathfrak{a}$-normalised Riemann period matrix is defined by

$$\begin{pmatrix} \mathbf{1} \\ \tau \end{pmatrix} = \Pi \mathcal{A}^{-1} = \begin{pmatrix} \mathbf{1} \\ \mathcal{B}\mathcal{A}^{-1} \end{pmatrix},$$

and is equivalent to calculating the matrix of periods with the so-called $\mathfrak{a}$-normalised basis for the differentials, uniquely defined as $\omega\mathcal{A}^{-1}$. Thus the period matrix τ, and hence Θ-functions, do depend on the choice of homology basis but not on any choice of basis for the differentials.

Tretkoff and Tretkoff produced an algorithm in [31] which finds such a homology basis for a general surface presented as a plane algebraic curve. This was implemented under Maple by Deconinck and van Hoeij (practical considerations are described in [11]) and allows Maple to fully automatically compute a period matrix for an arbitrary curve.

However, not all curves are arbitrary and when they possess symmetries it is advantageous to use a homology basis where the action of those symmetries is neat. We will see that such a choice can produce dramatic simplifications in the resulting period matrix, and arguably even more reduction is possible in Θ-functions.

These symmetric calculations are currently entirely manual and some form of computer aid in working with symmetric homology bases would be desirable. The ultimate goal would obviously be a fully automated choice of the best homology basis available. This inevitably runs into the aesthetic issue of which basis actually is best so more limited goals must be set. However, work has been done at least on finding *good* bases and determining what should be expected. Gilman considers in [17] and [18] the possible actions of automorphisms on canonical homology bases. Her results are only directly applicable when the order of the automorphism is a prime and are suggested by the Riemann-Hurwitz formula. Suppose $\hat{\Sigma}$ is a Riemann surface of genus $\hat{g}$ with an

automorphism $\phi : \hat{\Sigma} \to \hat{\Sigma}$ of order p. We can form the quotient surface $\Sigma = \hat{\Sigma}/\langle \phi \rangle$ with genus g and natural projection $\pi : \hat{\Sigma} \to \Sigma$. The Riemann-Hurwitz formula is

$$2\hat{g} - 2 = deg(\pi)(2g - 2) + \sum b_q.$$

where b_i is called the branching number and is given in this case by

$$\sum_{p \in \pi^{-1}(q)} (|\{\text{orbit of } p \text{ under } \phi\}| - 1) .$$

For all but a finite number of points (the so-called branch points, see later for more details) this is finite. Because the order of ϕ is a prime, the branching numbers are particularly easy to calculate here: the size of an orbit under ϕ is either 1 or p and so b_i is either $p - 1$ (for a fixed point) or 0 (for a regular point). So, as Gilman notes in [18], the formula can be rewritten

$$2\hat{g} - 2 = p(2g - 2) + (p - 1)f,$$

where f is the number of fixed points of ϕ; this can be suggestively rearranged to

$$2\hat{g} = 2gp + (f - 2)(p - 1).$$

Gilman shows that there is a canonical homology basis for $\hat{\Sigma}$ with $2gp$ well-behaved cycles (falling into $2g$ ϕ-orbits) and $(f - 2)(p - 1)$ anomalous cycles.

This result is obviously of particular interest if there are precisely 2 fixed points, since then there then are no anomalies to break the pattern. We will soon see that this actually occurs for Klein's curve (with an order 3 automorphism) and is responsible for a great deal of the beauty in the results obtained.

In principle Gilman's results could be automated. However, given that they are obtained via a detour into Fuchsian groups of surfaces and the requirement for an automorphism of prime order, it is unclear whether automation would be useful. Certainly ad-hoc methods of exploiting symmetry have proven adequate to the surfaces at hand.

Here we present a Maple library which is designed to be used in conjunction with the existing `algcurves` routines and provides the desired flexibility in homology choice.

2.2   Representation of Riemann surfaces

There are many ways to describe Riemann surfaces mathematically, and many ways to represent each of those descriptions computationally. In this section we will fix both of these parameters and justify the decisions made.

Some of the more obvious options for a mathematical description are:

1. An abstract 1 dimensional complex manifold, perhaps given as a set of coordinate change functions making up an atlas.

2. A quotient of the Poincaré hyperbolic disc with some Fuchsian group. This approach has been computationally taken by Buser and Silhol, for example in [8].

3. An algebraic variety in $\mathbb{P}^2$, that is, for a given homogeneous polynomial f in three variables

$$\Sigma : \{[x, y, z] \in \mathbb{P}^2 \mid f(x, y, z) = 0\}.$$

   The surface is then defined completely by the polynomial f, although for a complete understanding more work may be needed to resolve singularities.

4. Relatedly, we could choose $[x, y, 1]$ to represent (nearly) all points above. As a primitive description this amounts to, given a (not necessarily homogeneous) polynomial f,

$$\Sigma : \{(x, y) \in \mathbb{C}^2 \mid f(x, y) = 0\}.$$

   Compared to the previous option, this leaves the points at ∞ implicit, so we would have to revert to projective space to describe them. However, most other points become slightly simpler.

Our choice is guided by the existing Maple `algcurves` library. Most of its functions use the fourth representation above, taking three parameters f, x, y where f is an algebraic expression in the symbols x and y. In principle the symbols x and y could be derived from f, since it should only have two indeterminates. However, in use they are usually distinguished so they are specified explicitly. Since interoperability is a worthy goal we choose to follow this choice, with a slight modification. Maple provides a "Record" data-type to package related properties of what is conceptually a single object. We make use of this in our programming interface; a Riemann surface will be provided as a Record with three fields: `f, x, y`.

For example the Riemann surface

$$y^2 = x^4 + 1$$

may be given by the definition

```
> curve := Record('f'=y^2-x^4-1, 'x'=x, 'y'=y):
```

Of course variables may be renamed with impunity. The following would be completely equivalent

```
> curve := Record('f'=w^2-z^4-1, 'x'=z, 'y'=w):
```

However, as mentioned above, swapping the variables as in

```
> curve := Record('f'=w^2-z^4-1, 'x'=w, 'y'=z):
```

would be subtly different. It would still represent the same surface, but functions would operate on it in different ways.

With this convention normal `algcurves` functions can obviously still be called. For example we might write

```
> genus(curve:-f, curve:-x, curve:-y);
1
```

2.2.1 Sheets and branches

Having made this choice, we need to consider the consequences of the representation in more detail. In particular, how should we describe a point on the Riemann surface?

We start by defining the projection maps π_x and π_y from Σ to $\mathbb{C}$ by

$$\pi_x([x, y, 1]) = x,$$
$$\pi_y([x, y, 1]) = y.$$

Disregarding the singular points and points at infinity (i.e. those of the form $[x, y, 0]$ in projective space) – both finite in number – if given, say, an arbitrary $x = \alpha \in \mathbb{C}$ then the equation

$$f(\alpha, y) = 0$$

is a univariate polynomial in y with finitely many solutions. For generic α this will be a constant $\deg(f)$, and we say that the surface has $\deg(f)$ sheets (these notions

correspond to more general ideas from differential topology, the degree has the same name, but our sheets correspond to the topologists' fibres). In a neighbourhood of one of these regular points π_x itself is a valid coordinate chart on the curve as a manifold.

At some special points of $\mathbb{C}$, called branch points, some of these solutions will merge and there will be fewer than $\deg(f)$ solutions. Suppose x_0 is such a branch point; then π_x is not necessarily a valid chart: for at least one preimage P of x_0 we will have points Q_1 and Q_2 arbitrarily close to P such that $\pi_x(Q_1) = \pi_x(Q_2)$. That is, on no open neighbourhood of P does π give a bijection to (a subset of) $\mathbb{C}$.

In these cases deriving a coordinate chart requires further study. We will usually avoid such issues because our main concern, paths on the surface, can be deformed so that they never go near a branch point.

We usually label this set of branch points $\mathcal{B}$. There is a strongly related notion of a ramification point (essentially the set $\pi_x^{-1}(b)$ for a branch point), but we won't use this concept directly.

Note that in the more general topological setting branch points are referred to as "critical values". Literature abounds dealing with these issues from both sides, on the Riemann surface side Farkas and Kra have written the particularly good [15], while the more general setting is discussed by Lee in [25].

2.3 Representation of homology cycles

Now that we have fixed a representation for surfaces, we need to decide on the homology cycles that live on them. By definition $H_1(\Sigma, \mathbb{Z})$ is the set of equivalence classes of paths on a surface Σ under homological equivalence. Naturally enough we only demand a single path from the user to represent its entire equivalence class.

Unfortunately there will be a tension in the remaining decisions between allowing potential users expressive freedom to describe the paths they want and keeping computations manageable.

In principle paths could be as complex and pathological as desired; after all the only requirement is continuity. One could even devise curves that intersect at a given point if and only if the Goldbach conjecture is true [1]; naturally, restricting this freedom is a priority but where exactly to draw the line is a difficult issue. Considering piecewise sections of paths, there are three obvious classes to consider:

---

[1]Describing such a path for Maple would present a challenge, fortunately.

Figure 2.1: Treatment of semicircular arcs encountered in `algcurves` paths

1. Straight lines. Deciding whether two line segments intersect is trivial, but many paths won't naturally be viewed as straight lines so some kind of conversion is needed.

2. Arcs of circles. It is slightly more difficult to discover whether circles intersect and additionally the circle/line and circle/circle cases need to be considered separately. On the other hand they are used internally by `algcurves` to describe its homology cycles, so some handling will be essential for interoperability.

3. The rest. No general way to determine whether arbitrary paths intersect.

Initially the first two were permitted because the alternative of computationally splitting an arbitrary circle into an equivalent path made from line segments is intricate and could also lead to less efficient analysis. This facility tended to be used only for compatibility with `algcurves` since a human is quite capable of splitting a nominally curved path into visual segments. Fortunately it was noted that `algcurves` only uses semicircular arcs centred on branch points with no others inside so any circle emerging from `algcurves`' algorithms could be trivially replaced by two line segments as in Figure 2.1.

Now suppose our path is given by two functions $x, y : [0, 1] \to \mathbb{C}$ satisfying the equation $f(x(t), y(t)) = 0$ for all t. Provided this path does not approach a branch point, only $y(0)$ is needed since continuity uniquely determines y for every other value.

This motivates the second constraint we put on users: no path is allowed to pass through a branch point – in fact they should stay far enough away that unique analytic continuation works even with the vicissitudes of floating point arithmetic.

The advantages from a computational point of view are obvious: with the earlier restriction to segments our paths will be described as in Figure 2.2 simply by $n + 1$ complex numbers $x_0, \ldots, x_n$ and y_0. For a closed cycle we would certainly demand $x_0 = x_n$ but also that the induced value of y at x_n matches y_0 – i.e. that we return on the same sheet.

Figure 2.2: Description of homology cycle for `extcurves`

Further, the algorithms we will employ are themselves greatly simplified by not having to consider whether paths pass through branch points. If we allowed paths through branch points then we would have to allow intersections there, which is a far more complicated issue: we would need knowledge of exactly how sheets come together (the monodromy parameters) to even make a start. Similarly, naïvely integrating even a holomorphic differential along a path through a branch point can fail. We would have to provide a true manifold coordinate in a neighbourhood of the branch in order to make the differential numerically finite. No matter how useful the demand that we avoid branch points is for implementation purposes, it is probably more onerous for users of the algorithms than the demand that all paths be given as straight segments. Many homology bases have traditionally been given as prescriptions for selected paths between branch points, for example in [3]. These have to be carefully modified before our algorithms become usable.

In the end, the representation chosen for a cycle is as follows. We define a data-type "Segment" whose constructor takes a beginning and an endpoint. We call the fundamental data type an `extpath`, and define it as a list with the following structure

- The first element is y_0.

- Subsequent elements are `Segment` objects, together giving the path through the x plane.

For example if `x[0]`,...,`x[3]` and `y0` are previously defined as complex numbers then

```
> cycle := [y0, Segment(x[0], x[1]),
            Segment(x[1],x[2]),
            Segment(x[2],x[3])]:
```

would construct a valid `extpath` (and possibly a cycle).

There is redundancy in this notation if it is being entered manually, but we will make things easier with utility functions in the programming interface. The advantage is that Segment can be a reasonably complicated object with some awareness of how Riemann surfaces behave and the sheet it's on which simplifies implementation of some algorithms.

In fact we will construct an entire utility program to mitigate the effects of these restrictions by presenting a visual interface for the construction of cycles and almost never have to deal with their internal representation (see Chapter 3).

Meanwhile, `extpath_is_closed` provides a useful consistency check. It determines whether a given extpath represents a *bona fide* homology cycle or actually just a path on Σ. If that function returns `false` then there is no hope of any other algorithms producing meaningful results.

2.4   Intersecting homology cycles

Now we have fixed upon a representation for both surfaces and cycles, we must turn our attention to the fundamental purpose of `extcurves`: calculating the intersection number of two cycles. Our restriction to piecewise linear paths has simplified matters greatly, but issues remain at the nodes of paths as we shall soon see.

In this section, given two paths γ and η we will denote the true intersection number with angle brackets, $\langle \gamma, \eta \rangle$, and the value computed algorithmically as $\mathrm{isect}(\gamma, \eta)$.

The first hurdle to overcome is the sheeted nature of our representation for Riemann surfaces. We can certainly decide whether our cycles intersect as piecewise linear paths in $\mathbb{C}$, but if the paths have differing values for y at that point they're in completely different areas of the Riemann surface – the putative intersection should be ignored.

With that in mind, a simplistic first attempt is given in Algorithm 1. The repeated analytic continuations may seem costly, indeed even in the real code they consume the bulk of the run-time of the algorithm, but the values can easily be cached to mitigate this redundancy.

In implementing this scheme three nontrivial matters arise

1. Performing the analytic continuation. It is actually very difficult to determine for certain whether this has been done correctly if the paths get close to a branch point. Fortunately this is one place where the existing `algcurves` comes to the

---

Algorithm 1 Simplistic intersection algorithm

Require: Two homology cycles: $\gamma = [y_0, Seg(x_0, x_1), \ldots, Seg(x_{n-1}, x_n)]$ and $\eta = [w_0, Seg(z_0, z_1), \ldots, Seg(z_{m-1}, z_m)]$, together with the surface they live on.

  for all $l_1 = Seg(a_1, b_1) \in \gamma$ **do**
    for all $l_2 = Seg(a_2, b_2) \in \eta$ **do**
      if l_1 and l_2 intersect at X **then**
        $Y \leftarrow$ analytic continuation of y_0 along γ to X
        $W \leftarrow$ analytic continuation of w_0 along η to X.
        if $Y = W$ **then**
          Add intersection with appropriate orientation.
        end if
      end if
    end for
  end for

---

rescue. Some of its own routines face the same problem, so it has an (internal) function `algcurves/Acontinuation` which does what we need.

2. Intersections could occur at the endpoints of segments as in Figure 2.3 which, depending on how we decide whether two segments intersect, may be counted twice. We might hope to counter this by considering our segments to be open at one end as sets, but special case code would still be needed to determine the orientation (and even existence) of an intersection at a node; see for example Figure 2.3b. If we had allowed branch points in a path this problem would have been even more intricate because the angles involved wouldn't be derived purely from the projection onto $\mathbb{C}$.

3. Worse, two segments could be entirely or partially coincident; this may seem pathological, but it actually occurs naturally when we want to consider two paths that are basically the same but on different sheets. This problem actually occurs in normal definitions of the intersection number, for example in the book by Griffiths and Harris [20]. One resolution is that perturbations of the curves remove the issue – any pair of cycles is homotopically equivalent to a pair with only transverse intersections (i.e. there is a continuous map $h : \mathbb{R} \times \mathbb{R} \to \Sigma$ such that $h(0, t) =$ original curve and $h(1, t) =$ transverse curve) and the intersection number is preserved under homotopical equivalence. We will adopt a slightly different approach here.

The latter two problems can be solved by extending the definition of isect (inherently given only for paths with well-defined tangents intersecting transversely at a point)

(a) Nodal intersection   (b) Nodal non-intersection

Figure 2.3: Problematic intersections

Figure 2.4: Split segments at any intersection

to these more complex cases where segments meet at their endpoints. Indeed the key insight is that each endpoint of coincident segments can be considered independently and assigned a value $\pm\frac{1}{2}$ (as indicated in Figure 2.5); when these half-integers are added up the true intersection number is obtained.

As a very first step we make sure that either all segments involved in an intersection terminate there or none do. If there is an intersection point where a segment ends then we split every other segment passing through that point as in Figure 2.4. This transformation is trivially valid and reduces the number of special cases that need to be considered.

Next we want a result classifying the types of intersection that can occur

Theorem 1. *Each nodal intersection point falls into one of the categories in Figure 2.5, where coincident lines have been replaced by parallel ones and the single intersection point has been expanded to a circle for clarity.*

Proof. This is simple exhaustion. We potentially have four lines meeting at this point:

- If all four segments are distinct then one of the diagrams in Figure 2.5a or 2.5b

(a) 4 directions, intersecting (b) 4 directions, no intersection   (c) 3 directions

(d) 3 directions   (e) 2 directions, double lines (f) 2 directions, double lines

(g) 2 directions, triple line   (h) single direction

Figure 2.5: Types of intersection and contribution at a node

applies.

- If there are three distinct lines then one must be doubled as in Figure 2.5c or 2.5d.

- With just two distinct segments they either fall into two pairs (Figures 2.5e and 2.5f) or a triple line and a single line (Figure 2.5g).

- Finally if all lines are coincident then only Figure 2.5h can apply.

$$\square$$

Now that we know what kinds of intersection can occur at the end of segments, we turn our attention to algorithmically calculating the intersection number from this information. The basic goal will be to successively reduce the general case to simpler and simpler intersections by altering the paths slightly, while keeping both $\langle \bullet, \bullet \rangle$ and $\mathrm{isect}(\bullet, \bullet)$ constant.

These alterations are reminiscent of the Reidemeister moves used for similar simplifications of in knot theory, but with a slightly more concrete numeric twist, suitable for our eventual implementation on a computer.

But first we must actually define isect as follows.

Definition 2. *Suppose two paths, γ and η, meet at a point. Without loss of generality (rotation does not affect intersection number) we may assume γ comes in at angle 0, out at γ_{out} and similarly η comes in at η_{in} and out at η_{out}. We demand $\gamma_{out}, \eta_{in}, \eta_{out} \in [0, 2\pi)$ as shown in Figure 2.6. Then the local contribution to* isect *is*

$$\frac{1}{2}\big[\, \mathrm{sgn}(\eta_{out})\,\mathrm{sgn}(\eta_{out} - \gamma_{out}) + \mathrm{sgn}(\eta_{in})\,\mathrm{sgn}(\gamma_{out} - \eta_{in})\,\big]. \tag{2.1}$$

This definition simply captures the intuitive notion of intersection number already annotated in Figure 2.5: 4-way intersections contribute ± 1 or 0; segments that double-back contribute 0 and other coincident segments contribute $\pm\frac{1}{2}$ or 0. In particular note that it changes sign under both $\gamma \leftrightarrow \eta$ and changing the direction of either path, as one would expect.

For example, changing the direction of η (for simplicity) induces the map

$$(\gamma_{out}, \eta_{in}, \eta_{out}) \mapsto (\gamma'_{out}, \eta'_{in}, \eta'_{out}) = (\gamma_{out}, \eta_{out}, \eta_{in}).$$

Figure 2.6: Angles in local definition of isect.

Then

$$\mathrm{sgn}(\eta'_{out})\,\mathrm{sgn}(\eta'_{out} - \gamma'_{out}) = \mathrm{sgn}(\eta_{in})\,\mathrm{sgn}(\eta_{in} - \gamma_{out})$$

$$= -\,\mathrm{sgn}(\eta_{in})\,\mathrm{sgn}(\gamma_{out} - \eta_{in}),$$

$$\mathrm{sgn}(\eta'_{in})\,\mathrm{sgn}(\gamma'_{out} - \eta'_{in}) = \mathrm{sgn}(\eta_{out})\,\mathrm{sgn}(\gamma_{out} - \eta_{out})$$

$$= -\,\mathrm{sgn}(\eta_{out})\,\mathrm{sgn}(\eta_{out} - \gamma_{out}),$$

and so half their sums simply changes sign as needed.

Formally, the main argument will be an induction on the number and types of bad intersections as in Figure 2.5. Two more definitions will be useful to make the argument concrete.

Definition 3. *Let the segments of γ be given by $\gamma = (\gamma_1, \ldots, \gamma_m)$, and similarly $\eta = (\eta_1, \ldots, \eta_n)$. The* nontransversality index *is the triple (d, c, b) where*

$$d = |\{\textit{Number of segments double-backing at an intersection}\}|,$$

$$c = |\{\textit{Number of coincident segments on different paths }\}|,$$

$$= |\{(i, j) : \gamma_i = \eta_j\}|$$

$$b = |\{\textit{Number of intersections at segment endpoints}\}|.$$

Note that since we have cut segments to guarantee that either all or no segments

terminate at a given intersection point, two segments are equal if and only if they have a coincident section. With a lexicographic ordering, this is the parameter we will perform the induction over. The second definition clarifies the type of alteration we will need for the induction to proceed.

Definition 4. *A* reduction *is a scheme for altering paths (not necessarily always applicable) such that, if it requires $\gamma \mapsto \gamma'$ and $\eta \mapsto \eta'$ then*

1. *The new paths are homologically equivalent to the originals; so $\langle \gamma, \eta \rangle = \langle \gamma', \eta' \rangle$.*

2. *The calculated intersection number does not change under the alteration; that is* $\operatorname{isect}(\gamma, \eta) = \operatorname{isect}(\gamma', \eta')$.

3. *The alteration strictly reduces the nontransversality index of the two paths.*

The theorem we would like to prove can then be stated as

Theorem 5. *Suppose that for any two paths with nonzero nontransversality index (d, c, b) we can provide an applicable reduction. Then the scheme presented in Definition 2 calculates the true intersection number. That is*

$$\forall \gamma, \eta \quad \langle \gamma, \eta \rangle = \operatorname{isect}(\gamma, \eta).$$

Proof. Consider two paths γ, η. If $(d, c, b) = (0, 0, 0)$ then the usual intersection rules apply and the intersection number is trivially correct.

If not, then some reduction $\gamma \mapsto \gamma'$, $\eta \mapsto \eta'$ applies. Then

$$\begin{aligned}
\langle \gamma, \eta \rangle &= \langle \gamma', \eta' \rangle \\
&= \operatorname{isect}(\gamma', \eta') \\
&= \operatorname{isect}(\gamma, \eta).
\end{aligned}$$

The first equality follows from the first requirement of a reduction; the second follows from the induction hypothesis (which is applicable because of the third requirement for a reduction); the third follows from the second requirement for a reduction. □

So now all that remains is to provide a reduction applicable to every case. Essentially we will treat each component of the nontransversality index individually and find a reduction that decreases it without affecting earlier ones.

Figure 2.7: Reducing an intersection with a double-backing segment to a simpler case.

2.4.1 Reduction to eliminate double-backs

Suppose γ and η have at least one segment that double-backs on itself at an intersection (i.e. an intersection from one of Figures 2.5d, 2.5e, 2.5g or 2.5h). Then the nonstransversality index is (d, c, b) with $d > 0$.

Consider the following transformation.

Reduction 1. *Given a common node p of γ and η where one segment retraces its steps there will only be finitely many rays meeting at this node, and hence a direction free of segments. Further, the closest branch point or other node will be some finite distance away from p. Extend the segments in question into this safe space so that the actual reversal occurs away from any intersections, as in Figure 2.7.*

Going through the requirements of a reduction in order

1. The section we are adding is homologically equivalent to 0, so both modified paths are equivalent to the originals as required.

2. In terms of contributions to isect, we have removed some number of double-backs each with contribution 0 and replaced each with two identical contributions in opposite directions which cancel out. Thus isect is preserved.

3. This transformation takes an nonstransversality index (d, c, b) to (d', c', b') where $d' < d$, a strict reduction.

Therefore this transformation is indeed a reduction, applicable whenever $d > 0$.

2.4.2 Reduction to eliminate coincident segments

The next transformation will apply to pairs γ, η with nonstransversality index $(0, c, b)$ and act to reduce the number, c, of coincident segments. This transformation isn't quite

Figure 2.8: Displacement transformation at a node

Original contribution	Transformed contribution	
	$\theta_\epsilon > 0$	$\theta_\epsilon < 0$
$-1/2$	0	-1
$+1/2$	$+1$	0
0	$+1/2$	$-1/2$

Table 2.1: Effect of transformation in Lemma 6 on isect.

so local, involving an entire segment including both endpoints; so a small lemma on transformations near a point is helpful

Lemma 6. *Suppose two coincident segments meet at a point p as in the initial diagram of Figure 2.8, both entering at angle 0 and leaving at γ_{out} and η_{out} (both in the range $[0, 2\pi)$) respectively.*

Transform this by moving the common segment of η to enter the circle at angle $\theta_\epsilon \in (-\pi, \pi]$, closer to angle 0 than either γ_{out} or η_{out} as in Figure 2.8.

Then the contribution to isect *is given by Table 2.1. This transformation has* no *effect on a true 4-way intersection at p provided θ_ϵ is small enough.*

Proof. In this case, in the notation of (2.1), $\eta_{in} = 0$ and the equation itself for the original intersection reduces to

$$\frac{1}{2} \operatorname{sgn}(\eta_{out} - \gamma_{out}).$$

After the modification, define $\eta'_{in} = \theta_\epsilon + 2\pi k$ so that it is in the range $[0, 2\pi)$. The local contribution becomes

$$\frac{1}{2}\big[\operatorname{sgn}(\eta_{out} - \gamma_{out}) + \operatorname{sgn}(\gamma_{out} - \eta'_{in})\big] = \frac{1}{2}\big[\operatorname{sgn}(\eta_{out} - \gamma_{out}) + \operatorname{sgn}(\theta_\epsilon)\big],$$

since if $\theta_\epsilon > 0$ then $\gamma_{out} > \eta'_{in}$ and vice-versa ($|\theta_\epsilon|$ is small).

Figure 2.9: Displacement to eliminate coincident segments.

Filling in Table 2.1 is now simple and it is trivial to see that this transformation does not affect the true intersection number. □

We are now in a position to give the reduction applicable in this case.

Reduction 2. *Given two paths γ, η with nontransversality index $(0, c, b)$ choose a common coincident segment L with endpoints p and q.*

Since there are only finitely many branch points and finitely many segments, there are radii r_p and r_q such that $B(p, r_p)$ and $B(q, r_q)$ contain no other significant points. We may further assume that these circles do not intersect.

Similarly, there is a radius r_L such that a parallel displacement of L by a distance less than r_L will intersect with no branches, segments or nodes except possibly within $B(p, r_p)$ or $B(q, r_q)$.

With these quantities displace L on η (but not γ) as in Figure 2.9.

Looking again at the points defining a reduction:

1. By choice of r_p, r_q and r_L this displacement has no effect on the homology of η and hence intersection number.

2. isect is unchanged, as will be shown in Lemma 7.

3. No double-backs are introduced and the number of coincident segments is strictly reduced by this transformation so the nontransversality index will become $(0, c', b')$ with $c' < c$, $b' > b$.

Lemma 7. *The transformation in Reduction 2 has no effect on the calculation of* isect.

Proof. Although we started considering just one local section of the path γ, our modification of η may have affected later or earlier parts of γ passing through either p or q. If the secondary effects are from an intersection not involving a copy of L then the rider to Lemma 6 applies and there is no effect on intersection number. On the other hand,

all intersections involving the common segment can be treated identically and we must consider the contribution from both endpoints, p and q, for the value of isect to remain fixed.

We wish to apply Lemma 6 at both p and q. We can freely rotate each node so that L enters at angle 0 with no effect on isect. We then have to change the direction of one or both paths so that they actually *enter* at that angle. This may change the sign of the total contribution to isect. But compared to p, q needs the direction of both paths changed to satisfy the conditions of the Lemma which means the sign-change will be the same at both p and q.

Now notice for example from Figure 2.9 that at one end of L, $\theta_\epsilon > 0$ and at the other $\theta_\epsilon < 0$. Referring to Table 2.1 this means that the modified contribution to isect comes from adding an entry from the $\theta_\epsilon > 0$ column and one (not necessarily distinct) from the $\theta_\epsilon < 0$ column. The original contribution to isect from both p and q was just the sum of the original contributions from two rows in Table 2.1.

It is easy to see that each possible combination leaves the isect fixed. For example an original $-1/2$ at p and $+1/2$ at q becomes either a -1 at p and a $+1$ at q or a 0 at both: in all cases the total contribution is 0. □

2.4.3 Reduction to eliminate intersections with no coincident segments

Finally we produce a reduction for the nontransversality index $(0, 0, b)$ which replaces any remaining nontrivial intersections with straight lines.

Reduction 3. *Choose an intersection point p with at least one 4-way intersection contributing to b. Let r be such that the disc $B(p, r)$ contains no branch points, no nodes of either path (except p itself) and no segments without an endpoint at p.*

Replace each pair of segments entering $B(p, r)$ at a and leaving at b (hence $a \to p \to b$) with the simple chord from a to b as in Figure 2.10.

Again this is a reduction

1. Because $B(p, r)$ contains no branches the paths are homologically equivalent.

2. Lemma 8 shows that isect remains fixed.

3. After replacement all intersections inside $B(p, r)$ occur on the interior of segments. Further, there are no intersections at all on the boundary of $B(p, r)$ (such an

Figure 2.10: Chord replacement at 4-way intersection

Order	Intersection number
$\gamma_{out} < \eta_{in} < \eta_{out}$	0
$\gamma_{out} < \eta_{out} < \eta_{in}$	0
$\eta_{in} < \gamma_{out} < \eta_{out}$	+1
$\eta_{in} < \eta_{out} < \gamma_{out}$	0
$\eta_{out} < \gamma_{out} < \eta_{in}$	-1
$\eta_{out} < \eta_{in} < \gamma_{out}$	0

Table 2.2: Dependence of intersection number on angle order.

intersection would originally have been a coincident segment), and for the same reason no coincident segments are introduced. Finally we have removed at least one intersection involving segment-ends from p and so the nontransversality index strictly decreases.

It remains to prove

Lemma 8. isect *is unchanged by the transformation in Reduction 3.*

Proof. The simplest method in this case is to simply enumerate all possible anticlockwise orders of η_{in}, η_{out} and γ_{out} as in Definition 2. It is easy to check for each row of Table 2.2 that the nominated intersection number is correct both before and after the transformation. □

These three reductions cover all possible nonzero nontransversality index, and so since a reduction is always possible Theorem 5 holds. An algorithm to implement this scheme is conceptually simple, if fiddly in the details.

The final (not really mathematical) issue encountered is the odd properties of floating point arithmetic. This can result in inconsistent decisions over whether two segments intersect depending on the exact order of calculations, particularly in the

already intricate nodal cases. This manifests itself almost exclusively in calculations based on the x coordinate since we have demanded all paths keep away from branch points which guarantees a certain minimal separation in y – well within the bounds of floating point precision.

In this case the solution is to perform all calculations in x with rational numbers, and in order to retain efficiency we make demands on denominator size which ensure only 32-bit precision is required. This guarantees native CPU instructions can be used for these calculations minimising the overhead in the switch.

The result is a function `isect` with typical use

```
> isect(curve, cycle1, cycle2);
        1
```

2.5 Development of useful functions

2.5.1 `find_homology_transform`

Now that we have a primitive function `isect`, this opens up many avenues of calculation. Most fundamentally, suppose we have two canonical bases for the homology group on a fixed surface Σ

$$\langle \gamma_i, \gamma_j \rangle = J_{ij},$$
$$\langle \eta_i, \eta_j \rangle = J_{ij}.$$

Then there is some symplectic matrix A which converts between the two

$$\gamma_i = A_{ij} \eta_j.$$

Finding the intersection of both sides of this equation with η_k yields

$$\langle \gamma_i, \eta_k \rangle = \langle A_{ij} \eta_j, \eta_k \rangle$$
$$= A_{ij} \langle \eta_j, \eta_k \rangle$$
$$= A_{ij} J_{jk}$$

or, multiplying by J

$$A_{ij} = - \langle \gamma_i, \eta_k \rangle J_{kj}.$$

In particular we can find the basis-change matrix solely with intersection calculations and matrix multiplication. `extcurves` implements this formula as the function `find_homology_transform` in `extcurves`. We will give examples of its use for reasonably substantial problems in later chapters, particularly those on Klein's and Bring's curves.

2.5.2   Maple's own homology

If we hope to interact with existing Maple calculations or even draw on its work, we must have some means of relating any basis we pick to the one it chooses. It is clear that, with `find_homology_transform` written, we need some method to convert Maple's homology basis to a list of `extpath`s.

Fortunately Maple's own function `monodromy` has an option which also returns the homology basis it used to derive that data. For example

```
> monodromy(y^2 - x^4 + 1, x, y, 'give paths');
  [..., [...], paths, r, rootof]
```

The first three entries in the list are the actual monodromy calculation and don't concern us here. The fourth entry is a list of (floating point approximations to) the branch points.

The fifth entry is the most interesting one for our uses. It is a table describing the paths Maple uses to get from its chosen base point to each of the branches. Together with information in the `homology` calculation (which gives the chosen homology basis in terms of how branches should be encircled) this is enough to convert any element of Maple's basis into an extpath.

The only possible problem is that Maple uses circular arcs to describe some sections of its cycles. Fortunately the circles are carefully chosen by Maple to just encircle one branch point and are only specified semicircles so in every case they can be replaced by two straight line segments (again as in Figure 2.1).

The result is the function `extpath_from_homology` which takes a curve and an index, returning the extpath equivalent of the desired Maple cycle. A convenience function `from_algcurves_homology` is also provided to act on the entire basis at once. Specific programming details can be found in the `extcurves` documentation in Appendix A.

This function is not explicitly used in later sections but forms the foundation of the next function.

2.5.3 `periodmatrix`

The package `algcurves` already provides a function `periodmatrix` for calculating the matrix of periods for a surface Σ

$$\Pi_{ij} = \int_{\gamma_i} \omega_j,$$

where $\{\gamma_i\} = \{\mathfrak{a}_1, \ldots, \mathfrak{a}_g, \mathfrak{b}_1, \ldots, \mathfrak{b}_g\}$ is a canonical homology basis and $\{\omega_1, \ldots, \omega_g\}$ is a basis for the space of holomorphic differentials on Σ.

The choice of a basis for the holomorphic differentials ω_i is largely unimportant since most uses for the matrix of periods are via the so-called Riemann period matrix which is obtained by *changing* the basis of differentials so that

$$\int_{\mathfrak{a}_i} \omega_j = \delta_{ij},$$

i.e. in block form

$$\Pi = \begin{pmatrix} \mathbf{1} \\ \tau \end{pmatrix}.$$

In particular it is τ that appears in expressions for the Θ-functions on Σ.

However, we would very much like to be able to control which homology basis is chosen, particularly if Σ has interesting symmetries. `Algcurves` itself does not provide this facility, but fortunately we don't have to reinvent the `periodmatrix` code to add it.

Suppose we have a basis $\{\eta_i\}$ we would prefer to use, and that it is related to the `algcurves` choice γ_i by

$$\eta_i = A_{ij}\gamma_j.$$

Then the new period matrix $\hat{\Pi}$ is given by

$$\begin{aligned}
\hat{\Pi}_{ij} &:= \int_{\eta_i} \omega_j \\
&= \int_{A_{ik}\gamma_k} \omega_j \\
&= A_{ik} \int_{\gamma_k} \omega_j \\
&= A_{ik}\Pi_{kj}.
\end{aligned}$$

Or, as matrices

$$\hat{\Pi} = A\Pi.$$

So, by putting together our own existing code to relate `algcurves`' basis to our preferred one and `periodmatrix` from `algcurves` we can easily compute the equivalent period matrix in our own basis.

`Extcurves` provides its own `periodmatrix` function, taking a curve and a list of `extpath`s representing a homology basis and does precisely this. In fact it does slightly more: the list does not have to form a basis and our version simply calculates the integrals along any closed path given which is occasionally helpful.

For example if `cycle1` and `cycle2` are homology cycles then

```
> periodmatrix(curve, [cycle1, cycle2]);
```

will calculate a matrix containing the integrals of each holomorphic differential (as specified by `algcurves/differentials`) along `cycle1` and `cycle2`. We will see particular examples of this in the following chapters.

2.5.4   `transform_extpath`

Finally, there is a much more fragile function for applying given transformations to an `extpath`. The problem is that by selecting the transformation carefully, a segment can be turned into an arbitrarily complex path on the surface. This would force us to deal with the very issues we had hoped to avoid by demanding paths be presented as segments. For example, Figure 2.11 arose in Klein's curve when we tried to calculate the image of a segment under an order 4 automorphism; notice that the straight segment became a path wrapping tightly around a branch point. Great care would be needed to accurately represent this situation by increasing the number of segments produced near this kink.

On the other hand many transformations, useful in practice, are actually very simple. Later we will have reason to apply automorphisms of surfaces which are linear in both components

$$(x, y) \mapsto (\alpha x, \beta y).$$

Such transformations should raise no significant issues because they send segments to segments. Any algorithm we choose to segmentise a generically curved path, no matter how unreliable in general, should perform well in this instance. Only slightly more

Figure 2.11: Difficult kinks arising from natural transformations

complicated, but still common, are transformations that act as Möbius functions on each coordinate (or sometimes just x).

The fundamental problem is taking the push-forwards of a segment under some transformation, and trying to split this curved path into an equivalent set of straight segments, twining around each branch point correctly. Additional goals are to stay away from branch points and to minimize the final number of segments produced (for computational efficiency on the result).

In the end we chose an algorithm that is rather unstable; however, provided it is supervised and used within the (simple) cases it handles well, it can save significant time. Even outside these bounds its output may need to be corrected but it has proven useful as a first approximation.

The algorithm first tries to approximate the final curve by a single segment. To test whether this is plausible, it calculates the winding number about each branch point of the (closed in $\mathbb{C}$) combined path. If it is zero then the paths probably take the same route around that branch, otherwise we need a better approximation.

The complex issue is deciding how to split the single segment if needed. We adopt the naïve approach of recursive bisection, which is nevertheless sufficient for many applications. In particular linear transformations are guaranteed to produce optimal results and Möbius transformations (under which a segment can only become a circular arc at worst) usually work.

Despite these limitations, this function will also see substantial use in dealing with Klein's and Bring's curves later when, in fact, many automorphisms will satisfy the stringent requirements, and it will prove useful enough to save time on others.

3

CYCLEPAINTER: VISUALLY REPRESENTING HOMOLOGIES

Figure 3.1: Traditional algebraic representation of paths on a surface

3.1 Motivation and visual style

The latest source code for the software described in this chapter can be downloaded from `http://gitorious.org/riemanncycles/CyclePainter`; a snapshot is on the included CD.

With all the restrictions we have placed on the format of homology cycles for use in `extcurves`, some kind of visual interface becomes almost essential.

The fundamental problem is that we have a surface

$$\Sigma \; : \; \{(x, y) \in \mathbb{C}^2 \mid f(x, y) = 0\}, \tag{3.1}$$

and a path described by two functions, say from the unit interval in $\mathbb{R}$

$$x, y \; : [0, 1] \to \mathbb{C}.$$

In the absence of four-dimensional displays this path cannot simply be plotted. The usual solution, shown for example in Figure 3.1 from [10] (with thanks to Antonella d'Avanzo for the use), is to simply plot the function $x(t)$ on the plane and exploit the fact that at any given x, there are only finitely many possibilities for y. Recall from Section 2.2.1 that by considering f, Σ is an n-sheeted cover of $\mathbb{C}$. The style of line then indicates which actual sheet the path is on. In Figure 3.1 for example, there are 3 sheets indicated by solid, dashed and dotted lines.

Figure 3.2: `CyclePainter` representation of paths

This idea can be applied algorithmically, although since many sheets may be needed (up to 7 in examples presented later), finding enough different styles of dashing proves impractical and we instead represent sheets by different colours.

With the choices made for `CyclePainter` and a fairly simple engine to output metapost, we can produce – almost automatically and from a basis created for functionality rather than looks – the images in Figure 3.2.

Clearly compromises have been made. In addition to those discussed above:

- Diagrams have to be split up more to remain clear as it is not really practical to stop cycles interfering with each other.

- Paths are less uniform.

However, the same information is represented and with little loss of clarity.

3.2 Representation of sheets

3.2.1 Colouring in surfaces

The biggest problem to be solved at this stage is an algorithmic assignment of colours.

Mathematically, Σ is given by (3.1) with a natural projection $\pi_x : \Sigma \to \mathbb{C}$, as in Section 2.2.1 given by $\pi_x(x, y) = x$. Then $\pi_x^{-1}(x)$ has at most $\deg(f)$ members (the number of sheets of this map). So if we are given a set of n colours, χ, our (unattainable) goal is to find a function

$$s : \Sigma \to \chi$$

with the following properties

- s restricted to $\pi_x^{-1}(x)$ is injective for all $x \in \mathbb{C}$. This condition says that each sheet above a given point will have a different colour, essential if we are to unambiguously display the sheet a path is on with colour.

- s should be as continuous as possible, in some sense. Complete continuity cannot be attained since χ is a finite set so s would be constant, but we do want to minimise the number of surprising changes of colour that occur in paths.

- Discontinuities of s should occur together in the affine projection as far as possible. That is, if s is discontinuous at (x, y) then if possible we should try to place discontinuities on other sheets at points of $\pi_x^{-1}(x)$. This condition will make the diagrams less cluttered by limiting the number of x coordinates where paths can change colour.

The first property is an absolute requirement, the subsequent ones are more guidelines, and indeed compromises will have to be made in the name of automation.

It is usual to actually only define the function s for most points on Σ. In particular branch points are often excluded because s would necessarily be discontinuous there and the actual sheet of any paths can be inferred by continuity since each branch point is isolated.

The normal resolution to finding the function s, at least on hyperelliptic curves, is via branch cuts as in Figure 3.3, representing a simple elliptic curve. Lines between various branch points are cut from $\pi_x(\Sigma)$ to give a disconnected set $\bar{\Sigma}$. Each component of $\bar{\Sigma}$ can then be given a different colour and visually paths will only change colour when they cross one of these cuts.

The problem with this approach from a computational point of view is that these branch cuts cannot be chosen arbitrarily; if we require each branch point to only be involved in one cut they may not even exist if there are more than two sheets. For example consider the curve

$$y^3 = x^4 - 1.$$

It has branch points at $\pm 1, \pm i$ (and ∞). If we made the obvious branch cut from 1 to -1 then we would implictly be asserting that a loop around these two in x was actually a closed path on the surface, which is false. In fact it is a simple exercise of `CyclePainter` to see that no branch cuts are valid for this curve: a closed path around any pair of branch points manifestly comes back a different colour.

Figure 3.3: Branch cuts on a simple elliptic curve $y^2 = x^4 - 1$

A slight generalisation to the above procedure, however, is appropriate for implementation. Instead of taking cuts directly between branch points, we designate a preferred central point $p_c \in \mathbb{C}$ and make a cut from p_c to each branch.

For the sake of simplicity in `CyclePainter` we demand

- that p_c is not a branch point.

- that a ray from p_c will reach each branch point without passing through another.

- that if ∞ is a branch point the left-pointing horizontal ray from p_c does not pass through any branch points on $\mathbb{C}$.

The result of these requirements is that straight line cuts radiating from p_c can be used unambiguously. Since there are only finitely many branch points such a p_c can always be found.

We will use this in the following form

Theorem 9. *Given a surface Σ which is an n-sheeted cover of $\mathbb{C}$, and a base point p_c as above, we can remove straight line cuts from p_c to each branch point,*

$$C = \bigcup_{b \in \mathcal{B}} \{p_c t + b(1 - t) : t \in [0, 1]\}.$$

where $\mathcal{B}$ is the set of all branch points. If ∞ is a branch then we take its ray to be from

Figure 3.4: How to colour a surface to represent sheets.

p_c to the left. Now define

$$\Sigma_c = \Sigma \setminus \pi_x^{-1}(C).$$

Σ_c has n path-connected components and π_x^{-1} restricted to $\mathbb{C} \setminus C$ maps each x to precisely n preimages, one from each component.

Proof. Trivial. See for example Hirsch [21]. $\square$

We can perhaps see more clearly what is going on if we refer to Figure 3.4. In reality, of course, the surface wouldn't be self-intersecting past the branch-points. But we can schematically see that a path looping around a branch point before returning to its initial (projected) position, may not be closed. Removing the cuts (where the colour changes discontinuously) is a strong enough change to disallow this possibility.

The computational effect of this result is that we only need to change the colour of a path when it *does* cross one of our cuts; further, the colour-change is fixed by continuity for the entire length of the cut. So, when initialising data-structures for a specified surface, `CyclePainter` will precompute how sheets change when a path crosses one of the cuts and then simply look this data up later on as paths are constructed. This significantly reduces the runtime cost that would be incurred by routine analytic continuation and makes realtime update of paths feasible. It is possible this technique could be adapted to `extcurves` as well to reduce runtime costs significantly, though with slight loss of flexibility (realistically, paths would have to cross cuts transversely and avoid p_c).

3.2.2   Precomputing the effect of crossing a cut

In addition to p_c we will fix a preferred sector (call it the primary sector) of $\mathbb{C} \setminus C$, on which we will initially define the colour map. Other sectors will be numbered in an anticlockwise manner. In practice this sector will be specified by a point $p_s \in \mathbb{C}$ lying

Figure 3.5: Special points on affine representation of a curve.

inside (as in the proof above). In summary, referring to Figure 3.5 (a representation of $y^2 = x^4 - 1$ where large dots are branches, solid lines are cuts) note:

$p_c = $ Centre from which straight cuts radiate to the branch points.

$p_s = $ Point at which initial assignment of colours is made.

A slight modification of Maple's builtin `monodromy` function gives us, for a specified p_c

- A preferred order $(y_1, \ldots, y_n)$ of the sheets $\pi_x^{-1}(p_c)$ at p_c.

- The analytic effect of circling any branch point once anticlockwise. Essentially, with the set of branch points $\mathcal{B}$ as before, it provides us with a function

$$m : \mathcal{B} \times \{y_1, \ldots, y_n\} \to \{y_1, \ldots, y_n\},$$
$$m_b(y_i) = \text{End point of loop from } (p_c, y_i) \text{ anticlockwise around } b.$$

We assign colours from the set χ to each connected component based on this preferred order and the primary sector. Points close to (p_c, y_i) in the sector of p_s will be assigned the ith colour, giving rise to a map onto the set of colours, χ,

$$c_1 : \{y_1, \ldots, y_n\} \to \chi,$$

which can be extended by continuity to the rest of Σ_c.

Thinking about two separate points, p_c and p_s is an unfortunate necessity because, although Maple provides a preferred sheet ordering at p_c this is not part of any of the connected components of Σ_c that we assign colours to. It should be clear that the actual

Figure 3.6: The analytic and colouring effect of crossing a cut

colours will depend on the choice of preferred sector: crossing a cut will by definition change colours even if points are analytically close.

Now consider sector 2, separated from the primary sector by a cut leading to the branch point b_1 as in Figure 3.6. The function m_{b_1} provided by `monodromy` is exactly what we need to find the function c_2 describing how colours are assigned near p_c in sector 2: a loop around b_1 starting near (p_c, y_i) in the primary sector and returning in sector 2 will arrive at the point $(p_c, m_{b_1}(y_i))$. At the beginning its mapping from analytic points to colours will be described approximately by c_1 and at the end (since it has crossed no cuts) by c_2. Thus

$$c_1(y_i) = c_2(m_{b_1}(y_i)),$$

or

$$c_2 = c_1 \circ m_{b_1}^{-1}.$$

From this we deduce that the effect of crossing b_1's cut in an anticlockwise manner is given by first finding out the analytic point corresponding to the sector-1 colour h,

then finding which sector-2 colour this corresponds to. Symbolically the colour-change l_1 induced by crossing the cut to b_1 is

$$l_1 : \chi \to \chi,$$
$$l_1(h) = c_2(c_1^{-1}(h))$$
$$= (c_1 \circ m_{b_1}^{-1} \circ c_1^{-1})(h).$$

Adopting the obvious numeric labelling of branches to correspond with sectors the generalisation

$$c_{i+1} = c_i \circ m_{b_i}^{-1},$$
$$l_i = c_{i+1} \circ c_i^{-1}$$
$$= c_i \circ m_{b_i}^{-1} \circ c_i^{-1},$$

is easily obtained for finite branch points. This system can be implemented as a recursion scheme as in Algorithm 2.

Algorithm 2 Precomputing effects of crossing cuts.

Require: Monodromy list *analyticMons* sorted by argument anticlockwise from primary sector.
Ensure: *cuts* is a list of the effect of crossing each cut.
    $sectorCol \leftarrow$ canonical colour map for primary sector.
    for all *analyticMon* in *analyticMons* **do**
      $cuts \leftarrow cuts, sectorCol \circ analyticMon \circ sectorCol^{-1}$.
      $sectorCol \leftarrow sectorCol \circ analyticMon^{-1}$.
    end for

Finally, a point at ∞ may be a branch point. l_∞ is simply chosen so that a small loop around p_c in $\mathbb{C}$ – which should certainly be closed whatever sheet it's on – does indeed come back to the same colour. Specifically we demand

$$\prod_{i=1}^{|\mathcal{B}|} l_{b_i} = \mathbf{1}.$$

3.3 User interface

The startup screen for `CyclePainter` is shown in Figure 3.7 with the key regions annotated. The numbering of each region corresponds to subsection numbers in the following discussion.
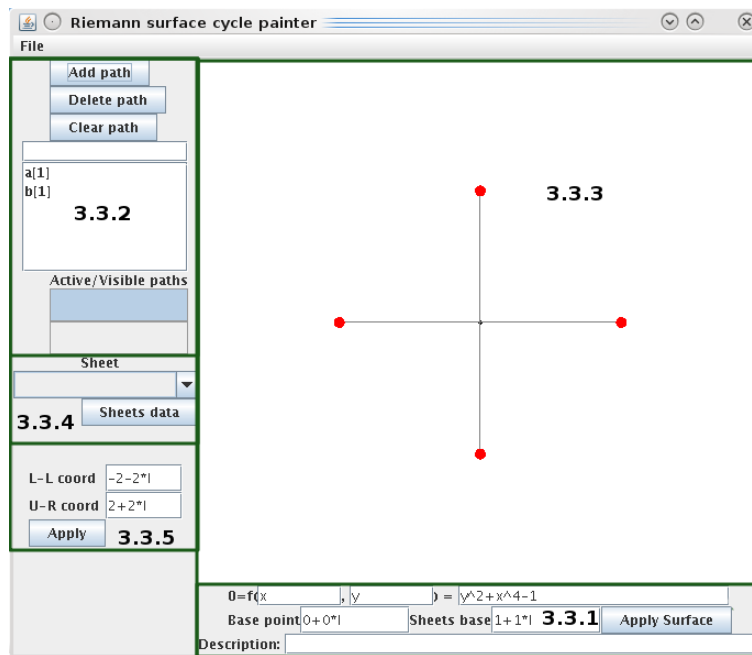
Figure 3.7: Normal `CyclePainter` user interface, with regions corresponding to subsections labelled in green.

3.3.1   Surface details

This section is used to specify precisely which surface we should be considering; a change here will be reflected in the display area (see 3.3.3) and will obviously have a profound effect on the rendering of paths.

The first line is self-explanatory, allowing the user to specify an affine representation of the Riemann surface. The second line specifies the two points necessary to determine sheets. "Base point" is p_c, the point from which our cuts will radiate. "Sheets base" is p_s, an arbitrary choice of where to assign initial colours. In principle `CyclePainter` could quite easily algorithmically derive a valid point for "Sheets base", however occasionally a user may want to move it to a more convenient location so the value is editable.

Finally, the "Description" line simply allows human-readable notes to be attached to a file.

It is usually a semantic error to change the surface but retain paths, however it can be valid if the change is small enough so `CyclePainter` does not attempt to second-guess the user and leaves the list of paths unchanged when a new surface is entered.

3.3.2   Paths configuration

This area serves two related purposes
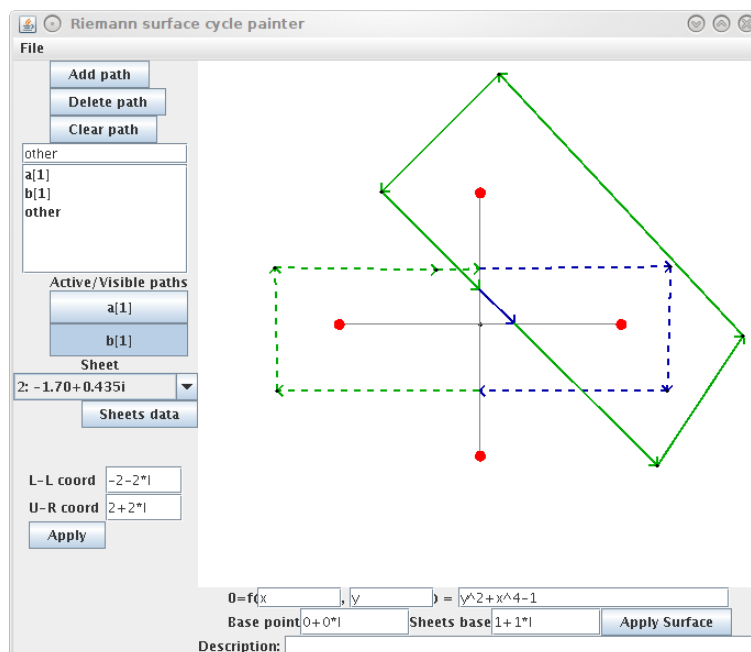
Figure 3.8: Demonstration of path selection UI

- It allows paths to be added to and removed from the file. Although a surface of genus g naturally has a symplectic homology basis $\mathfrak{a}_1, \ldots, \mathfrak{a}_g, \mathfrak{b}_1, \ldots, \mathfrak{b}_g$, forcing just these paths with this naming convention on a user would be counterproductive; there are frequent cases when we want to consider either a larger or more limited set of paths.

- It selects which paths should currently be displayed and editable in the viewport area.

Paths are created by first entering a (valid Maple) name, for example "other" in Figure 3.8 and then clicking "Add path". Deletion is by selecting a path in the list box and (unsurprisingly) clicking "Delete path".

The more interesting decision is how paths are selected for displaying. This is via drag and drop from the list box (containing `a[1]` and `b[1]` in Figure 3.7) to either of the two buttons below "Active/visible paths". The buttons themselves show which paths will be visible in the main display area (none in Figure 3.7 but `a[1]` and `b[1]` in Figure 3.8).

The highlighted button (`b[1]` in Figure 3.8) shows which path will be modified by any editing actions – the active path. That path is displayed as a solid rather than dashed line in the surface display (again visible in Figure 3.8).

3.3.3   Surface display

This is where most of the work is done editing paths. The key operations supported are

- A left-click either starts a new path or begins extending the existing path. Nodes added will snap to nearby existing nodes, so creating a closed path does not require pixel-perfect clicking.

- A right click will end the current drawing operation without placing a node under the cursor.

- When not drawing, nodes can be selected by clicking on them; the selected node is highlighted with a red box.

- A selected node can be moved by dragging.

- A selected node can be removed by pressing the "Delete" key.

This interface is rather unconventional, but not completely divorced from standard expectations. It was chosen simply for easy implementation since the primary concern of this package is mathematical.

3.3.4   Sheets configuration

The drop-down box selects the sheet (either in terms of number or y value) of the initial point of the currently active path.

The button "Sheets data" opens a second window which displays the correspondence between sheet colour, sheet number and y value at two points: the specified "Sheets base", and the selected node if it exists.

3.3.5   Viewport configuration

This section simply allows the displayed area of the complex plane to be set in terms of its lower-left and upper-right coordinates. It effectively acts as a zoom control.

3.3.6   File menu

Finally, as would be expected, the usual Save/Load abilities are here. Of slightly more interest is a "Write Metapost" option, which partially automates producing diagrams like those in Figure 3.2 by producing Metapost code which can draw each cycle.

3.4   Integration with `extcurves`

`CyclePainter`'s primary save and load functions (in the "File" menu) produce files that are valid Maple code and provide all the data needed to reconstruct the `extpaths` drawn. The conventional filename extension is `.pic`, and this is reflected in the names of functions interfacing with `extcurves`.

However, a couple of utility functions in `extcurves` make the interoperation much more painless.

`read_pic` takes a filename and returns a sequence with three elements

1. A curve record, as described in the previous section.

2. A list of `extpaths` defined by the `CyclePainter` file.

3. A list of strings, giving the name each `extpath` has in `CyclePainter`.

For example

```
> curve, homology, names := read_pic("somefile.pic"):
```

Essentially this gives all the information needed by other `extcurves` routines, in the format they are needed, and without polluting the global namespace by executing a `read` statement in Maple.

A partial inverse is also provided, in case the need to write `CyclePainter` files from Maple arises. In its simplest use, the function `write_pic` accepts a filename, curve and list of `extpath` s. For example

```
> write_pic("somefile.pic", curve, homology):
```

More options are allowed for finer control: see Appendix A for details.

4

EXAMPLE: KLEIN'S CURVE

4.1   Introduction

We will now apply the techniques and software developed above to Klein's curve. This is the genus 3 Riemann surface with maximal symmetry group so we should expect results to be correspondingly simple. First we'll review the basic facts about Klein's curve and its symmetry group. In fact one of the symmetries is surprisingly elusive and describing it with sufficient precision to be useful will occupy us for a while. After that it will be easy to find a suitable homology basis (a short calculation based on the tools created). However, proving that this homology basis actually gives the desired period matrix will be slightly more involved.

We then study Rauch and Lewittes' hyperbolic model of Klein's curve. They produced a period matrix in that setting in [29], and by constructing a reasonably precise correspondence between the hyperbolic and algebraic models we show that the two matrices are indeed equivalent. We give an explicit algebraic analogue of Rauch and Lewittes' basis and provide the symplectic transformation relating it to our symmetric basis.

Finally we calculate the vector of Riemann constants for this curve, an essential ingredient to any possible applications in integrable systems.

4.2   Algebraic representations of Klein's curve

Klein's quartic curve is expressed algebraically in projective space $\mathbb{CP}^3$ as

$$x^3 y + y^3 z + z^3 x = 0.$$

The affine projection can obviously be written as

$$x^3 y + y^3 + x = 0,$$

and has 9 branch-points in the x-plane. They are at 0, ∞ and the points of a regular septagon centred on 0 as in Figure 4.1. If we make the birational transformation

Figure 4.1: Finite branch points of Klein's curve in quartic representation.

mentioned in [22]

$$\phi : (x, y) \mapsto (t, s) = \left(1 + \frac{x^3}{y^2}, -\frac{x}{y}\right), \tag{4.1}$$

$$(t, s) \mapsto (x, y) = \left(\frac{t-1}{s^2}, \frac{1-t}{s^3}\right), \tag{4.2}$$

then the coordinates (t, s) satisfy a septic equation

$$s^7 = t(t-1)^2.$$

This representation has 3 branch points on the t plane, located at 0, 1 and ∞.

A final useful form of the curve is obtained by applying a fractional-linear transformation which simply sends these branch-points to the cube roots of unity. If we denote $\rho = \exp(2\pi i/3)$ then the desired transformation is

$$\tilde{\phi} : (t, s) \mapsto (z, w) = \left(\frac{t + \rho^2}{\rho t + \rho^2}, \frac{s(\rho^2 - 1)}{\rho t + \rho^2}\right), \tag{4.3}$$

$$(z, w) \mapsto (t, s) = \left(\frac{\rho^2(z-1)}{1 - \rho z}, \frac{w}{1 - \rho z}\right), \tag{4.4}$$

and the curve now obtained is

$$w^7 = (z-1)(z-\rho)^2(z-\rho^2)^4.$$

4.3   Holomorphic differentials on Klein's curve

We want to pick a basis of holomorphic differentials in each of the spaces so that disruption is minimised when moving between the pictures. To do this we define the basis in the quartic form of the curve and pull back twice to obtain corresponding bases in the two septic cases. Fortunately the pullback is rather simple and the only difference from the natural choice in each case is a scalar multiple, as we'll see.

In the quartic case we pick the following (ordered) basis

$$\{\omega_1, \omega_2, \omega_3\} = \left\{ \frac{x\mathrm{d}x}{x^3 + 3y^2}, \frac{y\mathrm{d}x}{x^3 + 3y^2}, \frac{\mathrm{d}x}{x^3 + 3y^2} \right\}. \tag{4.5}$$

The actual order is unimportant for most of the work, but that chosen will make the action of the order 2 and 4 automorphism on the differentials simpler later.

It is equivalent to the following simple basis in the (t, s) version

$$(\phi^{-1})^*\omega_1 = \frac{1}{7}\frac{(t-1)\mathrm{d}t}{s^5},$$
$$(\phi^{-1})^*\omega_2 = \frac{1}{7}\frac{(t-1)\mathrm{d}t}{s^6},$$
$$(\phi^{-1})^*\omega_3 = \frac{1}{7}\frac{\mathrm{d}t}{s^3}.$$

Finally, moving to the rotationally symmetric (z, w) plane each differential has a *different* scale-factor from the natural choice, but is still just a multiple

$$(\tilde{\phi}^{-1} \circ \phi^{-1})^*\omega_1 = \frac{\rho - 1}{7}\frac{(z-\rho)(z-\rho^2)^2\mathrm{d}z}{w^5},$$
$$(\tilde{\phi}^{-1} \circ \phi^{-1})^*\omega_2 = \frac{-1 - 2\rho}{7}\frac{(z-\rho)(z-\rho^2)^3\mathrm{d}z}{w^6},$$
$$(\tilde{\phi}^{-1} \circ \phi^{-1})^*\omega_3 = \frac{2 + \rho}{7}\frac{\mathrm{d}z}{w^3}.$$

4.4   Symmetries

It is well-known that Klein's curve possesses a holomorphic symmetry group of order 168. Making use of this will allow us determine the period matrix almost completely, but first we must study many of the symmetries in enough depth to determine their action on both the space of holomorphic differentials and (when we have chosen one) on the homology basis.

The following is presented in roughly increasing order of complexity. The final

Representation		Action of symmetry	
0	$=$	$x^3y + y^3z + z^3x$	$[x,y,z] \mapsto [y,z,x]$
0	$=$	$x^3y + y^3 + x$	$(x,y) \mapsto \left(\frac{y}{x}, \frac{1}{x}\right)$
s^7	$=$	$t(t-1)^2$	$(t,s) \mapsto \left(1 - \frac{1}{t}, \frac{t-1}{s^3}\right)$
w^7	$=$	$(z-1)(z-\rho)^2(z-\rho^2)^2$	$(z,w) \mapsto \left(\rho^2 z, \frac{-\rho^{\frac{5}{3}}(z-1)(z-\rho)(z-\rho^2)^2}{w^3}\right)$

Table 4.1: Order 3 automorphism

automorphism is not actually used to determine the period matrix, so it is given purely for completeness and in less detail than the others.

4.4.1 Antiholomorphic involution

Since the curve is real it possesses an antiholomorphic involution given in $[x,y,z]$ coordinates simply by $[x,y,z] \mapsto [\bar{x}, \bar{y}, \bar{z}]$.

The only coordinates where this automorphism has nontrivial action in are (z,w). In this case the presence of the complex number ρ complicates matters and the relevant expression is

$$(z,w) \mapsto \left(\frac{1}{\bar{z}}, -\rho\frac{\bar{w}}{\bar{z}}\right).$$

The symmetry simply acts as complex conjugation on the differentials, for example

$$\frac{\mathrm{d}x}{x^3 + 3y^2} \mapsto \frac{\mathrm{d}\bar{x}}{\bar{x}^3 + 3\bar{y}^2}.$$

4.4.2 Order 3 cyclic automorphism

This symmetry is easiest to see in the projective quartic coordinates. Permuting x,y,z cyclically obviously fixes the curve. Choose $[x,y,z] \mapsto [y,z,x]$ as a favoured permutation, and the transformations in each coordinate system are given by Table 4.1.

The effect of the pullback on the holomorphic differentials is as follows

$$\omega_1 \mapsto \omega_2,$$

$$\omega_2 \mapsto \omega_3,$$

$$\omega_3 \mapsto \omega_1.$$

As will be the case for subsequent automorphisms, the action is by definition identical on the corresponding differentials in the other representations.

4.4.3   Order 7 automorphism

There is an obvious order 7 automorphism in (z, w) and (t, s) spaces; namely if we define $\zeta = \exp(2\pi i/7)$ then

$$(z, w) \mapsto (z, \zeta w)$$

is a symmetry. The same formula holds in (t, s), however in the (x, y) and $[x, y, z]$ quartic coordinates this automorphism is

$$(x, y) \mapsto (\zeta^5 x, \zeta^4 y),$$

$$[x, y, z] \mapsto [\zeta^5 x, \zeta^4 y, z].$$

The pullback action on differentials is

$$\omega_1 \mapsto \zeta^2 \omega_1,$$

$$\omega_2 \mapsto \zeta^1 \omega_2,$$

$$\omega_3 \mapsto \zeta^4 \omega_3.$$

4.4.4   Holomorphic involution

This symmetry is really a square of the following order 4 automorphism presented by Egan at [14], but since it is far simpler to express and is more directly useful in calculating the period matrix, it has its own section. The techniques used to describe it are also just simplified versions of those used by Egan for the order 4 automorphism.

The involution is most easily expressed in the projective $[x, y, z]$ coordinates since it is effectively a rotation there. The idea is that if we look at the set of real solutions in the correct coordinates it has an obvious rotational symmetry of order 2; this extends to the space of complex solutions and yields the desired involution.

So we proceed in stages from the original $x^3y + y^3z + z^3x = 0$ form. Since this equation is homogeneous the set of real solutions forms a cone ((x, y, z) a solution implies (kx, ky, kz) a solution). The first step is a rotation so that the z axis aligns with the direction $(1, 1, 1)$ in the original coordinates – the "centre" of the cone. It

Figure 4.2: Section of real cone at $z' = 1$ (so after rotation).

corresponds to a rotation of $\tan^{-1} \sqrt{2}$ about $(1, -1, 0)$. Explicitly

$$R_1 = \frac{2}{3} \begin{pmatrix} \sqrt{3} + 3 & \sqrt{3} - 3 & -2\sqrt{3} \\ \sqrt{3} - 3 & \sqrt{3} + 3 & -2\sqrt{3} \\ 2\sqrt{3} & 2\sqrt{3} & 2\sqrt{3} \end{pmatrix}$$

is the rotation mapping $(1, 1, 1)$ to $(0, 0, \sqrt{3})$ and takes us to coordinates (x', y', z').

In these coordinates, the curve can be written rather simply in cylindrical polars, (r, θ, z') as

$$0 = 4z'^4 + 6r^2 z'^2 - 3r^4 - 2\sqrt{14} r^3 z' \cos[3(\theta - \theta_0)],$$

where

$$\theta_0 = \frac{\pi}{4} - \frac{1}{3} \tan^{-1}(3\sqrt{3}).$$

This has an obvious threefold symmetry on rotating about z'. See Figure 4.2 for a cross-section of the rays making up the cone. The involution is a rotation of the cone through π radians about an axis through the origin, perpendicular to z' and with angle $\theta_0 + \frac{\pi}{2}$ from x' (indicated by the dotted line in Figure 4.2). In words, we conjugate a rotation of π radians $(a, b, c) \mapsto (-a, b, -c)$ with a rotation of θ_0 about the z' axis which lines up the symmetry of the cone. In (x', y', z') coordinates the involution has matrix form

$$\begin{pmatrix} -\sin\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) & -\cos\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) & 0 \\ -\cos\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) & \sin\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) & 0 \\ 0 & 0 & -1 \end{pmatrix}.$$

It is now a simple, if detailed, matter to express this involution in the original

(x, y, z) coordinates. The final expression is

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} \alpha & \beta & \gamma \\ \beta & \gamma & \alpha \\ \gamma & \alpha & \beta \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}, \tag{4.6}$$

where

$$3\alpha = \cos\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) - \sqrt{3}\sin\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) - 1,$$

$$3\beta = -2\cos\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) - 1, \tag{4.7}$$

$$3\gamma = \cos\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) + \sqrt{3}\sin\left(\frac{2}{3}\tan^{-1}(3\sqrt{3})\right) - 1.$$

At this stage we could calculate what this involution is on either (t, s) or (z, w) space. However the expression gets progressively more complicated (the transformation is not a projectivity in either of the other systems) and is not very illuminating.

It is, however, essential to calculate the effect on the differentials. This is made much simpler by the fact that the symmetry is expressed as a linear function in projective space. We know that

$$\omega_3 \mapsto A\omega_1 + B\omega_2 + C\omega_3$$
$$= (Ax + By + C)\omega_3$$

for some fixed numbers A, B, C. We also know that for any automorphism ϕ,

$$\phi^*(f(x, y)\omega) = \phi^*(f)(x, y)\phi^*(\omega),$$

which leads to

$$\omega_1 = x\omega_3 \mapsto \frac{\alpha x + \beta y + \gamma}{\gamma x + \alpha y + \beta}(Ax + By + C)\omega_3,$$

$$\omega_2 = y\omega_3 \mapsto \frac{\beta x + \gamma y + \alpha}{\gamma x + \alpha y + \beta}(Ax + By + C)\omega_3.$$

Together these force $A = k\gamma, B = k\alpha, C = k\beta$ for some constant k. This would be trivial if x and y were independent and we could work in the polynomial ring $\mathbb{C}[x, y]$: we would know that, since ω_1 maps to a holomorphic differential, there are some complex

numbers D, E, F such that

$$\frac{\alpha x + \beta y + \gamma}{\gamma x + \alpha y + \beta}(Ax + By + C) = (Dx + Ey + F)$$

everywhere. Equating coefficients would immetiately give the result. Fortunately the same holds true here, though with a slightly more involved proof. Quite generally:

Lemma 10. *Let Greek letters denote fixed nonzero complex numbers. Suppose*

$$(Ax + By + C)(\alpha x + \beta y + \gamma) = (Dx + Ey + F)(\delta x + \epsilon y + \zeta) \tag{4.8}$$

for all (x, y) on Klein's curve. Then there exists $k \in \mathbb{C}$ such that $A = k\delta, B = k\epsilon, C = k\zeta, D = k\alpha, E = k\beta$ and $F = k\gamma$.

Proof. (With thanks to Leo Butler for this much-simplified version). Equation (4.8) can be considered the assertion that the determinant of a matrix of linear forms,

$$M = \begin{pmatrix} a_1 & a_2 \\ b_1 & b_2 \end{pmatrix},$$

where $a_1 = Ax + By + C$, $a_2 = Dx + Ey + F$, $b_1 = \alpha x + \beta y + \gamma$ and $b_2 = \alpha x + \beta y + \gamma$ vanishes.

In this view, one of the suppositions of the Lemma is that $\det M|_\Sigma = 0$. Now, if $\det M$ is nontrivial, it defines a quadric which must therefore contain Σ, absurd since Σ is given by an irreducible equation of degree 4.

Alternatively $\det M$ vanishes identically. This puts us back in the situation mentioned before this Lemma, and the result is trivial. $\square$

Lemma 10 tells us that an automorphism which is linear in projective coordinates has essentially the same (linear) action on differentials as those coordinates. All that remains is to determine the scale factor k.

Because this is an involution, squaring the transformation must yield the identity which tells us $k^2 = 1$. So $k = \pm 1$, and we have to determine the sign. We can do this with just one term in the expansion of the holomorphic differentials near 0. We switch to y as a coordinate in that neighbourhood and restate ω_3 as

$$\omega_3 = -\frac{\mathrm{d}y}{1 + 3x^2 y}$$

Then it is easy to see that to first order the involution maps (via pullback)

$$y \mapsto \frac{\alpha}{\beta} + y\left(\frac{\gamma\beta - \alpha^2}{\beta^2}\right) + \cdots,$$

$$x \mapsto \frac{\gamma}{\beta} + \cdots,$$

$$\mathrm{d}y \mapsto \left(\frac{\gamma\beta - \alpha^2}{\beta^2}\right)\mathrm{d}y(1 + \cdots).$$

Thus

$$\omega_3 \mapsto -\frac{\gamma\beta^2 - \alpha^2\beta}{\beta^3 + 3\alpha\gamma^2}\mathrm{d}y(1 + \cdots)$$

$$= -\beta\mathrm{d}y(1 + \cdots)$$

$$= \beta\omega_3 + \cdots$$

and since neither ω_1 or ω_2 have a constant term at $y = 0$ this suffices to determine $k = 1$. So in full

$$\omega_1 \mapsto \alpha\omega_1 + \beta\omega_2 + \gamma\omega_3,$$

$$\omega_2 \mapsto \beta\omega_1 + \gamma\omega_2 + \alpha\omega_3,$$

$$\omega_3 \mapsto \gamma\omega_1 + \alpha\omega_2 + \beta\omega_3.$$

4.4.5   Order 4 automorphism

In principle this is not much more difficult than the involution above. After the final coordinate-changing rotation, instead of a simple rotation by π we perform the slightly more complicated transformation

$$(a, b, c) \mapsto (\mathrm{i}a, b, -\mathrm{i}c).$$

This produces many more terms in the resulting matrix, but essentially it's still a linear transformation on projective space (in fact its square is easily seen to be the involution above).

The argument above for the differentials carries through here as well; however, since it is not actually necessary for the period matrix we won't go into more detail.

Branch point z	Effect on sheet k
0	$k \mapsto k+1 \pmod 7$
ρ	$k \mapsto k+2 \pmod 7$
ρ^2	$k \mapsto k+4 \pmod 7$

Table 4.2: Monodromy from base point 0 around each branch in a positive direction

4.5   Canonical homology basis

We try to construct a homology basis so that its transformation under most of these symmetries is as simple as possible. The curve has genus 3, and so we are looking for three $\mathfrak{a}$-cycles and $\mathfrak{b}$-cycles. Thus a natural choice is to specify $\mathfrak{a}_1$ and $\mathfrak{b}_1$ and then define the others as images of the order 3 symmetry.

This is most easily accomplished in the (z, w) space, where there are only three branch-points to choose from in constructing our paths and the automorphism is a simple rotation in the z-plane.

To have even a hope of constructing such a path, we need to know about the monodromy around the three branch points. We label the sheets by the w value at $z = 0$, i.e. for $k = 0, \ldots, 6$ "sheet k" is $(z, w) = (0, \exp(\frac{\pi i}{21}(6k - 1)))$. The monodromy about this base point is (for each branch point) a simple constant shift of sheets so defined. The exact change is given in Table 4.2.

With this in mind, the obvious place to start is some number of circuits around two branches. So we define $\mathfrak{a}_1$ to start at sheet 0 and proceed clockwise 3 times around $z = 1$ (taking us to sheet -3) and then once clockwise around $z = \rho^2$ (taking us back to sheet $-7 = 0$). If we then define the remaining $\mathfrak{a}$-cycles by applying the order 3 automorphism in Table 4.1 we discover they do not intersect each other and so can potentially form part of a canonical basis.

We now try to find appropriate $\mathfrak{b}$-cycles. One of the simplest approaches is to shift the sheet of the $\mathfrak{a}$-cycles. Fortunately, this works: if we simply start $\mathfrak{b}_1$ on sheet 2 instead of 0 and then derive the others by the order 3 symmetry again then the basis obtained is canonical.

See Figure 4.3 for illustration, where the top row gives the $\mathfrak{a}$-cycles, and the bottom row gives essentially the same paths on different sheets for $\mathfrak{b}$-cycles.

●Sheet 1●Sheet 2●Sheet 3●Sheet 4●Sheet 5●Sheet 6●Sheet ′

Figure 4.3: Homology basis in (z, w) coordinates. Along the top are $\mathfrak{a}_1, \mathfrak{a}_2, \mathfrak{a}_3$ from left to right, and $\mathfrak{b}_i$ on the bottom.

4.5.1   Propagating homology basis to other coordinates

Now that we have a functional homology basis in the (z, w) coordinates, we should find out what it corresponds to in other spaces.

In principle this could be done automatically by `extcurves`, but the resulting paths tend to have more nodes than are strictly necessary with corresponding penalties in performance for future routines. This would be compounded when we calculate the effects of the other automorphisms and so we spend some time here making sure the results are as simple as possible.

Accordingly the general pattern will be

- Computationally push each path forward into the next coordinate system. This uses a fairly simple script wrapping `transform_extpath`.

- Inspect the output to remove any extraneous points added, ensure the path doesn't go too close to any branch-points and generally create a neater picture of the path.

Naturally the result is simple in (t, s) coordinates, since the transformation between z and t is simply Möbius. See Figure 4.4.

In the (x, y) coordinates it is more complicated. See Figure 4.5.

●Sheet 1 ●Sheet 2 ●Sheet 3 ●Sheet 4 ●Sheet 5 ●Sheet 6 ●Sheet 7

Figure 4.4: Homology in (t, s) coordinates. Along the top are $\mathfrak{a}_1, \mathfrak{a}_2, \mathfrak{a}_3$ from left to right, and $\mathfrak{b}_i$ on the bottom.

●Sheet 1 ●Sheet 2 ●Sheet 3

Figure 4.5: Homology in (x, y) coordinates. Along the top are $\mathfrak{a}_1, \mathfrak{a}_2, \mathfrak{a}_3$ from left to right, and $\mathfrak{b}_i$ on the bottom.

4.5.2   Action of symmetries on homology basis

Having defined our homology basis and the symmetries in each of the pictures, it is a simple matter to apply one to the other and write down the action of each automorphism on the homology.

The action is most simply given as a matrix, so that if $\{\gamma_i\} = \{\mathfrak{a}_1, \mathfrak{a}_2, \mathfrak{a}_3, \mathfrak{b}_1, \mathfrak{b}_2, \mathfrak{b}_3\}$ is a canonical homology basis and ϕ and automorphism then

$$\phi_* : \gamma_i \mapsto M_{ij}\gamma_j.$$

Most of the symmetries have a simple action in at least one of the coordinate systems. For these finding the matrix is a few simple Maple commands. The involution is more complicated. Even in the simplest (x, y) coordinates the transformation is slightly too complicated for automatic processing.

The results are

- Antiholomorphic involution. Easiest to calculate in the t-plane where it fixes all branch-points and does not produce singularities from finite values. The command sequence is fairly natural (`conjugate` and `map2` are built-in Maple functions)

```
> curve, hom, names := read_pic("TS_tidy.pic"):
> invol := (t,s) -> [conjugate(t), conjugate(s)];

                                  ---    ---
                  invol := (t, s) -> [(t), (s)]


> new_hom := map2(transform_extpath, curve, hom, invol):
> find_homology_transform(curve, hom, new_hom);
```

and the last line produces the matrix

$$M = \begin{pmatrix} 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 \end{pmatrix}.$$

- Order 3 automorphism. No computation is needed here. By definition of the homology this cyclically permutes the $\mathfrak{a}_i$ and $\mathfrak{b}_i$ independently.

$$M = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}.$$

- Order 7 automorphism. This acts simply on paths in either septic space, just shifting sheets. A very similar procedure to the antiholomorphic involution produces the fairly complicated result

$$M = \begin{pmatrix} 1 & 0 & -1 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ -1 & 0 & -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 0 & 1 & 0 \\ 0 & -1 & -1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}. \tag{4.9}$$

However, since we know that $\mathfrak{b}_i$ *is* simply $\mathfrak{a}_i$ shifted by some sheet number, we know that some power of this automorphism will map $\mathfrak{a}_i$ to $\mathfrak{b}_i$. Indeed M itself takes $\mathfrak{a}_2$ to $\mathfrak{b}_2$; M^2 takes $\mathfrak{a}_1$ to $\mathfrak{b}_1$; and M^4 takes $\mathfrak{a}_3$ to $\mathfrak{b}_3$.

- Order 2 involution. Computing the matrix in this case is rather more complicated than before. `transform_extpath` cannot cope with the track directly because it sends the original straight lines to complex paths which are very close to certain branch points. We use the same approach as transforming the homology between the coordinates: create an approximate path automatically and manually tidy the output by inspecting the (curved) image of each straight segment in turn.

After that the transformation matrix can be calculated just as easily as before,

and the following surprisingly simple result is obtained

$$
M = \begin{pmatrix}
0 & 0 & -1 & 0 & 0 & 0 \\
0 & -1 & 0 & 0 & 0 & 0 \\
-1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & -1 \\
0 & 0 & 0 & 0 & -1 & 0 \\
0 & 0 & 0 & -1 & 0 & 0
\end{pmatrix}.
$$

4.6   Period matrix

Many of the symmetries of the surface lead directly to simplifications of its period matrix via the following basic result of integration. Suppose we have a holomorphic function of surfaces $\sigma : \Sigma \to \Sigma$; then

$$
\int_{\sigma_*(\gamma)} \omega = \int_\gamma \sigma^*(\omega).
$$

This is a simple restatement of how to change variables in an integration, but when σ is an automorphism these integrals are related to entries in the matrix of periods.

Since integration of this kind is linear both in the integrand and the path followed, this equation can be rewritten using the matrices M and L that give the action of an automorphism on homology cycles and differentials respectively.

Suppose we have a fixed automorphism σ acting via pullback on the holomorphic differentials $\{\omega_i\}$

$$
\sigma^*(\omega_i) = \omega_j L_{ji},
$$

and via pushforward on a canonical basis $\{\gamma_i\} = \{\mathfrak{a}_1, \ldots, \mathfrak{a}_g, \mathfrak{b}_1, \ldots, \mathfrak{b}_g\}$

$$
\sigma_*(\gamma_i) = M_{ij}\gamma_j.
$$

The matrix of periods is then

$$
\Pi = \begin{pmatrix} \mathcal{A} \\ \mathcal{B} \end{pmatrix},
$$

where

$$\mathcal{A}_{ij} = \int_{\mathfrak{a}_i} \omega_j, \qquad\qquad \mathcal{B}_{ij} = \int_{\mathfrak{b}_i} \omega_j.$$

And we can calculate that

$$(M\Pi)_{ik} = M_{ij} \int_{\gamma_j} \omega_k = \int_{M_{ij}\gamma_j} \omega_k$$

$$= \int_{\gamma_i} \omega_j L_{jk}$$

$$= \left(\int_{\gamma_i} \omega_j \right) L_{jk} = (\Pi L)_{ik}.$$

Or

$$M \begin{pmatrix} \mathcal{A} \\ \mathcal{B} \end{pmatrix} = \begin{pmatrix} \mathcal{A} \\ \mathcal{B} \end{pmatrix} L, \tag{4.10}$$

which in our case (with L and M known) is set of equations constraining the matrix of periods.

4.6.1   Order 3 symmetry

This symmetry constrains the $\mathfrak{a}$-periods and $\mathfrak{b}$-periods separately, but in an identical manner. We'll omit further discussion of the $\mathfrak{b}$-periods here since later automorphisms will firmly fix them in terms of the $\mathfrak{a}$s.

We have previously calculated the action on differentials and our homology basis; using those results the matrix relation (4.10) in this case becomes

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{pmatrix} \mathcal{A} = \mathcal{A} \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix},$$

or

$$\begin{pmatrix} \mathcal{A}_{21} & \mathcal{A}_{22} & \mathcal{A}_{23} \\ \mathcal{A}_{31} & \mathcal{A}_{32} & \mathcal{A}_{33} \\ \mathcal{A}_{11} & \mathcal{A}_{12} & \mathcal{A}_{13} \end{pmatrix} = \begin{pmatrix} \mathcal{A}_{12} & \mathcal{A}_{13} & \mathcal{A}_{11} \\ \mathcal{A}_{22} & \mathcal{A}_{23} & \mathcal{A}_{21} \\ \mathcal{A}_{32} & \mathcal{A}_{33} & \mathcal{A}_{31} \end{pmatrix}.$$

This means we can rewrite all $\mathcal{A}_{ij}$ in terms of the three values $\mathcal{A}_{1i}$, or even more simply

just X, Y, Z

$$\mathcal{A} = \begin{pmatrix} X & Y & Z \\ Y & Z & X \\ Z & X & Y \end{pmatrix}.$$

4.6.2   Antiholomorphic involution

This determines the $\mathfrak{b}$-periods in terms of the $\mathfrak{a}$-periods. Essentially the same argument holds, although the fact that the symmetry is not holomorphic means its effect on differentials will not be given by a complex matrix multiplication; we will perform the calculation explicitly. Calling the symmetry σ, we have

$$\mathcal{B}_{ij} = \int_{\mathfrak{b}_i} \omega_j = -\int_{\sigma(\mathfrak{a}_i)} \omega_j = -\int_{\mathfrak{a}_i} \sigma^*(\omega_j) = -\int_{\mathfrak{a}_i} \overline{\omega}_j = -\overline{\int_{\mathfrak{a}_i} \omega_j} = -\overline{\mathcal{A}}_{ij}.$$

4.6.3   Order 7 automorphism

This symmetry sends $\mathfrak{a}_2$ to $\mathfrak{b}_2$, so it tells us

$$\mathcal{B}_{21} = \zeta^2 \mathcal{A}_{21}, \qquad\qquad \mathcal{B}_{22} = \zeta \mathcal{A}_{22}, \qquad\qquad \mathcal{B}_{23} = \zeta^4 \mathcal{A}_{23}.$$

Or, in terms of X, Y, Z

$$-\overline{Y} = \zeta^2 Y, \qquad\qquad -\overline{Z} = \zeta Z, \qquad\qquad -\overline{X} = \zeta^4 X.$$

This allows us to write all three numbers in essentially polar form. For example suppose $X = r_1 \exp(i\theta)$. Then

$$-r_1 e^{-i\theta} = \left(e^{2\pi i/7} \right)^4 r_1 e^{i\theta},$$

or

$$\theta = -\frac{\pi}{14} + \pi k.$$

The integer k is irrelevant here, just corresponding to r_1 being positive or negative in the end. The solutions given below happen to result in $r_i > 0$ but that's purely an aesthetic post-hoc choice. We obtain

$$X = r_1 \exp(-\pi i/14), \qquad Y = r_2 \exp(-11\pi i/14), \qquad Z = r_3 \exp(-9\pi i/14). \qquad (4.11)$$

4.6.4   Holomorphic involution

As a first step (4.10) gives the relation (in terms of the trigonometric expressions (4.7) from Section 4.4.4)

$$
-\begin{pmatrix} Z & X & Y \\ Y & Z & X \\ X & Y & Z \end{pmatrix} = \begin{pmatrix} X & Y & Z \\ Y & Z & X \\ Z & X & Y \end{pmatrix} \begin{pmatrix} \alpha & \beta & \gamma \\ \beta & \gamma & \alpha \\ \gamma & \alpha & \beta \end{pmatrix} \tag{4.12}
$$

which naively corresponds to three independent equations. However the α, β, γ are not independent and

$$
\alpha\gamma = \beta(\beta + 1),
$$
$$
\beta^2 = (\alpha + 1)(\gamma + 1),
$$

together with permutations of those relations. So if we (say) multiply the first equation in (4.12)

$$
\alpha X + \beta Y + \gamma Z = -Z
$$

or equivalently

$$
\alpha X + \beta Y + (\gamma + 1)Z = 0
$$

by $(\beta + 1)/\alpha$ we in fact get the second equation

$$
(\beta + 1)X + \gamma Y + \alpha Z = 0.
$$

Similarly the remaining equation is actually equivalent to both of these. Thus there is only one independent (complex) equation in (4.12).

However, as we only have three real parameters r_1, r_2, r_3 left, this is enough to determine (say) r_1 and r_2 in terms of r_3 and hence the matrix of periods up to an overall real multiple.

This in turn *completely* fixes the Riemann form of the period matrix, and we'll see in the next section that even the scalar parameter is explicitly calculable.

Take the first entry of (4.12) as our starting point;

$$\alpha X + \beta Y + \gamma Z = -Z.$$

This can be solved algorithmically after substituting (4.7) and (4.11), and greatly simplified. Trigonometric simplification is more of an art than an algorithm. But we know that each of α, β, γ are algebraic numbers and simplification over algebraic extensions of $\mathbb{Q}$ is just a matter of time rather than insight; a computer can deal with the problem with no difficulty. Mathematica in particular includes a function `RootReduce` which takes an arbitrary algebraic expression of algebraic numbers (specified by minimal polynomial and root number) and produces its minimal polynomial. The application of this function to the present case is given as a notebook on the associated CD.

It tells us that

$$r_1 = \mu r_3, \qquad\qquad\qquad r_2 = \nu r_3;$$

where

$$\mu^3 + \mu^2 - 2\mu - 1 = 0, \qquad\qquad \nu^3 - 2\nu^2 - \nu + 1 = 0.$$

Each of these have three real roots, and split over $\mathbb{Q}(\zeta)$. The solution we desire is seen to be

$$\mu = \zeta + \zeta^{-1}, \qquad\qquad \nu = 1 + \zeta + \zeta^{-1}.$$

4.6.5   Final free parameter

The final free parameter can be computed directly from the integral in the (t, s) space. The cycle $\mathfrak{a}_3$ can be deformed to a path traversing $t = 0$ to $t = 1$ and then back again on a different sheet. Along this segment there is a sheet where s remains real, and all other

sheets are simply related by multiplication with some power of ζ. Thus for example

$$
\begin{aligned}
Y &= \int_{\mathfrak{a}_3} \omega_3 \\
&= \frac{1}{7} \int_{\mathfrak{a}_3} \frac{\mathrm{d}t}{s^3} \\
&= \frac{\zeta^k}{7} \int_0^1 \frac{\mathrm{d}t}{t^{3/7}(t-1)^{6/7}} + \frac{\zeta^l}{7} \int_1^0 \frac{\mathrm{d}t}{t^{3/7}(t-1)^{6/7}} \\
&= \frac{B(4/7,1/7)}{7}(\zeta^k - \zeta^l).
\end{aligned}
$$

All that remains is to determine the integers k and l. This is a trivial matter of analytically continuing the path $\mathfrak{a}_3$ over to somewhere on the segment $(0,1)$ and noting the argument of the resulting s. We discover

$$
Y = \frac{B(4/7,1/7)}{7}(\zeta^4 - \zeta).
$$

Using this to express r_2 in terms of Γ-functions rather than B-functions

$$
r_2 = \sqrt{2 + 2\cos(\pi/7)} \frac{\Gamma(8/7)\Gamma(4/7)}{\Gamma(5/7)}.
$$

4.7 Riemann period matrix

The $\mathfrak{a}$-normalised Riemann period matrix is

$$
\tau = \mathcal{B}\mathcal{A}^{-1}.
$$

Because $\mathcal{A}$ and $\mathcal{B}$ have cyclic symmetry, we can express them as

$$
\mathcal{A} = PWDW^{-1},
$$
$$
\mathcal{B} = PWD'W^{-1},
$$

where

$$
P = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \qquad\qquad W = \begin{pmatrix} 1 & 1 & 1 \\ 1 & \rho & \rho^2 \\ 1 & \rho^2 & \rho \end{pmatrix}.
$$

The diagonal matrices D, D' are the eigenvalues of $P\mathcal{A}$ and $P\mathcal{B}$ respectively, thus

$$D_{11} = X + Y + Z,$$

$$D_{22} = X + \rho Y + \rho^2 Z,$$

$$D_{33} = X + \rho^2 Y + \rho Z,$$

$$D'_{11} = -\bar{X} - \bar{Y} - \bar{Z},$$

$$D'_{22} = -\bar{X} - \rho\bar{Y} - \rho^2\bar{Z},$$

$$D'_{33} = -\bar{X} - \rho^2\bar{Y} - \rho\bar{Z}.$$

So we may write

$$\tau = \mathcal{B}\mathcal{A}^{-1}$$

$$= PWD'W^{-1}WD^{-1}W^{-1}P$$

$$= PWD'D^{-1}W^{-1}P$$

$$= PWD''W^{-1}P.$$

And we see that τ has the structure

$$\tau = \begin{pmatrix} \tau_1 & \tau_2 & \tau_3 \\ \tau_3 & \tau_1 & \tau_2 \\ \tau_2 & \tau_3 & \tau_1 \end{pmatrix}.$$

We can calculate the entries of D'' in terms of

$$X = r_3 \exp(-\pi i/14)(\zeta + \zeta^{-1}),$$

$$Y = r_3 \exp(-\pi i/14)(-1 - \zeta - \zeta^2),$$

$$Z = r_3 \exp(-\pi i/14)\zeta^{-2}.$$

In the ratios that occur r_3 will cancel as we knew it should. After simplification we find

$$D''_{11} = \frac{3 + i\sqrt{7}}{4},$$

$$D''_{22} = \frac{-3 + i\sqrt{7}}{4},$$

$$D''_{33} = \frac{-3 + i\sqrt{7}}{4}.$$

This can be inverted trivially to give

$$\tau = \begin{pmatrix} \frac{-1+i\sqrt{7}}{4} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{-1+i\sqrt{7}}{4} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{-1+i\sqrt{7}}{4} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} e & 1 & 1 \\ 1 & e & 1 \\ 1 & 1 & e \end{pmatrix}, \tag{4.13}$$

where $e = \frac{-1+i\sqrt{7}}{2}$.

4.8 Relation to period matrix of Rauch and Lewittes

In [29], Rauch and Lewittes have already described a canonical homology basis for this curve and calculated the associated Riemann period matrix. We would like to relate the result they obtained (in particular their homology basis) to our own above.

Recall that changing the basis for the homology group produces a natural action on the period matrix τ; if the symplectic matrix giving the basis change is written as

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

satisfying

$$A^T D - C^T B = \mathbf{1},$$
$$A^T C = C^T A,$$
$$D^T B = B^T D,$$

then it acts naturally on the matrix of periods by simple matrix multiplication

$$\begin{pmatrix} \mathbf{1} \\ \tau \end{pmatrix} \mapsto \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} \mathbf{1} \\ \tau \end{pmatrix} = \begin{pmatrix} A + B\tau \\ C + D\tau \end{pmatrix} \equiv \begin{pmatrix} \mathbf{1} \\ (C + D\tau)(A + B\tau)^{-1} \end{pmatrix}.$$

This clearly induces the natural action on the Riemann period matrix τ,

$$\tau \mapsto (C + D\tau)(A + B\tau)^{-1}.$$

In its most immediate form, then, our problem is that we have two period matrices τ and τ' arising from different canonical homology bases. We want a symplectic

transformation such that

$$\tau' = (C + D\tau)(A + B\tau)^{-1}.$$

This gives rise to a set of linear and quadratic equations in A, B, C, D

$$\tau' B\tau + \tau' A - D\tau - C = 0,$$
$$A^T D - C^T B = \mathbf{1},$$
$$A^T C = C^T A,$$
$$D^T B = B^T D,$$

which must be solved over $\mathbb{Z}$. The first of these is linear in each entry of the symplectic matrix, and can be easily solved by converting the linear system to Smith normal form. However the remaining equations are quadratic in nature.

The general solution to a system of quadratic equations is unknown. In fact Britton shows in [7] that if linear terms are allowed then there is *no* algorithmic solution. Of course, these equations are not completely general and may well fall into a subset for which an algorithm can be produced; but a naïve approach will not work.

This problem is actually solvable via another route. Siegel gave a fundamental domain for the space of period matrices under the action of the symplectic group, described by Klingen in [24]. Unfortunately the algorithm for putting a period matrix into canonical form relies on Minkowski reduction of a lattice, for which there is no known algorithm that executes in polynomial time. Thus the algorithm consumes time exponential in the genus g and so practical calculations, as in this thesis, must rely on alternate means to find basis transformations.

It should be noted that Maple's algcurves library includes a function `Siegel` modelled after this algorithm. However, it substitutes a Minkowski reduction with the (fast but approximate) Lenstra–Lenstra–Lovász lattice reduction algorithm. As a result the output of the algorithm, while lying in a restricted subset of the space of all possible period matrices, does not necessarily lie in the fundamental domain and so is not canonical.

For example our period matrix for Klein's curve is

$$\tau_1 = \begin{pmatrix} \frac{-1+i\sqrt{7}}{4} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{-1+i\sqrt{7}}{4} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{-1+i\sqrt{7}}{4} \end{pmatrix},$$

and Rauch and Lewittes' is

$$\tau_2 = \begin{pmatrix} \frac{-1+3i\sqrt{7}}{8} & \frac{-1-i\sqrt{7}}{4} & \frac{-3+i\sqrt{7}}{8} \\ \frac{-1-i\sqrt{7}}{4} & \frac{1+i\sqrt{7}}{2} & \frac{-1-i\sqrt{7}}{4} \\ \frac{-3+i\sqrt{7}}{8} & \frac{-1-i\sqrt{7}}{4} & \frac{7+3i\sqrt{7}}{8} \end{pmatrix}.$$

These should be equivalent, but applying Maple's `Siegel` algorithm results in the different matrices (truncated to 3 d.p.)

$$\tau_1' \cong \begin{pmatrix} -0.125 + 0.992i & 0.375 - 0.331i & -0.375 + 0.331i \\ 0.375 - 0.331i & -0.125 + 0.992i & -0.375 + 0.331i \\ -0.375 + 0.331i & -0.375 + 0.331i & -0.125 + 0.992i \end{pmatrix},$$

$$\tau_2' \cong \begin{pmatrix} -0.125 + 0.992i & 0.375 + 0.331i & +0.375 + 0.331i \\ 0.375 + 0.331i & 0.125 + 0.992i & -0.375 - 0.331i \\ 0.375 + 0.331i & -0.375 - 0.331i & 0.125 + 0.992i \end{pmatrix}.$$

The situation is actually rather worse than suggested just by this example. In this specific case, while the matrices aren't identical they are clearly related and a little work would allow us to go between τ_1' and τ_2'. Unfortunately, not only is there no reason for this to be true in general, but it is expected that increasing the complexity (genus) makes the correspondence progressively more tenuous.

Indeed, applying the `Siegel` to two period matrices for Bring's curve (the `algcurves` and Riera matrix, see next chapter for details) we obtain something much less useful,

$$\tau_1' \cong \begin{pmatrix} .2500 + 1.0348i & -.2500 + .1015i & .2500 - .1015i & .2500 + .6696i \\ -.2500 + .1015i & .2500 + 1.0348i & .2500 - .1015i & .2500 + .6696i \\ .2500 - .1015i & .2500 - .1015i & .2500 + 1.0348i & -.2500 - .6696i \\ .2500 + .6696i & .2500 + .6696i & -.2500 - .6696i & 0 + 1.3392i \end{pmatrix},$$

$$
\begin{array}{ccccccc}
e_1 & e_3 & e_5 & e_7 & e_9 & e_{11} & e_{13} \\
\updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow \\
e_6 & e_8 & e_{10} & e_{12} & e_{14} & e_2 & e_4
\end{array}
$$

Table 4.3: Side identifications in Klein's curve

and

$$
\tau_2' \cong \begin{pmatrix}
-.5000 + .8685\mathrm{i} & .5000 - .2678\mathrm{i} & .5000 - .2678\mathrm{i} & 0 - 0.649e - 1\mathrm{i} \\
.5000 - .2678\mathrm{i} & -.5000 + .8685\mathrm{i} & 0 - 0.649e - 1\mathrm{i} & .5000 - .2678\mathrm{i} \\
.5000 - .2678\mathrm{i} & 0 - 0.0649e\mathrm{i} & .5000 + .8685\mathrm{i} & -.5000 - .2678\mathrm{i} \\
0 - 0.649e - 1\mathrm{i} & .5000 - .2678\mathrm{i} & -.5000 - .2678\mathrm{i} & -.5000 + .8685\mathrm{i}
\end{pmatrix}.
$$

With this in mind, our approach to Rauch and Lewittes' basis will be much more direct. If we can describe the two bases in the same (algebraic) setting then the tools provided by `extcurves` will allow us to easily write down the symplectic transformation between them.

So our task is to build a strong enough correspondence between the hyperbolic and algebraic models of Klein's curve that we can carry out this procedure.

4.8.1 Hyperbolic model of Klein's curve

As described in [22], the hyperbolic model arises from the quotient of Poincaré's hyperbolic disc with a discrete Fuchsian group of symmetries and amounts to a regular 14-gon (shown in Figure 4.6) centred in the disc with the identification of vertices and pairs of sides as given by Table 4.3. As can be seen in Figure 4.6, it may be tiled by 336 triangles with angles $(2\pi/7, \pi/3, \pi/2)$, each of which can form a fundamental domain of the symmetry group (extended by the antiholomorphic involutions). This group is then manifest as rotations about any triangular vertex together with reflections in any geodesic line consisting of triangular edges.

Rauch and Lewittes' homology basis is then described in terms of paths back and forth between P_1 and P_2 along prescribed edges. Taking positive numbers to indicate anticlockwise traversal (about the centre) and negative the reverse, the basis is explicitly

$$
\begin{array}{lll}
\mathfrak{a}_1 = 1 - 4 - 7 - 9, & \mathfrak{a}_2 = -4 - 9, & \mathfrak{a}_3 = -4 - 5, \\
\mathfrak{b}_1 = 2 + 3 + 4 + 5, & \mathfrak{b}_2 = -3 + 7, & \mathfrak{b}_3 = 3 - 5.
\end{array}
$$

Figure 4.6: Hyperbolic disc model of Klein's curve.

(So for example $\mathfrak{a}_1$ consists of moving from P_1 to P_2 along side 1, back to P_1 along 7, to P_2 again along 4 and finally back to the start along 9 – there is ambiguity in the order of edges taken, but all resulting paths are homologous).

It will be our goal to express this as a path in some coordinate space so that we can apply the tools developed in `extcurves` and `CyclePainter` to relate this basis to the earlier one.

4.8.2   Identification of two models

We now wish to identify the two models of Klein's curve so that we can express Rauch and Lewittes' homology basis as a path in some coordinate plane. The first step here is to realise that each of our coordinates are essentially meromorphic functions on the surface. However, in this light, any symmetry is seen to send one meromorphic function to a different one, so there is no unique choice of how to map our coordinates to the disc.

We can turn this restriction to our advantage, however, by picking a correspondence that has nice properties on the disc. We work in (t, s) coordinates for reasons that will soon become apparent.

Consider the subgroup of the full automorphism group generated by

$$
\begin{aligned}
Z &: (t, s) \mapsto (t, \zeta s), \\
R &: (t, s) \mapsto \left(1 - \frac{1}{t}, \frac{t - 1}{s^3}\right), \\
c &: (t, s) \mapsto (\bar{t}, \bar{s}).
\end{aligned}
$$

These satisfy the relations $RZR^{-1} = Z^4$, $cRc^{-1} = R$, $cZc^{-1} = Z^{-1}$ which allow us to express any element of the group as $c^i Z^j R^k$ and hence the group has order 42 (the explicit representation above tells us that each of Z, R, c are nontrivial, so the relations simply serve to limit the size of the group).

It is easy to verify that the symmetry group of Klein's curve (which is isomorphic to $PGL(2, 7)$) has eight subgroups of order 42, all of which are conjugate. So if we can find an appropriate isomorphic group on the hyperbolic model (one which interacts well with Rauch and Lewittes' basis) then by applying some automorphism we may assume without loss of generality that this corresponds to our group $\langle Z, R, c \rangle$.

Z must correspond to some rotation of (an integer multiple of) $2\pi/7$ about the

centre of one of the septagons. The natural choice is, of course, the central septagon. We note that any such rotation will fix precisely 3 points. In the hyperbolic model it is easy to see that these are P_0, P_1, P_2 in Figure 4.6. In fact a single rotation can equally be considered as a rotation about any of these points.

At this stage we can see it would be prudent to use either (z, w) or (t, s) coordinates. In either case P_0, P_1, P_2 would be identified with the branch points (which are fixed by Z), and so Rauch and Lewittes' homology could be expressed as traversing specified paths between branch points – easy to convert to our notation.

Now we note that R cyclically permutes the branch points in coordinate space, so it must permute P_0, P_1, P_2 in hyperbolic space. There are 14 vertices on the disc which possess order three rotations fixing $\{P_i\}$. These are indicated as $R_\bullet$ in Figure 4.6 and come in pairs: a rotation of $\frac{2\pi}{3}$ about one such point is equally a rotation of $\frac{4\pi}{3}$ about one of the others. We must include all (seven) of these rotations in our hyperbolic group of order 42 (they are conjugate under Z), but by using that conjugacy we will favour the rotation about R_0 and make it correspond to the automorphism R.

Exactly which rotation fixing R_0 we need is determined by the relation $RZR^{-1} = Z^4$, since if we pick the wrong one we will discover $R^2 Z R^{-2} = Z^2$. This requires us to identify R with a rotation of $2\pi/3$ anticlockwise about R_0.

Finally we consider c. It should obviously correspond to a reflection (as it reverses the orientation on the surface). Further it fixes P_0, P_1, P_2 so it must be reflection in a line through (say) the centre. Again all possible choices are in the group, and all choices have the correct commutation with a central rotation; however, by demanding that R correspond to a rotation about R_0 we have already fixed c to be the reflection in the dashed line of Figure 4.6 – other reflections do not commute with rotation about R_0.

The fixed points of c make the choice between (z, w) and (t, s) coordinates obvious. In (t, s) coordinates, the fixed points correspond simply to the real axis, whereas in (z, w) they form the circle $|z| = 1$. The real axis will be easier to work with, so we use (t, s) coordinates.

In summary, given any identification $\sigma : H \to \mathbb{C}^2$ such that $\sigma(p) = (t(p), s(p))$, Z induces the transformation $\bar{Z} : H \to H$ given by

$$\bar{Z}(p) = (\sigma^{-1} \circ Z \circ \sigma)(p).$$

If $\bar{Z}$ is not initially a rotation about the centre of the disc, we replace σ by $\sigma \circ g$ so

that it is. The arguments above showed that by replacing the new σ yet again by $\sigma \circ \bar{Z}^k$ for some k we could demand that $\bar{R} = (\sigma^{-1} \circ R \circ \sigma)(p)$ is a rotation about R_0. Considerations of the relations among these automorphisms then showed in addition that

- $\bar{R}$ is actually a rotation of $2\pi/3$ anticlockwise.

- $\bar{c}$ is the reflection in the dashed line of Figure 4.6.

- $\{\sigma(P_0), \sigma(P_1), \sigma(P_2)\} = \{(0,0), (1,0), (\infty,\infty)\}$.

Moving on, if we replace σ by $\sigma \circ \bar{R}^k$ all of the above properties are unchanged, but in addition we can ask that

$$\sigma(P_0) = (\infty, \infty).$$

We then deduce that $\sigma(P_1) = (1,0)$ and $\sigma(P_2) = (0,0)$.

At this stage we know that traversing side $2/11$ is equivalent to travelling from $t = 0$ to $t = 1$ along the real axis in both t and s. If we knew that $\bar{Z}$ corresponded to a central rotation by $2\pi j/7$ then we could deduce the phase of s along any numbered edge in terms of j. Specifically, along edge $2k$

$$\sigma(p) \in \mathbb{R} \times \zeta^{(k-1)j^{-1}} \mathbb{R},$$

where the inverse is taken mod 7. The odd edges are obtained by their identification with even edges. This would be enough to completely determine any path expressed, as Rauch and Lewittes do, by which numbered edges should be traversed in the hyperbolic model.

So our final task is to identify j, or equivalently which central hyperbolic rotation our coordinate transformation Z corresponds to. The solution is provided by the hyperbolic structure near the point P_1. The idea is that P_1 is the centre of a septagon and if you go towards it on some edge and away on another then in some sense the angle between these paths corresponds to how many branch cuts you would cross doing the same thing in coordinate space. Requiring the same phase for s from both processes is enough to fix j.

More precisely, we can deduce from Figure 4.6 how the numbered edges are laid out around P_1 and P_2. For example the bottom vertex shows us that an anticlockwise rotation of $2\pi/7$ about P_1 will take us from edge 1 to edge 14. Putting all this information

Figure 4.7: How edges come together near P_1

Edge	2/11	4/13	6/1	8/3	10/5	12/7	14/9
Phase	1	ζ^4	ζ	ζ^5	ζ^2	ζ^6	ζ^3

Table 4.4: s phase on each numbered edge

together, we obtain Figure 4.7.

Now notice the unlabelled "spokes" in the diagram. We can think of these as branch cuts in coordinate space. If a branch cut was some ray from $t = 1$ on all sheets, then in hyperbolic space these would be represented by a single line and its images under the sheet-changing operator Z (or some power). These lines won't necessarily be geodesic, but they *will* have a well-defined direction emerging from P_1 and be related by a $2\pi/7$ rotation in Figure 4.7. Thus (for example by considering a small neighbourhood around P_1) we may as well consider the unlabelled "spokes" as the branch cuts.

Now we can put these two pictures together. Suppose we start with both t and s real. If we go around the branch point $t = 1$ once anticlockwise we discover that $s \in \zeta^2 \mathbb{R}$ (cf. Table 4.2 pushed forwards to (t, s) coordinates). In doing so we've crossed just one branch cut.

In the hyperbolic picture what we've done is start on edge 2 and go anticlockwise crossing one spoke. So we are on edge 10. But this corresponds to $s \in \zeta^{4j-1} \mathbb{R}$. For these two statements to be consistent $j = 2$, Z corresponds to a central rotation of $4\pi/7$ and the argument of s along each edge is as in Table 4.4.

Note that in the above there was an ambiguity over the direction of paths – an implicit assumption that anticlockwise was the same in both models. This corresponds to a choice of orientation. Group theoretically if we conjugate by c it takes us between these two choices (because $cZc^{-1} = Z^{-1}$). The choice made gives us a symplectic
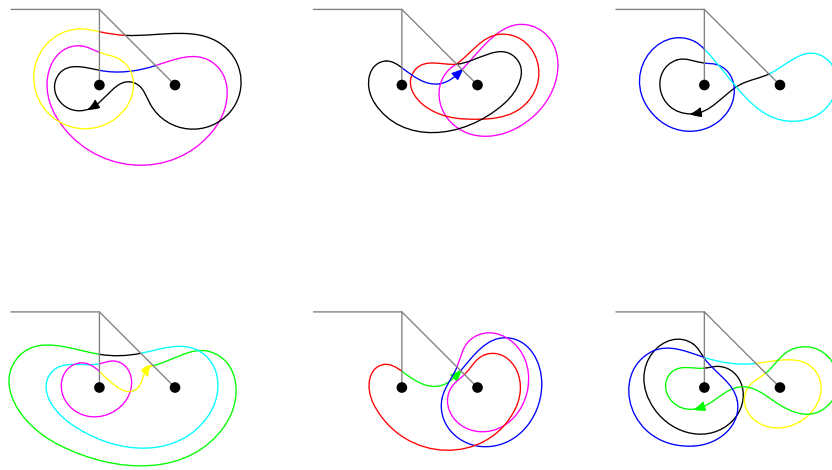
●Sheet 1 ●Sheet 2 ●Sheet 3 ●Sheet 4 ●Sheet 5 ●Sheet 6 ●Sheet 7

Figure 4.8: Rauch and Lewittes' homology basis in (t, s) coordinates.

transformation between the two homology bases, rather than antisymplectic.

4.8.3 Rauch and Lewittes' homology in coordinates

With the above identification, we can simply read off which sheets Rauch and Lewittes' homology basis uses in its repeated journeys between $t = 0$ and $t = 1$. From that information we put together paths that make the same journeys and so draw the basis. We obtain Figure 4.8.

We can algorithmically check that these paths form a canonical homology basis, and indeed we discover the correct intersection matrix.

4.8.4 Symplectic transformation

It is now a simple matter to calculate the symplectic transformation moving from our basis to this one. If our basis is γ_i and the Rauch-Lewittes one γ_i' then the transformation (obtained by applying `find_homology_transform`) is $\gamma_i' = M_{ij}\gamma_j$ where

$$
M = \begin{pmatrix}
1 & -1 & 0 & 1 & -1 & 0 \\
0 & -1 & 1 & 0 & -1 & 1 \\
-1 & -1 & 0 & 0 & 0 & 1 \\
0 & 1 & 0 & 0 & 0 & -1 \\
0 & 0 & 0 & 0 & 0 & 1 \\
-1 & 0 & 0 & -1 & 0 & 0
\end{pmatrix}.
$$

As a final check we note that if

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

in block diagonal form then the action of the symplectic transformation on the Riemann period matrix τ is $\tau \mapsto (C + D\tau)(A + B\tau)^{-1}$. Applied to (4.13) this indeed gives the Rauch-Lewittes period matrix

$$\tau = \begin{pmatrix} \frac{-1+3\mathrm{i}\sqrt{7}}{8} & \frac{-1-\mathrm{i}\sqrt{7}}{4} & \frac{-3+\mathrm{i}\sqrt{7}}{8} \\ \frac{-1-\mathrm{i}\sqrt{7}}{4} & \frac{1+\mathrm{i}\sqrt{7}}{2} & \frac{-1-\mathrm{i}\sqrt{7}}{4} \\ \frac{-3+\mathrm{i}\sqrt{7}}{8} & \frac{-1-\mathrm{i}\sqrt{7}}{4} & \frac{7+3\mathrm{i}\sqrt{7}}{8} \end{pmatrix}.$$

4.9 Vector of Riemann constants

A similar mixture of analysis and numeric work allows us to determine the vector of Riemann constants with respect to our chosen homology basis. Recall that the Abel map based at a point Q of a Riemann surface Σ is the function essentially defined by integrating the g holomorphic differentials from Q to P, that is

$$A_Q(P) = \left(\int_Q^P \omega_1, \ldots, \int_Q^P \omega_g \right).$$

Of course, this is ill-defined as a map $\Sigma \to \mathbb{C}^g$ since it depends on the basis chosen for differentials and the path taken from Q to P. The solution is to define the Abel map onto the Jacobian of Σ instead, considering it well-defined only up to the addition of periods, for any closed path γ, we consider points differing by

$$\left(\int_\gamma \omega_1, \ldots, \int_\gamma \omega_g \right)$$

to be equivalent. Since any such period can be expressed in terms of a homology basis, the result is a quotient of $\mathbb{C}^g$ with a $2g$ dimensional lattice of points – a torus.

The other issue, of what happens when a different basis is chosen for ω_i is less important, it's simply viewed as taking a different basis for the Jacobian. Naturally, a coordinate free abstract definition can be given for the Jacobian (see for example Griffiths and Harris in [20] or, for a more readable and immediately applicable exposition Miranda's book [27]), but for our purposes this concrete version is more useful.

Given this Abel map, recall that Θ-functions can be defined as general meromorphic maps from this Jacobian back to $\mathbb{C}$. (for details, see Farkas and Kra's book [15]). These two types of function then have the property that

$$P \mapsto \Theta(A_Q(P) - e)$$

has precisely g zeroes $P_1, \ldots, P_g$ and there is a vector K_Q, called the vector of Riemann constants and independent of e, such that

$$A_Q(P_1) + \cdots + A_Q(P_g) + K_Q = e.$$

This relation is frequently used in integrable systems to give explicit solutions to differential equations in terms of these theta functions. See for example Babelon and Talon's book [2].

4.9.1   Theory of constants up to a half-period

Symmetries can be most easily applied to finding the vector of Riemann constants via the well-known relation

$$-2K_Q = A(\mathcal{K}_\Sigma), \tag{4.14}$$

where $\mathcal{K}_\Sigma$ is the canonical divisor, i.e. the divisor of a meromorphic differential. Any differential will serve and in all cases will have degree $2g - 2$. If we select a holomorphic differential then the divisor will be positive. It can be taken to consist of the $2g - 2$ points $P_1, \ldots, P_{2g-2}$ and (4.14) may be written as

$$-2K_Q = \sum_{i=1}^{2g-2} A_Q(P_i). \tag{4.15}$$

The symmetries of Klein's curve will completely determine $-2K_Q$ and thus the vector of constants up to adding an arbitrary half-period.

We proceed as follows. Suppose $\phi : \Sigma \to \Sigma$ is a symmetry of Klein's curve and acts on the basis of $\mathfrak{a}$-normalised holomorphic differentials, ω, via the matrix F

$$\phi^*(\omega) = F\omega.$$

We work first on the Abel map itself and discover

$$
\begin{aligned}
FA_Q(P) = F \int_Q^P \omega \\
= \int_Q^P \phi^*(\omega) \\
= \int_{\phi(Q)}^{\phi(P)} \omega \\
= \int_{\phi(Q)}^Q \omega + \int_Q^{\phi(P)} \omega \\
= A_Q(\phi(P)) + \int_{\phi(Q)}^Q \omega.
\end{aligned}
$$

Applying this to (4.15) (noting that $\phi(\mathcal{K}_C) \sim \mathcal{K}_C$, still a canonical divisor) yields

$$
F(-2K_Q) = -2K_Q + (2g - 2) \int_{\phi(Q)}^Q \omega.
$$

If we choose the base, Q, of the Abel map to be a fixed point of ϕ then the second term vanishes and we are left with the even more simple relation

$$
(F - \mathbf{1})(-2K_Q) = 0, \tag{4.16}
$$

where this equality is of course in the Jacobian and hence numerically will only hold up to the addition of a period.

4.9.2 Constraints from order 7 automorphism

We are now in a position to apply the above results directly to Klein's curve. We favour the order 7 automorphism by choosing to base the Abel map at $(0,0)$ in (x, y) coordinates, calling the resulting vector of constants simply K_0. Thus we have

$$
\phi(x, y) = (\zeta^5 x, \zeta^4 y),
$$

and the pullback on our chosen differentials (rather than the $\mathfrak{a}$-normalised ones) is given by the matrix $\hat{F}$

$$
\phi^*(\omega) = \hat{F}\omega = \begin{pmatrix} \zeta^2 & 0 & 0 \\ 0 & \zeta & 0 \\ 0 & 0 & \zeta^4 \end{pmatrix} \omega.
$$

This means that the action on the 𝔞-normalised basis implicit in the Abel map is given by

$$F = \mathcal{A}^{-1}\hat{F}\mathcal{A}.$$

In $\mathbb{C}^6$ rather than the Jacobian, the symmetry constraint (4.16) becomes

$$(F - 1)(-2K_0) = \begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{k},$$

for some vector of integers $\mathbf{k}$, or (since 1 is not an eigenvalue of $\hat{F}$ and so $F - 1$ is invertible)

$$-2K_0 = (F - 1)^{-1}\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{k}. \tag{4.17}$$

At first glance, every choice of $\mathbf{k}$ will give different possible value for $-2K_0$, however we are still only interested in determining $-2K_0$ up to a period and in fact many choices of $\mathbf{k}$ give equivalent vectors. We want to know when $\mathbf{k}$ and $\mathbf{k} + \mathbf{a}$ give equivalent values for $-2K_0$ so our task is to find out which $\mathbf{a} \in \mathbb{Z}^6$ make the equations

$$(F - 1)^{-1}\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{a} = \begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{b},$$

solvable for $\mathbf{b} \in \mathbb{Z}^6$. Multiplying by $F - 1$ again and applying the result (4.10) from Section 4.6 that $F\begin{pmatrix} 1 & \tau \end{pmatrix} = \begin{pmatrix} 1 & \tau \end{pmatrix}M$, we obtain

$$\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{a} = \begin{pmatrix} 1 & \tau \end{pmatrix}(M - 1)\mathbf{b},$$

which, since $\mathbf{a}$ and $\mathbf{b}$ are in $\mathbb{Z}^6$, is equivalent to the set of wholly integer linear equations

$$\mathbf{a} = (M - 1)\mathbf{b}.$$

Calculating the Smith normal form of $M - 1$ we discover unimodular matrices U, V

such that

$$M - 1 = U \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 7 \end{pmatrix} V,$$

giving

$$U^{-1}\mathbf{a} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 7 \end{pmatrix} (V\mathbf{b}).$$

Since V is also invertible over the integers this has a solution provided the final entry of $U^{-1}\mathbf{a}$ is a multiple of 7. That is, $(F - 1)^{-1} \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{a}$ is a period if and only if

$$(U^{-1}\mathbf{a})_6 = -3a_1 + a_2 + 2a_3 - 4a_4 - a_5 - 2a_6 \equiv 0 \pmod 7.$$

In particular any choice of $\mathbf{k}$ in (4.17) leads to $-2K_0$ equivalent to one obtained from the restricted set $\mathbf{k} \in \{(n00000) : n \in \{0, \ldots, 6\}\}$, which means

$$-2K_0 = (F - 1)^{-1} \begin{pmatrix} 1 & \tau \end{pmatrix} \begin{pmatrix} n & 0 & 0 & 0 & 0 & 0 \end{pmatrix}^T$$
$$= n(F - 1)^{-1}\mathbf{e}_1$$

for some $n \in \{0, \ldots, 6\}$.

4.9.3   Using the involution to determine n

Now that we have a very limited number of possibilities for $-2K_0$, and in particular a discrete set, we can use numeric methods to determine n above.

The only other symmetry which fixes $(0, 0)$ is the antiholomorphic involution; a calculation shows that this unfortunately gives us no new information: any n satisfies the constraint derived from it. Of the remaining symmetries, the order 3 automorphism is slightly more complicated to apply numerically since it sends $(0, 0)$ to (∞, ∞). Therefore

we will apply the holomorphic involution again.

The involution sends $(0,0)$ to $(\gamma/\beta, \alpha/\beta)$, and we call its action on the dual basis of differentials T. Then the constraint imposed is

$$(T-1)(F-1)^{-1}n\mathbf{e}_1 = 4\int_{(0,0)}^{(\gamma/\beta,\alpha/\beta)} \omega + \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{l}.$$

In principle this could be checked algebraically: the integral would have to be evaluated fully, but since it is a $\frac{1}{14}$th-period this isn't a huge barrier.

However, it is easier to simply evaluate numerically. We discover fairly quickly that $n = 3$ is the only possible solution and hence

$$K_0 = \frac{3}{2}(F-1)^{-1}\mathbf{e}_1 + \frac{1}{2}\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{m},$$

where $\mathbf{m} \in \{0,1\}^6$.

4.9.4   Final half-period

The value of $\mathbf{m}$ can be determined numerically by its relation to the zeroes of the Θ-function. We pick three points P_1, P_2 and P_3. Then we can directly calculate

$$\mathbf{e} = A_0(P_1) + A_0(P_2) + A_0(P_3) + \frac{3}{2}(F-1)^{-1}\mathbf{e}_1 + \frac{1}{2}\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{m}.$$

Finally we check whether (say) $\Theta(A_Q(P_1) - e) = 0$ for each $\mathbf{m}$. We discover that $\mathbf{m} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 \end{pmatrix}$ is the only possible solution and so

$$\begin{aligned}
K_0 &= \frac{3}{2}(F-1)^{-1}\mathbf{e}_1 + \frac{1}{2}\begin{pmatrix} 1 & \tau \end{pmatrix}\begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 \end{pmatrix}^T \\
&= \frac{3}{2}(F-1)^{-1}\mathbf{e}_1 + \begin{pmatrix} \frac{3+i\sqrt{7}}{8} & \frac{3}{4} & \frac{3}{4} \end{pmatrix}^T.
\end{aligned}$$

With a little fiddling (adding full periods to make the result look neat) we can finally obtain the simple result

$$K_0 = \frac{i}{\sqrt{7}}\begin{pmatrix} 3 & -1 & 5 \end{pmatrix}^T.$$

4.10   Summary

In this chapter we have combined two very different representations of Klein's curve, one as an algebraic plane curve and the other as a quotient of the hyperbolic disc. These

ideas originated in very different fields of mathematics for different reasons, and have largely remained isolated until now.

However, by building a bridge between these two models we were able to transfer results from one setting to the other. In particular we took a homology basis on the hyperbolic side that gives a particularly beautiful form to the period matrix and, by understanding the group structure of this curve were able to present an equivalent basis on the algebraic side, in a form suitable for further calculation.

We then made use of this by calculating the vector of Riemann constants for Klein's curve, an essential component to explicit calculations arising from the application of these ideas to integrable systems. See, for example Babelon and Talon's book [2].

CHAPTER

5

EXAMPLE: BRING'S CURVE

5.1   Introduction

Bring's curve is a Riemann surface of genus 4 admitting the automorphism group S_5, proven to be the maximum available for a genus 4 surface in [5] by Breuer, for example. The fundamental definition, used by Bring in [6] to study solutions of the quintic, is as a subset of $\mathbb{P}^4$ given by

$$x_1 + x_2 + x_3 + x_4 + x_5 = 0,$$
$$x_1^2 + x_2^2 + x_3^2 + x_4^2 + x_5^2 = 0,$$
$$x_1^3 + x_2^3 + x_3^3 + x_4^3 + x_5^3 = 0.$$

The S_5 symmetry is manifest as permutations of the coordinates x_i.

As with Klein's curve, we will study plane algebraic and hyperbolic representations of Bring's curve. The first representation comes from Dye's paper [13], where he explicitly gives a sextic plane curve and proves its equivalence to Bring's. The remarkable fact about this representation is that of the full S_5 symmetry group, A_5 is generated by projectivities in $\mathbb{P}^2$.

In [9], Craig studies the rational points of a second genus 4 sextic which possesses at least A_5 as a symmetry group. In fact this curve is very closely related to Dye's representation and we will show that it too is equivalent to Bring's curve by giving an explicit transformation of $\mathbb{P}^2$ mapping Dye's curve to Craig's. Craig's representation will be seen to be more useful for our purposes since it has more obvious real structure and simpler branching properties.

On the hyperbolic side Riera and Rodríguez, in [30], studied a representation much like that of Klein's curve and produced a very simple period matrix of the form

$$\tau = \tau_0 \begin{pmatrix} 4 & 1 & -1 & 1 \\ 1 & 4 & 1 & -1 \\ -1 & 1 & 4 & 1 \\ 1 & -1 & 1 & 4 \end{pmatrix},$$

for a determined $\tau_0 \in \mathbb{C}$. The period matrix already exhibits much of the symmetry implicit in the automorphism group and we won't attempt to improve on this result, but we relate this representation to Craig's plane algebraic one and provide an equivalent homology basis in that setting.

Finally we compute the vector of Riemann constants for Bring's curve in Riera and Rodríguez's homology basis. We first use an almost purely numeric method to obtain a result suitable for computations, then we apply a hybrid algebraic-numeric approach which allows us to derive an algebraic form much as in the Klein case.

5.2   Dye's sextic

Let $j = \frac{1+\sqrt{5}}{2}$, i.e. a root of $j^2 = j + 1$. In [13], Dye studies the plane sextic curves given by

$$
\begin{aligned}
\mathcal{D}_\lambda(x,y,z) := {} & (x + jy)^6 + (x - jy)^6 + (y + jz)^6 \\
& + (y - jz)^6 + (z + jx)^6 + (z - jx)^6 + \lambda(x^2 + y^2 + z^2)^3 = 0.
\end{aligned}
$$

For generic $\lambda \in \mathbb{C}$ the curve has genus 10, but if λ is chosen to be $-\frac{78+104j}{5}$ then the genus drops to 4 and the resulting curve is shown in [13] to be equivalent to Bring's. We correspondingly define

$$
\mathcal{D}(x,y,z) := \mathcal{D}_{-\frac{78+104j}{5}}(x,y,z).
$$

5.2.1   Automorphisms of Dye's sextic

$\mathcal{D}(x,y,z)$ has the obvious order three cyclic symmetry

$$
b' : (x,y,z) \mapsto (y,z,x),
$$

as well as the less obvious order two symmetry

$$
a' : \begin{pmatrix} x \\ y \\ z \end{pmatrix} \mapsto \begin{pmatrix} -j & 1 & j^2 \\ 1 & -j^2 & j \\ j^2 & j & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix},
$$

both presented by Dye in his paper [13]. The primes in these equations are for convenience later; it will become apparent that we really want to consider a slightly different set of generators to those given by Dye.

Note that there are also obvious order-2 symmetries produced by changing the sign of one or more of x, y, z. However, these only generate a group of order 24 so Dye

needed to look further for the richer group.

It is easy to check that these a' and b' are the classical generators for A_5:

$$a'b' = \begin{pmatrix} j^2 & -j & 1 \\ j & 1 & -j^2 \\ 1 & j^2 & j \end{pmatrix} \tag{5.1}$$

which has order 5 (taking into account the projective nature of the space). So

$$a'^2 = b'^3 = (a'b')^5 = 1.$$

Recall that the entire symmetry group of Bring's curve is S_5 so there are more symmetries to be found. However these need not be expressible as a simple matrix action on the projective coordinates (x, y, z). Indeed, as Dye notes, the surprising fact is that $A_5 \leq S_5$ *is* realized this simply.

5.3 Craig's sextic

Craig describes in [9] a sextic equation, also of genus 4 and admitting A_5 as a symmetry group:

$$\mathcal{C}(\tilde{x}, \tilde{y}, \tilde{z}) := \tilde{x}(\tilde{y}^5 + \tilde{z}^5) + (\tilde{x}\tilde{y}\tilde{z})^2 - \tilde{x}^4\tilde{y}\tilde{z} - 2(\tilde{y}\tilde{z})^3 = 0. \tag{5.2}$$

In this case an order 5 symmetry is obvious and we may take

$$\widetilde{ab} : (\tilde{x}, \tilde{y}, \tilde{z}) \mapsto (\zeta^2\tilde{x}, \zeta^4\tilde{y}, \tilde{z}), \tag{5.3}$$

where $\zeta = e^{2\pi i/5}$.

There is also a corresponding order two symmetry described by Craig

$$\widetilde{a} : \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix} \mapsto \begin{pmatrix} 1 & 2 & 2 \\ 1 & \zeta + \zeta^{-1} & \zeta^2 + \zeta^{-2} \\ 1 & \zeta^2 + \zeta^{-2} & \zeta + \zeta^{-1} \end{pmatrix} \begin{pmatrix} \tilde{x} \\ \tilde{y} \\ \tilde{z} \end{pmatrix}.$$

Together these generate A_5 again since $\widetilde{a}\widetilde{ab} =: \widetilde{b}$ has order 3 (hence the slightly unusual choice of notation for $\widetilde{ab}$ above).

This representation will turn out to be the most convenient for later work so it is worth spending some time on its detail, particularly its desingularisation.

5.3.1   Special points in Craig's representation and desingularisation

The points at infinity for Craig's curve (5.2) are naïvely given by $[0, 1, 0]$ and $[1, 0, 0]$, however the latter is singular. In fact the singularities of Craig's curve are $[1, 0, 0]$ and $[\zeta^k, \zeta^{2k}, 1]$ for $k \in \{0, \ldots, 4\}$ so we must work out expansions nearby in order to form a compact curve.

We will follow the procedure detailed by Kirwan, for example, in the very readable book [23]. Briefly, the idea is to enumerate possible series expansions near the singular points. This can be done either as so-called Puiseux series involving fractional powers of the usual variables or in terms of another, independent complex parameter (t, say) which acts as a true manifold coordinate.

Regardless of that detail, when we have the possible expansions, we will find that some cannot possibly be related by a diffeomorphism (i.e. coordinate change). There is a large theory behind this (much of it developed by Brieskorn, see for example his book with Knörrer [32]), but the result is that when you identify points on a desingularised space by these expansions, a compact nonsingular manifold structure can be imposed.

First the infinite singularity: consider the structure near $[1, 0, 0]$, say at points $[1, y, z]$ for small y, z. The curve reduces to

$$y^5 + z^5 + y^2 z^2 - yz - 2y^3 z^3 = 0,$$

so in the usual Puiseux construction (again, see [23] for details) we suppose $z = Ay^\alpha + \cdots$ (with $\alpha > 0$) and equate coefficients of the lowest order. Of course, which term is lowest order depends on α. There's a well-defined algorithm involving Newton polygons for determining candidates, but in this particular case it should be clear that regardless of α, the only candidates are yz, z^5 and y^5.

Equating lowest order terms in the valid combinations we get one of

- $A^5 y^{5\alpha} - Ay^{\alpha+1} = 0$ which implies $z = y^{1/4}$.

- $y^5 - Ay^{\alpha+1} = 0$ which implies $z = y^4 + \cdots$.

The second of these gives a single z for each y near 0, the first gives four different values for z. Together these make up the expected five sheets and so expansions after this point are unique. All expansions in the first equation are conjugate under a coordinate change, arising from points $[1, 0, 0] \sim [1, t^4, t]$ where t is a manifold coordinate. Hence

the point $[1, 0, 0]$ desingularises into precisely two points on the nonsingular curve:

$$[1, 0, 0]_1 \sim [1, t^4, t], \qquad\qquad\qquad y \sim x^{-3} \qquad\qquad (5.4)$$

$$[1, 0, 0]_2 \sim [1, t, t^4]. \qquad\qquad\qquad y \sim x^{3/4}$$

For the remaining singular points, we only need to explicitly investigate one and then note that the automorphism $[x, y, z] \mapsto [\zeta x, \zeta^2 y, z]$ will tell us how the other singularities behave (i.e. very similarly).

So we look at $[1, 1, 1]$. At first sight two of the sheets come together here. Consider $[1 + \epsilon, y, 1]$ near to $[1, 1, 1]$. To first order

$$y^5 - 2y^3 + y^2 - y + 1 = 0. \qquad\qquad (5.5)$$

This quintic has four distinct roots: two are complex, corresponding to nonsingular points and will play no role in future developments. There is a real negative root (approximately -1.7549) which also corresponds to a nonsingular point and will occur later. Finally 1 is a root, which gives us the expected singularity at $[1, 1, 1]$.

Expanding about this singular point, at the next order we discover

$$\begin{aligned} y &= 1 + \epsilon \frac{1 + \sqrt{5}}{2} + \cdots, \\ y &= 1 + \epsilon \frac{1 - \sqrt{5}}{2} + \cdots. \end{aligned} \qquad\qquad (5.6)$$

These are clearly distinct solutions and together with the nonsingular expansions exhaust the five possible nonsingular preimages near $x = 1$, so $[1, 1, 1]$ also desingularises to two distinct points.

5.3.2 Branched covers of $\mathbb{P}^1$ and real paths

We now consider Craig's curve as a branched cover of $\mathbb{P}^1$ with x as the coordinate. The equation obtained by setting $z = 1$ in (5.2) is

$$xy^5 + x + x^2 y^2 - x^4 y - 2y^3 = 0. \qquad\qquad (5.7)$$

There are 5 sheets above the generic x, with branch points at $0, \infty$ and

$$\left\{ \frac{\zeta^k}{4} \left(1674 \pm 870\mathrm{i}\sqrt{15} \right)^{1/5} : k \in \{0, \ldots, 4\} \right\}.$$

Naïvely there is also a double solution at $x = \zeta^k$ but these are precisely the singular points similar to $[1, 1, 1]$ we investigated before and after resolution the cover is regular there.

At $x = 0$ we have two preimages, one corresponding to $[0, 0, 1]$ with an expansion

$$y = \frac{1}{2^{1/3}} x^{1/3} + \cdots , \tag{5.8}$$

and the other to $[0, 1, 0]$ with expansion

$$y = \sqrt{2} x^{-1/2} + \cdots . \tag{5.9}$$

Similarly at $x = \infty$ we have two preimages after desingularisation, $[1, 0, 0]_1$ and $[1, 0, 0]_2$. The other branch points will play a much less critical role in the future so we don't go into details.

Real paths on Craig's curve

The real paths in this cover will be of particular interest later on so we will take some time to explore their nature now. Firstly, if the number of real roots of (5.7) considered as a polynomial in y changes then its discriminant

$$- x^3 (256 x^{20} - 1349 x^{15} + 5386 x^{10} - 7749 x^5 + 3456) \tag{5.10}$$

must vanish there. The only real roots of this equation are $x = 0, 1$, so we are reduced to considering the intervals $(-\infty, 0), (0, 1), (1, \infty)$.

- If $x < 0$ then there is just one real root.

- If $0 < x < 1$ then there are three real roots.

- If $x > 1$ then there are also three real roots.

Referring to the expansions (5.8) and (5.9) we see that a real path starting with $x < 0$ moving towards $x = 0$ must be approaching $[0, 0, 1]$ along the expansion $y = 2^{-1/3} x^{1/3} + \cdots$ (i.e. $y \to 0$ too). Continuity demands that when extended past $x = 0$ it should have y small and positive for small $x > 0$ too. We will call this path γ_0.

We now turn our attention to another real path approaching $x = 0$, this time from above. It must lie on the expansion $y = \sqrt{2} x^{-1/2} + \cdots$ and hence y is either large and

positive or large and negative; we will call these paths γ_+ and γ_-. In fact the expansion is telling us that γ_+ and γ_- meet at $[0, 1, 0]$ and we could form a single continuous path, but we will maintain the distinction for now.

In summary we have three real paths coming out of $x = 0$ along the positive axis, satisfying (for small $x > 0$),

$$y(\gamma_-) \ll 0 < y(\gamma_0) \ll y(\gamma_+).$$

Now we are ready to consider what happens at $x = 1$. On the desingularised curve there are three real points here (the two from desingularising $y = 1$ and the remaining real root of about -1.7549). Each of the curves coming out of $x = 0$ must pass through one of them. Further, the order of the y values among the paths must be the same approaching $x = 1$ as it was leaving $x = 0$ since otherwise they would have crossed in between and this would have shown itself in (5.10).

The three expansions near $x = 1$ in order of increasing y for $x < 1$ are

$$y \sim \alpha \qquad\qquad y \sim 1 + (x-1)\frac{1+\sqrt{5}}{2} \qquad\qquad y \sim 1 + (x-1)\frac{1-\sqrt{5}}{2},$$

where $\alpha < 0$ is the remaining real solution to (5.5). Thus the path that started $y \ll 0$ must pass through the first point, $y \sim 0$ must pass through the second and $y \gg 0$ the third. Significantly this means the latter two paths actually cross at $x = 1$ and for $x = 1 + \epsilon$ we have

$$y(\gamma_-) < y(\gamma_+) < 1 < y(\gamma_0).$$

Finally we consider the points at ∞. Recall the expansions (5.4). If $x \ll 0$ then naturally there is only one real path, which arrives at $[1, 0, 0]_1$ with small y. If $x \gg 0$ the situation is very similar to $x = 0$: two expansions with $|y| \gg 0$ arriving at $[1, 0, 0]_2$ and one in the middle with $y \sim 0$. As before the paths cannot have crossed between $x = 1$ and $x = \infty$ and so we are forced to conclude that γ_- has the expansion $y \sim -x^{3/4}$, γ_+ has the expansion $y \sim x^{-3}$ and γ_0 has the expansion $y \sim x^{3/4}$ near ∞.

Putting these facts together we can plot Figure 5.1 (the joined semicircular dots represent the same point on the curve, separated to show the distinct y values of paths entering them). We discover that all the paths ($\gamma_-, \gamma_0, \gamma_+$ and the $x < 0$ path) actually form part of one large closed loop.

$y \sim \sqrt{2}x^{-1/2}$

$x = 0$

$y \sim 2^{1/3}x^{1/3}$

$y \sim -\sqrt{2}x^{-1/2}$

γ_- γ_+ γ_0

$y \sim -x^{3/4}$

$y \sim x^{-3}$

$x = \infty$

$y \sim x^{3/4}$

Figure 5.1: Real paths on Craig's curve as a branched cover of $\mathbb{P}^1$.

5.4 Relating two sextics

Since these are two homogeneous sextic equations, if they both describe Bring's curve we might hope they are related by a simple linear transformation on the coordinates. It turns out that the symmetries of the curves constrain this supposed linear map completely and leave just one candidate isomorphism which can easily be tested. In conjunction with Dye's result this analysis proves that Craig's sextic is indeed another representation of Bring's curve.

5.4.1 Preliminary group theory

Suppose $\mathbf{x} \mapsto A^{-1}\mathbf{x}$ maps Dye's curve to Craig's, that is

$$\mathcal{D}(\mathbf{x}) = 0 \iff \mathcal{C}(A^{-1}\mathbf{x}) = 0. \tag{5.11}$$

The matrix A cannot be unique for two related reasons. First there is the intrinsically projective nature of the curves in question. If A satisfies (5.11) then any nonzero scalar multiple λA will too since

$$\mathcal{C}(\lambda^{-1}A^{-1}\mathbf{x}) = \lambda^{-6}\mathcal{C}(A^{-1}\mathbf{x}) = 0 \iff \mathcal{C}(A^{-1}\mathbf{x}) = 0.$$

In fact A and λA describe the same transformation on the projective spaces and this is just the relevant manifestation of that property.

But more freedom is granted by the automorphisms. If $\widetilde{\sigma}$ is an automorphism on Craig's curve and τ on Dye's then consider

$$A' := \tau^{-1} \circ A \circ \widetilde{\sigma}^{-1}.$$

A simple argument shows that A' will also act to transform between the representations. We will use this fact to find a candidate for A.

Also, given an isomorphism A we can conjugate any automorphism on Dye's curve to one on Craig's. In more concrete terms, if $\widetilde{\tau}$ is an automorphism of Dye's curve then

$$A^{-1}\widetilde{\tau}A$$

will be an automorphism on Craig's.

Combining this with the previous result, we obtain the fact that for any $\widetilde{\sigma}$ in the automorphism group of Craig's curve

$$\widetilde{\sigma}^{-1}(A^{-1}\tau A)\widetilde{\sigma}$$

is also an automorphism of Craig's curve. And hence, if we have a favourite (conjugate) automorphism on each side we can demand that our isomorphism A sends one to the other by replacing it with $A\sigma$ if necessary.

5.4.2   First constraint on A

In A_5 there are two conjugacy classes for elements of order 5. In matrix representations, eigenvalues are preserved under conjugacy since

$$\det(BAB^{-1} - \lambda\mathbf{1}) = \det\left(B(A - \lambda\mathbf{1})B^{-1}\right) = \det(A - \lambda\mathbf{1}).$$

In our projective case, of course, eigenvalues will only be preserved up to a scalar multiplication. Equivalently, ratios of eigenvalues will be preserved but not necessarily the values themselves. Looking at (5.1) we find that $a'b'$ has eigenvalues

$$\left\{1 + \sqrt{5}, 1 + \mathrm{i}\sqrt{5 + 2\sqrt{5}}, 1 - \mathrm{i}\sqrt{5 + 2\sqrt{5}}\right\},$$

and the other conjugacy class (represented by $(a'b')^2$ say) in contrast[1] has eigenvalues

$$\left\{ 2(3+\sqrt{5}), -2(2+\sqrt{5}-i\sqrt{5+2\sqrt{5}}), -2(2+\sqrt{5}+i\sqrt{5+2\sqrt{5}}) \right\}.$$

The order 5 element $\widetilde{ab}$ on the Craig side given by (5.3) obviously has eigenvalues

$$\left\{ \zeta^2, \zeta^4, 1 \right\}.$$

Now, if we divide the eigenvalues of $(a'b')^2$ by $2(3+\sqrt{5})$ we obtain this same Craig set, so the ratios of eigenvalues in $(a'b')^2$ are compatible with it corresponding to Craig's $\widetilde{ab}$. The same is not true for $a'b'$ no matter how we try to arrange it. So at this stage we modify our generators for the Dye automorphism group slightly. We define

$$ab = (a'b')^2$$

$$= \begin{pmatrix} 2+2j & -2j & 2+4j \\ 2j & -2-4j & -2-2j \\ 2+4j & 2+2j & -2j \end{pmatrix}, \tag{5.12}$$

also of order 5. We need a replacement for a' too since $a'(ab)$ does not have order 3.

$$a := a'b'^2a'b'a'$$

$$\equiv \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}$$

suffices $(a^2 = b^3 = (ab)^5 = 1)$ and we can now demand that A conjugates ab to $\widetilde{ab}$, i.e.

$$A^{-1}(ab)A = \lambda(\widetilde{ab})$$

$$A^{-1} \begin{pmatrix} 2+2j & -2j & 2+4j \\ 2j & -2-4j & -2-2j \\ 2+4j & 2+2j & -2j \end{pmatrix} A = \lambda \begin{pmatrix} \zeta^2 & 0 & 0 \\ 0 & \zeta^4 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

---

[1]The fact that these are incompatible incidentally proves that we would have to look further afield than projectivities for the full S_5 symmetry group. If this representation could be extended then these two elements would be conjugate and hence have the same eigenvalues.

for some λ. In fact this is simply a matter of diagonalising ab and a solution is given by the columns of A taking on the eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ of ab in (5.12)

$$A = \langle \mathbf{v}_1 | \mathbf{v}_2 | \mathbf{v}_3 \rangle \, .$$

It is easy to see that this solution is essentially unique and the only freedom remaining is to scale each column of A independently giving possible isomorphisms

$$A = \langle p_1 \mathbf{v}_1 | p_2 \mathbf{v}_2 | p_3 \mathbf{v}_3 \rangle \, .$$

5.4.3   Second constraint on A

To identify these scaling factors p_i we must look at another automorphism. In A_5 with the standard generators $a^2 = b^3 = (ab)^5$ there are precisely 5 elements, x, of order 2 with the property that xab has order 3. Specifically these are a and its conjugates under powers of ab. Since this conjugation leaves ab fixed we are still free to use it in our search for the matrix A; it allows us to demand, in addition to what we already know, that

$$A^{-1} a A = \mu \widetilde{a}$$

(for some μ) rather than any other element of order 2 in Craig's representation. Solving the resulting equations yields (scalar multiples of)

$$A = \begin{pmatrix} j & 1 & 1 \\ 0 & -\mathrm{i}\sqrt{2+j} & \mathrm{i}\sqrt{2+j} \\ 1 & -j & -j \end{pmatrix} . \tag{5.13}$$

So we now have a candidate A. If any matrix is going to provide the transformation between the two curves, this will. In fact it would be very odd if it failed at this stage: it certainly induces a correspondence between the two representations of A_5 involved. An algebraic calculation confirms that

$$\mathcal{D}(A\mathbf{x}) = -960(9 + 4\sqrt{5})\mathcal{C}(\mathbf{x}),$$

and hence

$$\mathcal{D}(\mathbf{x}) = 0 \iff \mathcal{C}(A^{-1}\mathbf{x}) = 0,$$

and so Craig's curve is just another representation of Bring's.

5.5 Riera and Rodríguez hyperbolic model

5.5.1 Introduction to H

Riera and Rodríguez, in [30], give Bring's curve as a quotient, H, of the hyperbolic disc. They then proceed to calculate a period matrix taking account of the symmetries of the curve.

The essential features of the model can be seen in Figure 5.2. The surface is seen to be a 20-gon with edges identified as shown in the table below the figure. This leads to vertices of the polygon falling into three equivalence classes, also annotated in the figure. Naturally, this surface has genus 4.

For future calculations it will also be very useful to know exactly how the splintered points P_2 and P_3 come together as two regular points on a manifold. This can be reconstructed quite easily from Figure 5.2. For example, start near P_1 in the bottom right quadrant on edge 2/9. Make a small arc around P_1 proceeding anticlockwise and you will next reach edge 1/14. Repeating at edge 14 tells us that we next meet 6/13. If this procedure is continued we obtain Figure 5.3.

The polygon can be tiled by 120 double triangles (one can take a sector of the central pentagon as a fundamental domain). Now consider the automorphism group. Let d be a rotation of $\frac{\pi}{2}$ about a vertex of the central pentagon and c be a rotation of π about the midpoint of an adjacent pentagon edge. Then clearly $c^2 = d^4 = 1$. But it is also easy to see that cd is a rotation of $\frac{2\pi}{5}$ about the centre and hence $(cd)^5 = 1$. c and d are thus the classical generators of S_5 and this describes the entire automorphism group of Bring's curve.

Riera and Rogríguez give the homology basis for this model by prescribing which edges of the polygon to traverse. We are going to construct an equivalent basis in Craig's model by understanding an isomorphism

$$f : H \to \left\{ (x, y, z) \in \mathbb{C}^3 : \mathcal{C}(x, y, z) = 0 \right\} \tag{5.14}$$

Edge identifications								
1	$\leftrightarrow 14$	5	$\leftrightarrow 18$	9	$\leftrightarrow 2$	13	$\leftrightarrow 6$	17	$\leftrightarrow 10$
3	$\leftrightarrow 12$	7	$\leftrightarrow 16$	11	$\leftrightarrow 20$	15	$\leftrightarrow 4$	19	$\leftrightarrow 8$

Figure 5.2: Riera and Rodríguez hyperbolic model, H, of Bring's curve

Figure 5.3: Conformal structure of P_1, P_2 and P_3

well enough to determine the precise values each edge of the polygon in Figure 5.2 maps to. Once this is achieved converting the homology basis will be a purely mechanical affair as in the Klein case. Along the way we will gain some understanding of how f acts on the automorphism group by conjugation.

5.5.2 Riera and Rodríguez's basis

We start by recapitulating the hyperbolic basis we will be interested in. Riera and Rodríguez begin with a simple non-canonical basis. They first define in [30]

$$\alpha_1 = e_1 + e_2,$$

$$\alpha_2 = e_3 + e_4$$

(in edge traversal notation). Next they act on these by rotations of $\frac{2\pi k}{5}$ to obtain their initial basis. So essentially

$$\alpha_i = e_{2i-1} + e_{2i}.$$

Next they specify (by fiat) a matrix which transforms these α_i into a canonical basis

and proceed to derive further basis change to make use of the symmetries. The end result is the following basis-change matrix (implicit in [30])

$$
\begin{pmatrix}
1 & 0 & 0 & 0 & -2 & 0 & 1 & 0 \\
1 & -1 & 0 & -1 & -1 & 1 & 1 & 1 \\
1 & -1 & 0 & 0 & -1 & 2 & 1 & -1 \\
0 & -1 & 0 & 0 & 1 & 2 & 0 & 0 \\
1 & -1 & 1 & 0 & -1 & 1 & -1 & -1 \\
1 & -1 & 1 & -1 & 0 & 0 & -1 & 1 \\
1 & -1 & 1 & -1 & 0 & 1 & 0 & 1 \\
0 & -1 & 0 & -1 & 1 & 1 & 1 & 2
\end{pmatrix},
\tag{5.15}
$$

which sends the initial $\alpha_1, \ldots, \alpha_8$ homology basis to another which is not only canonical but behaves well with respect to the symmetry group of the curve. As a result they prove that the period matrix can be written as

$$
\tau = \tau_0
\begin{pmatrix}
4 & 1 & -1 & 1 \\
1 & 4 & 1 & -1 \\
-1 & 1 & 4 & 1 \\
1 & -1 & 1 & 4
\end{pmatrix},
$$

where $\tau_0 \cong -0.5 + 0.186676i$ is defined in terms of j-invariants by

$$
j(\tau_0) = -\frac{29^3 \times 5}{2^5},
$$
$$
j(5\tau_0) = -\frac{25}{2}.
$$

Riera and Rodríguez, in [30], swap these two equations. However, we believe this to be a typographical error.

5.5.3 Understanding the isomorphism f

We now turn our attention to the isomorphism, f, mentioned in (5.14). Actually there clearly won't be a single isomorphism since (as was the case on Craig and Dye's curves) if a is an automorphism of H and σ of Craig's representation then $\sigma \circ f \circ a$ will also be an isomorphism from the hyperbolic model H to Craig's representation. Once again we exploit this fact rather than become discouraged by it.

There are two keys to this process. First is the rotation of the entire hyperbolic polygon about its centre by $2\pi/5$ (the automorphism cd above). This automorphism allows us to express all twenty of the polygon's edges in terms of just four, a great simplification of our problem. If we knew the values of f on four edges, and the matrix representing $f_*(cd)$ then

$$f(\text{edge } k+4) = f\left((cd)(\text{edge } k)\right)$$
$$= f_*(cd)f(\text{edge } k),$$

which allows us to compute the values of f on the remaining 16 edges.

Second is a geodesic reflection on the hyperbolic disc, denoted by the dashed lines in Figure 5.2; the line starts at P_3, goes through C to P_1, along edge 1 to P_2 and along edge 3 back to P_3. If we knew how this acted in Craig's model, we would know its fixed points correspond in some manner to edges e_1/e_{14} and e_3/e_{12}, and the marked diameter.

Identifying points P_1, P_2 and P_3 on the Craig representation would then complete the picture by dividing this fixed line up into just the intervals needed to draw homology paths around known branch points.

Starting with the central rotation cd on the hyperbolic model and some isomorphism f to Craig's representation, since all order 5 elements of S_5 are conjugate there is a Craig-automorphism $\widetilde{\sigma} \in S_5$ such that

$$\widetilde{\sigma} f_*(cd)\widetilde{\sigma}^{-1} = Z^k,$$

where $k \in \{0, \ldots, 4\}$ and

$$Z : [x, y, z] \mapsto [\zeta x, \zeta^2 y, z].$$

We are being flexible about which power of Z occurs here because later choices (specifically rotations about R in Figure 5.2) will disrupt any decision we make at this stage. But then

$$(\sigma \circ f)_*(cd) = \sigma_*(f_*(cd))$$
$$= \sigma f_*(cd)\sigma^{-1}$$
$$= Z^k.$$

So the isomorphism $\sigma \circ f$ from the hyperbolic model to Craig's sends cd to Z^k.

Now consider a rotation about R in Figure 5.2 which cyclically permutes the fixed points of cd. The fixed points on the hyperbolic side are C, P_1, P_2, P_3 and on the Craig side $[0, 1, 0], [0, 0, 1], [1, 0, 0]_1, [1, 0, 0]_2$. Let integers i and n be defined by the equations

$$P_i = (\sigma \circ f)^{-1}([0, 0, 1]), \qquad\qquad R^n(P_i) = C.$$

Then

$$(\sigma \circ f \circ R^{-n})(C) = (\sigma \circ f \circ R^{-n})(R^n(P_i))$$
$$= (\sigma \circ f)(P_i)$$
$$= [0, 0, 1],$$

and further

$$(\sigma \circ f \circ R^{-n})_*(cd) = (\sigma \circ f)_* \left(R_*^{-n}(cd) \right)$$
$$= (\sigma \circ f)_* \left((cd)^j \right)$$
$$= Z^{jk}$$
$$= Z^m,$$

for some integers j and more importantly m. Since we haven't fixed the power of Z up to now this means that $\sigma \circ f \circ R^{-n}$ serves our purposes just as well as $\sigma \circ f$ did.

We used most freedom to constrain the interaction between f, Z and C, but we actually still have the ability to apply a central rotation if it would help since that would disrupt neither of the above properties.

Consider complex conjugation in Craig's model. It is a symmetry that reverses orientation (and so not part of the S_5 symmetry group). It fixes an entire line (the real axis) including the fixed points of Z. In the hyperbolic picture this means it must be a reflection about some diameter. We use our final remaining freedom to demand that it is reflection about the dashed diameter in Figure 5.2, i.e. that the real axis in Craig's model corresponds to these dashed edges (and diameter).

We now have two tasks remaining:

- Find out what P_1, P_2 and P_3 become in Craig's model so we can describe edge 1/14 as the real path from P_2 to P_1 and edge 3/12 as the real path from P_2 to P_3.

- Find out what power of Z the central rotation of $2\pi/5$ becomes so we can describe (for example) edge 4/15 as Z^k applied to the real path from P_2 to P_3.

The second task is actually easier to accomplish at this stage. Consider the structure near $[0, 0, 1]$ (which we demanded was the centre of the polygon, C, hyperbolically); there are three sheets coming together at this branch so unwrapping it will effectively divide angles by 3. Mathematically this means that any set of manifold coordinates $\phi : C \to \mathbb{C}$ centred on $[0, 0, 1]$ will satisfy

$$\phi([x, y, 1])^3 = \alpha x + O(x^2).$$

In these coordinates, since $[0, 0, 1]$ is a fixed point $Z : [x, y, z] \mapsto [\zeta x, \zeta^2 y, z]$ acts locally as a rotation. Denoting this action in the coordinates ϕ by Z_ϕ this means

$$Z_\phi(t) = \beta t + O(t^2),$$

where β is characteristic of Z and independent of ϕ. Now, on the one hand

$$\begin{aligned}
Z_\phi\left(\phi([x, y, 1])\right)^3 &= \phi\left(Z([x, y, 1])\right)^3 \\
&= \phi([\zeta x, \zeta^2 y, 1])^3 \\
&= \alpha \zeta x + O(x^2),
\end{aligned}$$

but also

$$\begin{aligned}
Z_\phi\left(\phi([x, y, 1])\right)^3 &= \left(\beta \phi([x, y, 1]) + O(\phi^2)\right)^3 \\
&= \beta^3 \phi([x, y, 1])^3 + O(\phi^4) \\
&= \beta^3 \alpha x + O(x^2).
\end{aligned}$$

So $\beta^3 = \zeta$, or

$$\beta = \exp\left(\frac{2\pi i}{15} + \frac{2\pi i k}{3}\right)$$

for some $k \in \{0, 1, 2\}$. But since Z has order 5 we also know that $\beta^5 = 1$, which in terms of k means that

$$\frac{2\pi i}{3} + \frac{10\pi i k}{3} = \frac{2\pi i}{3}(1 + 5k) \in 2\pi i \mathbb{Z},$$

or $\beta = \exp(\frac{4\pi i}{5})$ and at last we can conclude that Z corresponds to a rotation of $2 \times \frac{2\pi i}{5}$ about C in the hyperbolic model.

Intuitively we have unwrapped the three sheets coming together at $[0,0,1]$ to obtain Figure 5.4 in x. We know that Z sends (say) $[\epsilon, y, 1]$ to $[\zeta\epsilon, y', 1]$ on some sheet y', which makes it one of the labelled destinations. But only one of these gives an order 5 transformation so we know Z completely.

Figure 5.4: Intuitive action of Z near $[0,0,1]$.

Using this information together with our knowledge that complex conjugation in Craig's model is a dashed reflection in Figure 5.2 allows us to deduce the outline structure in Figure 5.5. The dots are the branch-points of Craig's model and the grey lines are the images of the hyperbolic polygon's edges under the isomorphism to Craig's model. It remains to establish which parts (and sheets) of each spoke in Figure 5.5 correspond to which hyperbolic edges (for example, does edge $1/14$ correspond to $x > 0$ or $x < 0$, and what about y?)

Similar analysis of the other fixed points of Z will allow us to actually identify the remaining P_i. We first discover

- Near $[0,1,0]$, Z is a rotation of $3\left(\frac{2\pi}{5}\right)$.

- Near $[1,0,0]_1 \sim [1, t^4, t]$, Z is a rotation of $4\left(\frac{2\pi}{5}\right)$.

- Near $[1,0,0]_2 \sim [1, t, t^4]$, Z is a rotation of $\frac{2\pi}{5}$.

But hyperbolically, it is easy to see that a rotation of $2\left(\frac{2\pi}{5}\right)$ about C (which Z is) is the same as one of $4\left(\frac{2\pi}{5}\right)$ about P_1, $3\left(\frac{2\pi}{5}\right)$ about P_2 or $\frac{2\pi}{5}$ about P_3 so we can deduce
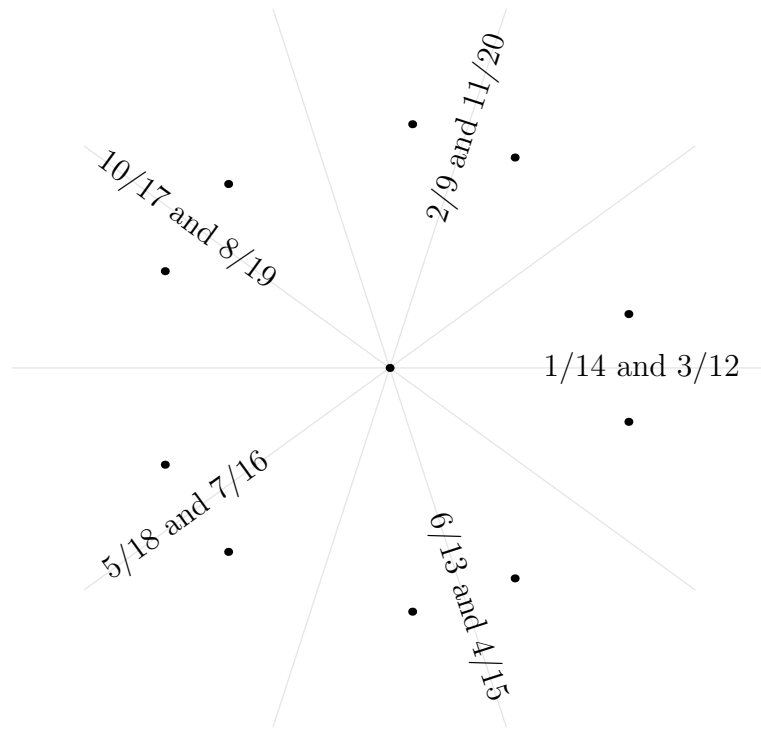
Figure 5.5: Hyperbolic polygonal edges in Craig's model

that $[0, 1, 0] \leftrightarrow P_2$, $[1, 0, 0]_1 \sim [1, t^4, t] \leftrightarrow P_1$ and $[1, 0, 0]_2 \sim [1, t, t^4] \leftrightarrow P_3$.

Therefore, edge $1/14$ corresponds to the real path from $[0, 1, 0]$ to $[1, 0, 0]_1 \sim [1, t^4, t]$; referring to Figure 5.1 we see that this is the path where y starts out large and positive near $x = 0$ (and remains positive). Edge 3 corresponds to the real path from $[0, 1, 0]$ to $[1, 0, 0]_2 \sim [1, t, t^4]$ which turns out to be the one starting out large and negative near $x = 0$ (and remaining negative).

The remaining paths (y small near $x = 0$) correspond to the diameter of the hyperbolic model and have no large role to play in describing the homology basis.

Other edges can now be obtained by applying a rotation of $2\pi/5$ on the hyperbolic side and Z^3 on the Craig side. The results are in Table 5.1.

$1/14$	$[\mathbb{R}_+, \mathbb{R}_+, 1]$	$2/9$	$[\zeta\mathbb{R}_+, \zeta^2\mathbb{R}_+, 1]$
$3/12$	$[\mathbb{R}_+, \mathbb{R}_-, 1]$	$4/15$	$[\zeta^4\mathbb{R}_+, \zeta^3\mathbb{R}_-, 1]$
$5/18$	$[\zeta^3\mathbb{R}_+, \zeta\mathbb{R}_+, 1]$	$6/13$	$[\zeta^4\mathbb{R}_+, \zeta^3\mathbb{R}_+, 1]$
$7/16$	$[\zeta^3\mathbb{R}_+, \zeta\mathbb{R}_-, 1]$	$8/19$	$[\zeta^2\mathbb{R}_+, \zeta^4\mathbb{R}_-, 1]$
$10/17$	$[\zeta^2\mathbb{R}_+, \zeta^4\mathbb{R}_+, 1]$	$11/20$	$[\zeta\mathbb{R}_+, \zeta^2\mathbb{R}_-, 1]$

Table 5.1: Values for $[x, y, 1]$ on hyperbolic edges

5.5.4   Riera and Rodríguez basis algebraically

We are now in a position to express the Riera and Rodríguez basis on this branched
cover. Recall that Riera and Rodríguez used

$$\alpha_i = (2i - 1) + (2i)$$

as a prescription on which edges to traverse in the hyperbolic model.

This becomes a specification to look up the relevant edges in Table 5.1, and construct
a path that has its main component in the specified regions (circling $x = 0$ and outside
all finite branch points enough times to reach the correct sheets). In fact, just like
Riera and Rodríguez we only need to construct α_1 and α_2 and then repeatedly apply
$(x, y) \mapsto (\zeta x, \zeta^2 x)$ to obtain the rest.

To be explicit and referring to Table 5.1, α_1 must go out along $x > 0$ with $y \gg 0$ near
0, loop around infinity until it can come back in to $x = 0$ along a ray with $\arg x = \frac{2\pi}{5}$
and $\arg y = \frac{4\pi}{5}$ before looping around 0 until it can join up with the beginning again. A
path conforming to this description is shown in Figure 5.6.

Similarly α_2 goes out along $x > 0$ with $y < 0$, loops and comes back with argument
of x as $-2\pi/5$ and argument of y as $6\pi/5$; it is also depicted in Figure 5.6.

Figure 5.6: α_1 and α_2 homology cycles in Craig's branched cover. Graphs of subsequent
α_i are rotations of these by $\frac{2\pi i}{5}$.

These are read into `extcurves` and converted into a full basis with the commands

```
> curve, hom, names := read_pic("homology.pic"):

> zeta := exp(2*Pi*I/5):

> trans := (x,y) -> [zeta^3*x,zeta*y]:

> for i from 1 to 3 do

    hom := [op(hom),
```

```
            transform_extpath(curve, hom[-2], trans),

            transform_extpath(curve, hom[-1], trans)];

  od:
```

An immediate check to this calculation is provided by calculating the intersection matrix of this constructed basis

```
> Matrix(8, (i,j) -> isect(curve, hom[i], hom[j]));
```

produces (with considerably less work and chance of error) precisely the matrix claimed by Riera and Rodríguez

$$\begin{pmatrix} 0 & 1 & -1 & 1 & -1 & 0 & 1 & -1 \\ -1 & 0 & 1 & -1 & 1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 1 & -1 & 1 & -1 & 0 \\ -1 & 1 & -1 & 0 & 1 & -1 & 1 & 0 \\ 1 & -1 & 1 & -1 & 0 & 1 & -1 & 1 \\ 0 & 0 & -1 & 1 & -1 & 0 & 1 & -1 \\ -1 & 0 & 1 & -1 & 1 & -1 & 0 & 1 \\ 1 & 0 & 0 & 0 & -1 & 1 & -1 & 0 \end{pmatrix}.$$

Finally we can calculate the period matrix and apply the transformation from (5.15) to reproduce the claimed result that

$$\tau = \tau_0 \begin{pmatrix} 4 & 1 & -1 & 1 \\ 1 & 4 & 1 & -1 \\ -1 & 1 & 4 & 1 \\ 1 & -1 & 1 & 4 \end{pmatrix}.$$

5.6 Vector of Riemann constants

We will now calculate the vector of Riemann constants for this surface in two slightly different ways. First we will apply a purely numeric approach suitable for any calculations needed regarding, for example, Θ-functions. Nevertheless this approach will allow us to guess the algebraic form of the vector, which itself could be very useful for enhancing the precision of any computations undertaken (if we can find a more precise vector, then the dubious methods used are, after all, not that important).

After that we will adopt a hybrid algebraic-numeric approach capable of actually deriving an algebraic form of the vector. This will naturally be harder work, but the results correspondingly more useful. This second version is very similar to that employed for Klein's curve, except the numerics involved will be more analytic than algebraic in nature, with corresponding increase in cost. Fortunately the problem is small enough that this is not important.

5.6.1   Purely numeric Riemann constants

Numerically calculating the vector of Riemann constants can now proceed in an almost identical manner to Klein's curve in Section 4.9.

1. Find a divisor for a holomorphic differential. We use Craig's representation and the differential
$$\frac{(y - x^2)\mathrm{d}x}{-5xy^4 - 2x^2y + x^4 + 6y^2}.$$

   It has a simple 0 at $[0, 0, 1]$, a double 0 at $[0, 1, 0]$ and a triple 0 at $[1, 0, 0] \sim [1, t, t^4]$ for the required total of $2g - 2 = 6$.

2. Choose a base-point Q and calculate $-2K_Q = A_Q(\mathcal{K}_C)$. We choose $[0, 0, 1]$ as the base and discover

$$-2K_0^{num} \cong \begin{pmatrix} -0.900002173641643699 + 1.30672399048697319\mathrm{i} \\ 1.89999072196101570 + 0.186685232667863965\mathrm{i} \\ 0.599990762181047632 + 1.12006518865340832\mathrm{i} \\ -1.59999547497843220 + 0.373348898781978522\mathrm{i} \end{pmatrix}.$$

3. Pick 4 generic points and add half-periods to $-A_0(\mathcal{K}_C)/2$ until we find the correct vanishing properties with respect to Θ functions. We discover

$$K_0^{num} \cong \begin{pmatrix} -0.299998896336445942 - 0.1866708344\mathrm{i} \\ -2.19999535149310521 + 0.373348534635175533\mathrm{i} \\ -0.299995378646934263 - 0.5600325922\mathrm{i} \\ 1.29999774009064328 - 0.1866744408\mathrm{i} \end{pmatrix}.$$

4. Add periods to make this result as neat as possible, obtaining

$$K_0^{num} \cong \begin{pmatrix} -0.299998890223328429 - 0.1866708287i \\ -0.199995316631002673 + 0.373348521868197392i \\ 0.200004641308583864 + 0.373349724864026844i \\ 0.299997742729032346 - 0.1866744176i \end{pmatrix}.$$

5. Identify the numbers appearing in terms of τ_0. We can see that it is very probable that

$$K_0 = \frac{1}{10} \begin{pmatrix} -3 \\ -2 \\ 2 \\ 3 \end{pmatrix} + \text{Im}\,(\tau_0) \begin{pmatrix} -1 \\ 2 \\ 2 \\ -1 \end{pmatrix} i. \tag{5.16}$$

5.6.2 Justifying the algebraic vector of constants

Our final goal will be to actually prove the guess in (5.16). If we were to follow the full procedure from Section 4.9 we would have to develop a much more complete understanding of

- The action of the automorphisms on holomorphic differentials.

- The action of the automorphisms on homology cycles.

- The structure of the full matrix of periods, Π.

Fortunately the problem is such that we can just touch on each of these issues and still emerge with the proof we desire. Recall that the fundamental way symmetries constrain the vector of Riemann constants is via the equation (expressed in $\mathbb{C}^g$ rather than the Jacobian and assuming the Abel map is based at a fixed point of the automorphism)

$$FK_0 = K_0 + \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{k}, \tag{5.17}$$

where F is the action of some symmetry on the holomorphic differentials, expressed in the basis dual to our chosen canonical homology basis. If we can use this, as at the beginning of Klein's result, to constrain the vector of constants to a finite set of algebraic possibilities then numeric methods will suffice to justify the ansatz made above.

In the Bring case the useful simplification is that everything follows from a deep study of the single (order 5) automorphism given in Craig's representation by

$$\phi : [\tilde{x}, \tilde{y}, \tilde{z}] \mapsto [\zeta^2 \tilde{x}, \zeta^4 \tilde{y}, \tilde{z}].$$

The first and easiest calculation is deriving its action on the differentials. We fix the ordered basis

$$\omega_1 = \frac{(\tilde{y}^3 - \tilde{x})\mathrm{d}\tilde{x}}{\partial_y \mathcal{C}(\tilde{x}, \tilde{y}, 1)}, \qquad\qquad \omega_2 = \frac{\tilde{y}(\tilde{x}^2 - \tilde{y})\mathrm{d}\tilde{x}}{\partial_y \mathcal{C}(\tilde{x}, \tilde{y}, 1)},$$

$$\omega_3 = \frac{(\tilde{y} - \tilde{x}^2)\mathrm{d}\tilde{x}}{\partial_y \mathcal{C}(\tilde{x}, \tilde{y}, 1)}, \qquad\qquad \omega_4 = \frac{(\tilde{y}^2 \tilde{x} - 1)\mathrm{d}\tilde{x}}{\partial_y \mathcal{C}(\tilde{x}, \tilde{y}, 1)}.$$

It is easy to check that

$$\phi^*(\omega_1) = \zeta \omega_1, \qquad\qquad \phi^*(\omega_2) = \zeta^2 \omega_2,$$

$$\phi^*(\omega_3) = \zeta^3 \omega_3, \qquad\qquad \phi^*(\omega_4) = \zeta^4 \omega_4,$$

the remarkable (and useful) fact being that there is no invariant differential. In terms of matrices this takes the form $\phi^*(\omega_i) = \hat{F}_{ij}\omega_j$ where

$$\hat{F} = \begin{pmatrix} \zeta & 0 & 0 & 0 \\ 0 & \zeta^2 & 0 & 0 \\ 0 & 0 & \zeta^3 & 0 \\ 0 & 0 & 0 & \zeta^4 \end{pmatrix}. \tag{5.18}$$

To apply (5.17), we will need to express this action in the $\mathfrak{a}$-normalised basis, which requires knowledge of $\mathcal{A}_{ij} = \int_{\mathfrak{a}_j} \omega_i$. We can use ϕ itself again to obtain this knowledge. The first step is deriving its action on homology cycles: $\phi_*(\gamma_i)$. With `extcurves` on hand this is a simple computational matter, complicated only slightly by the noncanonical

nature of the paths we obtained in Section 5.5.4. We obtain $\phi_*(\gamma_i) = M_{ij}\gamma_j$ where

$$
M = \begin{pmatrix}
0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
-1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\
0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -1 & 1 & -1 & 1 \\
0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -1 & 0
\end{pmatrix}. \tag{5.19}
$$

5.6.3 The Abel map of the canonical divisor

Since the action of ϕ on the differentials has no invariants we know, as in the Klein case, that $F - 1$ is invertible and

$$
-2K_0 = (F - 1)^{-1} \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{k},
$$

up to some period. Again there are only finitely many essentially unique values for $\mathbf{k}$, the rest simply altering $-2K_0$ by a period. In this case $\mathbf{k}$ and $\mathbf{k} + \mathbf{a}$ give results differing by a period if

$$
(F - 1)^{-1} \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{a} = \begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{b}
$$

has a solution in $\mathbf{b}$ (again $\mathbf{a}, \mathbf{b} \in \mathbb{Z}^8$), or, rearranging and using the fact that $F \begin{pmatrix} 1 & \tau \end{pmatrix} = \begin{pmatrix} 1 & \tau \end{pmatrix} M$,

$$
\begin{pmatrix} 1 & \tau \end{pmatrix} \mathbf{a} = \begin{pmatrix} 1 & \tau \end{pmatrix} (M - 1)\mathbf{b},
$$

which is equivalent to

$$
\mathbf{a} = (M - 1)\mathbf{b}.
$$

Calculating the Smith normal form of $M - 1$ gives us unimodular matrices U, V such that

$$
M - 1 = U \operatorname{diag}(1, 1, 1, 1, 1, 1, 5, 5)V
$$

or

$$U^{-1}\mathbf{a} = \mathrm{diag}(1,1,1,1,1,1,5,5)V\mathbf{b},$$

which has a solution provided the final two entries of $U^{-1}\mathbf{a}$ are multiples of 5. That is, $(F-1)^{-1}\begin{pmatrix} 1 & \tau \end{pmatrix}\mathbf{a}$ is a period if and only if

$$-a_5 + a_6 - a_7 - 4a_8 \equiv 0 \pmod 5,$$

$$-a_1 + 2a_2 - 3a_3 - a_4 - 11a_5 + 6a_6 - a_7 - 34a_8 \equiv 0 \pmod 5.$$

This gives 25 possible unique candidates for $-2K_0$; we can vary k_i arbitrarily (by adding an appropriate $\mathbf{a}$) without essentially changing $-2K_0$ for (say) $i = 2,3,4,6,7,8$ but then k_1 and k_5 are fixed. Explicitly, every $-2K_0$ is equivalent to one generated by

$$\mathbf{k} = \begin{pmatrix} k_1 & 0 & 0 & 0 & k_5 & 0 & 0 & 0 \end{pmatrix}.$$

5.6.4   Action of ϕ on dual differential basis

The final step before we can actually apply numerical methods is determining what the matrix F, acting on the differentials, actually is. This needs knowledge of the matrix of periods.

Applying the two equations (5.19) and (5.18) to the standard constraint on period matrices $\hat{F}\begin{pmatrix} \mathcal{A} & \mathcal{B} \end{pmatrix} = \begin{pmatrix} \mathcal{A} & \mathcal{B} \end{pmatrix} M$, we obtain the following information about $\mathcal{A}_{ij} = \int_{\mathfrak{a}_j} \omega_i$

$$\mathcal{A} = \begin{pmatrix} a_1 & 0 & 0 & 0 \\ 0 & a_2 & 0 & 0 \\ 0 & 0 & a_3 & 0 \\ 0 & 0 & 0 & a_4 \end{pmatrix} \begin{pmatrix} 1 & -1-\zeta^4 & 1+\zeta^4+\zeta^3 & \zeta \\ 1 & -1-\zeta^3 & 1+\zeta^3+\zeta & \zeta^2 \\ 1 & -1-\zeta^2 & 1+\zeta^2+\zeta^4 & \zeta^3 \\ 1 & -1-\zeta & 1+\zeta+\zeta^2 & \zeta^4 \end{pmatrix},$$

for some unknown a_i, which since $\hat{F}$ is diagonal are fortunately irrelevant. It follows that the representation of ϕ on the basis of differentials dual to our chosen canonical

homology basis is $F = \mathcal{A}^{-1}\hat{F}\mathcal{A}$, which can be simplified to

$$
F = \begin{pmatrix}
0 & -1 & 0 & 0 \\
0 & 0 & -1 & 0 \\
0 & 0 & 0 & -1 \\
1 & -1 & 1 & -1
\end{pmatrix}.
$$

We can finally apply numerical methods to obtain results that are nonetheless algebraic. Comparing each possible $k_{\{1,5\}}$ against K_0^{num} we discover that only $k_1 = k_5 = 3$ gives a result differing by a period. That is,

$$
-2K_0 = \frac{1}{5}\begin{pmatrix} -12 \\ -3 \\ 3 \\ -3 \end{pmatrix} + \tau_0 \begin{pmatrix} -6 \\ -6 \\ 3 \\ 0 \end{pmatrix},
$$

with the explicit relation to the numerically derived version given earlier

$$
-2K_0 = -2K_0^{num} + \begin{pmatrix} \mathbf{1} & \tau \end{pmatrix}\begin{pmatrix} 4 & -1 & -3 & 1 & 1 & 1 & -2 & 1 \end{pmatrix}^T.
$$

Now that we know $-2K_0$ the argument proceeds in an identical manner to the numeric case and we obviously obtain (5.16) again, with more justification this time.

5.7 Summary

In this chapter, we have considered three representations of Bring's curve which is known to be the unique curve of genus 4 with S_5 as an automorphism group, just as Klein's is the unique genus 3 curve with maximal symmetry.

I believe the two sextic representations were known to be equivalent previously: Dye certainly knows his curve is equivalent to Bring's (see the introduction to [13]). Craig does not explicitly state this result but suggests in [9] that he knows that A_5 is only part of the group which would result in equivalence to Bring's.

However, here we have produced a novel and explicit relationship between these two equations. This mapping is as simple as could possibly be expected and will allow future work on one curve to be transferred to the other, no matter how concrete the result seems.

The third hyperbolic representation used by Riera and Rodruíguez in [30] was different, but we have created a strong enough bridge between it and the algebraic equations that we could transfer their result on a homology basis to the other setting.

Finally, we presented two methods of calculating a vector of Riemann constants for this curve. In this case we had enough information that both were applicable. On other curves, only one or the other of these techniques may be practical.

CHAPTER

6

GENUS 2 CURVES COVERING GENUS 1

6.1   Introduction

6.1.1   Curves covering curves

Given a curve $\hat{\mathcal{C}}$ of genus $\hat{g}$ it is often useful to know curves of lower genus covered. For example this can allow us to write the Θ-functions of $\hat{\mathcal{C}}$ in terms of the (simpler) Θ-functions on the curve with lower genus, as described by Martens in [26].

Suppose $\pi : \hat{\mathcal{C}} \to \mathcal{C}$ is such a cover onto $\mathcal{C}$ of genus g. The Poincaré reducibility criterion gives a relation between the period matrices of the two curves, and can actually be used in reverse when the genus of $\hat{\mathcal{C}}$ is 2 to identify such covers. We derive the equation as follows. Suppose

$$\pi^*(\omega_i) = \hat{\omega}_j \lambda_{ji},$$
$$\pi_*(\hat{\gamma}_i) = M_{ij} \gamma_j,$$
$$\tag{6.1}$$

where γs represent a choice of homology bases, and ωs of one-forms; the matrices of periods will be denoted $\Pi_{ij} := \int_{\gamma_i} \omega_j$. Both λ and M have maximal rank:

- π_* is surjective: suppose we have an arbitrary cycle $\delta : [0, 1] \to \mathcal{C}$. We can choose a preimage of $\delta(0)$ and extend this by continuity to obtain $\pi_*^{-1}(\delta) : [0, 1] \to \hat{\mathcal{C}}$. (If δ passed through any branch points of π then further arbitrary choices would be required, but continuity could be maintained). The only barrier to $\pi_*^{-1}(\delta)$ being a preimage of δ under π_* (i.e. in $H_1(\hat{\mathcal{C}}, \mathbb{Z})$) is that it may not be a closed cycle. But this can be corrected by adjoining a path χ on $\hat{\mathcal{C}}$ whose push-forward, $\pi_*(\chi)$, is homotopic to 0 but which closes $\pi_*^{-1}(\delta)$. Then we would have $\pi_*^{-1}(\delta) + \chi$ a closed path satisfying

$$\pi_*(\pi_*^{-1}(\delta) + \chi) = \pi_*(\pi_*^{-1}(\delta)) + \pi_*(\chi) = \delta + 0.$$

  Such a path will always exist: the monodromy based at $\delta(0)$ can be calculated around all branch points of the cover and must allow access to all sheets (since $\hat{\mathcal{C}}$ is connected). Hence, on $\mathcal{C}$, paths that go out from $\delta(0)$ to any branch point, perform a small loop and return along the same track are homologous to zero and allow us to change to any sheet. For example Figure 6.1 shows a possible scenario where a genus 2 surface covers a torus with two branch points.

- π^* is injective since it is the adjoint of the surjective π_*.

Figure 6.1: Adding a path χ (shown dotted) to close a pullback.

Then

$$
\begin{aligned}
M_{ij}\Pi_{jk} = M_{ij}\int_{\gamma_j}\omega_k &= \int_{M_{ij}\gamma_j}\omega_k \\
&= \int_{\pi_*(\hat{\gamma}_i)}\omega_k = \int_{\hat{\gamma}_i}\pi^*(\omega_k) \\
&= \int_{\hat{\gamma}_i}\hat{\omega}_j\lambda_{jk} = \lambda_{jk}\int_{\hat{\gamma}_i}\hat{\omega}_j = \hat{\Pi}_{ij}\lambda_{jk}.
\end{aligned}
$$

That is, if $\hat{\mathcal{C}}$ covers another curve $\mathcal{C}$ then there exists

- an integer matrix M of size $2\hat{g} \times 2g$ with maximal rank,

- a complex matrix λ of size $\hat{g} \times g$,

such that

$$
M\Pi = \hat{\Pi}\lambda. \tag{6.2}
$$

For the $\hat{g} = 2, g = 1$ case Murabayashi proves the converse in [28]. With slight alterations to avoid introducing extra notation, his result is stated

Theorem 11 (Murabayashi). *J_C has a period matrix of the form*

$$
\begin{pmatrix}
z_1 & \frac{1}{k} & 1 & 0 \\
\frac{1}{k} & z_2 & 0 & 1
\end{pmatrix}, \; z_1, z_2 \in \mathfrak{h}
$$

if and only if C has a maximal map ψ of degree k to an elliptic curve X.

Here, J_C is the Jacobian of a Riemann surface C of genus 2 and $\mathfrak{h}$ is the upper-half complex plane.

As we will see later, if there is one solution to equation (6.2) with matrices having the required properties then there are more, and some of these will give $\hat{\Pi}$ precisely the form stated in Murabayashi's result (up to sign conventions). This tells us that there is some covering map whenever we can solve (6.2).

6.1.2   Simplification when $\hat{g} = 2$, $g = 1$

Here we specialise to the case where a genus 2 curve covers one with genus 1. I am indebted to my supervisor Harry Braden for most of the work on this simplification. (6.2) is equivalent to the single matrix equation (actually valid whenever it makes sense, i.e. $\hat{g} = 2g$) given in block form by

$$\begin{pmatrix} \hat{\Pi} & M \end{pmatrix} \begin{pmatrix} \lambda \\ -\Pi \end{pmatrix} = 0.$$

Now, the left-hand operand is a 4×4 matrix, so a necessary condition for a nontrivial solution to exist in M, λ, Π is that

$$\det \begin{pmatrix} \hat{\Pi} & M \end{pmatrix} = 0.$$

If we can find all solutions to this equation then since we're working with genus 2 and 1, the only barrier to solving (6.2) itself is whether a *bona fide* period matrix is induced from the eigenvector on the right; this is merely a matter of checking positive definiteness of the imaginary part of Π_{21}/Π_{11}.

The next thing to note about this equation is that if we have a solution M and

$$\hat{\Pi} = \begin{pmatrix} \mathcal{A} \\ \mathcal{B} \end{pmatrix}$$

then

$$\begin{aligned}
0 &= \det \begin{pmatrix} \mathcal{A} & M_1 \\ \mathcal{B} & M_2 \end{pmatrix} \\
&= \det \begin{pmatrix} \mathcal{A} & M_1 \\ \mathcal{B} & M_2 \end{pmatrix} \begin{pmatrix} \mathcal{A}^{-1} & 0 \\ 0 & \mathbf{1} \end{pmatrix} \\
&= \det \begin{pmatrix} 1 & M_1 \\ \hat{\tau} & M_2 \end{pmatrix}.
\end{aligned}$$

So we search for solutions to this latter equation

$$
0 = \det \begin{pmatrix} 1 & 0 & m_{11} & m_{12} \\ 0 & 1 & m_{21} & m_{22} \\ \hat{\tau}_{11} & \hat{\tau}_{12} & m_{31} & m_{32} \\ \hat{\tau}_{12} & \hat{\tau}_{22} & m_{41} & m_{42} \end{pmatrix}.
$$

Expanding this determinant we obtain

$$
0 = \begin{vmatrix} m_{31} & m_{32} \\ m_{41} & m_{42} \end{vmatrix} - \hat{\tau}_{11} \begin{vmatrix} m_{11} & m_{12} \\ m_{41} & m_{42} \end{vmatrix} + \hat{\tau}_{12} \left(\begin{vmatrix} m_{11} & m_{12} \\ m_{31} & m_{32} \end{vmatrix} - \begin{vmatrix} m_{21} & m_{22} \\ m_{41} & m_{42} \end{vmatrix} \right)
$$
$$
+ \hat{\tau}_{22} \begin{vmatrix} m_{21} & m_{22} \\ m_{31} & m_{32} \end{vmatrix} + (\hat{\tau}_{11}\hat{\tau}_{22} - \hat{\tau}_{12}^2) \begin{vmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{vmatrix}, \quad (6.3)
$$

which must be solved for the m_{ij} given known $\hat{\tau}_{ij}$.

Each of these minor determinants must be an integer first and foremost, thus in our search for a solution the first step is to define

$$
q_1 = \begin{vmatrix} m_{31} & m_{32} \\ m_{41} & m_{42} \end{vmatrix}, \qquad q_2 = - \begin{vmatrix} m_{11} & m_{12} \\ m_{41} & m_{42} \end{vmatrix}, \qquad q_3 = \begin{vmatrix} m_{11} & m_{12} \\ m_{31} & m_{32} \end{vmatrix} - \begin{vmatrix} m_{21} & m_{22} \\ m_{41} & m_{42} \end{vmatrix},
$$
$$
q_4 = \begin{vmatrix} m_{21} & m_{22} \\ m_{31} & m_{32} \end{vmatrix}, \qquad q_5 = \begin{vmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{vmatrix},
$$

and search for solutions to the reduced equation

$$
q_1 + q_2 \hat{\tau}_{11} + q_3 \hat{\tau}_{12} + q_4 \hat{\tau}_{22} + q_5 (\hat{\tau}_{11}\hat{\tau}_{22} - \hat{\tau}_{12}^2) = 0
$$

in integers q_i. We can then take each of these solutions and search for corresponding m_{ij}, which must exist by Murabayashi's results in [28]. A second constraint on the q_i is provided by the demand, for an order h cover, that

$$
hJ = M^T \hat{J} M, \tag{6.4}
$$

where J and $\hat{J}$ are the 2×2 and 4×4 (respectively) standard sympelctic matrices defined earlier.

Equation (6.4) can be derived in two stages by considering a suitably defined pseudo-

inverse for π_* acting on homology cycles. Let $\pi^{-1}(\gamma)$ be defined by taking all possible preimages of γ and joining the endpoints to form a closed path. That is

$$\pi^{-1}(\gamma) = \hat{\alpha}_1 + \cdots + \hat{\alpha}_h$$

where each of $\hat{\alpha}_i$ are (not usually closed) paths in $\hat{\mathcal{C}}$ pushing forwards to γ. Now we define the matrix Q to be the action of π^{-1} on a canonical homology basis

$$\pi^{-1}(\gamma_i) = Q_{ij}\hat{\gamma}_j.$$

We can calculate Q as follows

$$\begin{aligned}
Q_{ik}\hat{J}_{kj} &= Q_{ik}\langle \hat{\gamma}_k, \hat{\gamma}_j \rangle \\
&= \langle \pi^{-1}(\gamma_i), \hat{\gamma}_j \rangle,
\end{aligned}$$

which, by considering each intersection point on $\mathcal{C}$ and $\hat{\mathcal{C}}$ is

$$\begin{aligned}
&= \langle \gamma_i, \pi_*(\hat{\gamma}_j) \rangle \\
&= \langle \gamma_i, \gamma_k \rangle M_{jk} \\
&= J_{ik}M_{jk}.
\end{aligned}$$

This assumes transverse intersections occurring only at regular points of π, but these conditions can always be arranged by deforming the paths involved. The result is

$$Q = -JM^T \hat{J}.$$

Finally we eliminate Q again by noting that

$$\pi_* \circ \pi^{-1} = h\mathbf{1},$$

and calculating as follows

$$
\begin{aligned}
h J_{ij} &= h \langle \gamma_i, \gamma_j \rangle \\
&= \langle (\pi_* \circ \pi^{-1})(\gamma_i), \gamma_j \rangle \\
&= \langle \pi_*(Q_{ik} \hat{\gamma}_k), \gamma_j \rangle \\
&= Q_{ik} \langle \pi_*(\hat{\gamma}_k), \gamma_j \rangle \\
&= Q_{ik} M_{kl} \langle \gamma_l, \gamma_j \rangle \\
&= -J_{im} M_{mn}^T \hat{J}_{nk} M_{kl} J_{lj}.
\end{aligned}
$$

Or, upon multiplying left and right by J,

$$
h J = M^T \hat{J} M,
$$

as required. When written out for our genus and expressed in terms of the q_i this equation reduces to

$$
q_3^2 + 4(q_1 q_5 - q_2 q_4) = h^2,
$$

and will be the key to showing that only finitely many solutions exist for a given h.

6.1.3 Running example (tetrahedral monopole)

To illustrate typical results obtained while constructing this solution we will use a particularly symmetric curve which we know *a priori* to cover some elliptic curves. In [3], Braden and Enolski consider a hyperelliptic curve of particular relevance to monopoles. This curve is given by the equation

$$
y^2 - (x^3 + 5\sqrt{2})^2 - 4 = 0,
$$

and they prove that it has period matrix

$$
\hat{\tau} = \frac{1}{78} \begin{pmatrix} -20\rho + 47 & -24\rho + 33 \\ -24\rho + 33 & -60\rho - 15 \end{pmatrix},
$$

where $\rho = e^{2\pi i/3}$, a solution to $\rho^2 + \rho + 1 = 0$.

6.2   Solve for q_i

Presented with the period matrix for a genus 2 curve, we must first solve the equations

$$q_1 + q_2\hat{\tau}_{11} + q_3\hat{\tau}_{12} + q_4\hat{\tau}_{22} + q_5(\hat{\tau}_{11}\hat{\tau}_{22} - \hat{\tau}_{12}^2) = 0, \tag{6.5}$$

$$q_3^2 + 4(q_1q_5 - q_2q_4) = h^2. \tag{6.6}$$

The difficult equation here is (6.6) which is quadratic in the unknowns q_i. For now we will leave this and concentrate on (6.5).

6.2.1   Converting linear part to system of Diophantine equations

Our first step is to convert (6.5) into a linear system of Diophantine equations. Knowing the matrix $\hat{\tau}$, we can place the entries in an appropriate extension field $K : \mathbb{Q}$. We can then express (6.5) in terms of a basis for K over $\mathbb{Q}$ and use linear independence to obtain an equivalent set of linear equations with coefficients in $\mathbb{Q}$. Finally we clear denominators to give proper Diophantine equations.

Typically, $\hat{\tau}$ lives in a fairly uncomplicated algebraic extension to $\mathbb{Q}$ and the analysis is easy. In the running example, we have

$$0 = \rho^2 + \rho + 1, \tag{6.7}$$

$$\hat{\tau} = \frac{1}{78}\begin{pmatrix} -20\rho + 47 & -24\rho + 33 \\ -24\rho + 33 & -60\rho - 15 \end{pmatrix}.$$

The polynomial for ρ is irreducible over $\mathbb{Q}$ so we can work in the field $\mathbb{Q}(\rho)$ and in light of (6.7), the linear equation (6.5) becomes

$$\frac{1}{78}(78q_1 + 47q_2 + 33q_3 - 15q_4 - 31q_5) + \frac{2}{39}(-5q_2 - 6q_3 - 15q_4 - 5q_5)\rho = 0,$$

which vanishes if and only if both parts vanish separately. Simultaneously clearing fractional parts we obtain

$$78q_1 + 47q_2 + 33q_3 - 15q_4 - 31q_5 = 0,$$

$$-5q_2 - 6q_3 - 15q_4 - 5q_5 = 0,$$

which can be expressed as

$$A\mathbf{q} := \begin{pmatrix} 78 & 47 & 33 & -15 & -31 \\ 0 & -5 & -6 & -15 & -5 \end{pmatrix} \mathbf{q} = 0. \tag{6.8}$$

Note that in this case an equivalent set of equations would have been obtained by splitting (6.5) into real and imaginary parts. In general finer information is obtained by the use of field extensions as above but, particularly in cases where we only know $\hat{\tau}$ numerically, the real/imaginary split can be useful. We return to this topic in Section 6.7.1 with an explicit example.

6.2.2 Solving the linear Diophantine system

The solution to Diophantine equations of the form $A\mathbf{q} = 0$ is well known. The first step is to put A into Smith normal form using an integer variant of Gaussian elimination, to obtain

$$A = L\Lambda R$$

where L and R are unimodular matrices and Λ (not necessarily square) has the block form

$$\Lambda = \begin{pmatrix} D_n & 0 \\ 0 & 0 \end{pmatrix},$$

$$D_n = \text{diag}(d_1, \ldots, d_n)$$

and $d_i | d_{i+1}$. We then have the equation

$$L\Lambda R\mathbf{q} = 0.$$

As L is invertible over $\mathbb{Z}$ it can be immediately eliminated to yield

$$\Lambda(R\mathbf{q}) = 0,$$

which has general solution

$$\mathbf{q} = R^{-1} \begin{pmatrix} 0_n \\ \star \end{pmatrix}.$$

In the case of (6.8) immediately at hand we find the factorisation

$$A = \begin{pmatrix} 1 & 0 \\ -298 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & -39 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} -23 & 1 & 3 & 2 & 1 \\ 23 & 0 & -12 & -3 & -1 \\ 23 & -1 & 10 & 0 & 0 \\ 3 & 3 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

which leads to the general solution

$$\mathbf{q} = \begin{pmatrix} 3\alpha + 2\beta + \gamma \\ -12\alpha - 3\beta - \gamma \\ 10\alpha \\ \beta \\ \gamma \end{pmatrix} \tag{6.9}$$

as α, β, γ range over $\mathbb{Z}$.

6.2.3 The quadratic equation

Dietmann shows in [12] that there is an algorithm to determine whether a single quadratic diophantine equation (in an arbitrary number of variables) as we have in (6.6) is solvable. The algorithm actually produces a solution if one exists, but it is still of theoretical interest only in this particular case for two reasons:

- It is exponential in the coefficients involved.

- It only guarantees one solution will be found if it exists and says nothing about the character of others.

Fortunately (6.6) is not a generic quadratic equation, particularly once we have extracted some information from (6.5), so other techniques can be applied.

In this section we will describe the example first and then move on to more general theory. Substituting (6.9) into the quadratic equation (6.6) gives

$$100\alpha^2 + 4\gamma^2 + 12\beta^2 + 12\gamma\alpha + 12\gamma\beta + 48\beta\alpha = h^2.$$

The important thing to notice about this is that the left hand side can be written as

$$\mathbf{x}^T Q \mathbf{x} := \begin{pmatrix} \alpha & \beta & \gamma \end{pmatrix} \begin{pmatrix} 100 & 24 & 6 \\ 24 & 12 & 6 \\ 6 & 6 & 4 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix}$$

and is a positive definite quadratic form in α, β, γ. Its eigenvalues, λ_i are roughly 0.2, 9.2 and 106.6. Suppose it is diagonalisable by the orthogonal matrix T, then (quite generally)

$$\begin{aligned} \mathbf{x}^T Q \mathbf{x} &= (T\mathbf{x})^T (TQT^T)(T\mathbf{x}) \\ &= \lambda_1 (T\mathbf{x})_1^2 + \cdots + \lambda_n (T\mathbf{x})_n^2 \\ &\geq \min(\lambda_i)((T\mathbf{x})_1^2 + \cdots + (T\mathbf{x})_n^2) \\ &= \min(\lambda_i)||T\mathbf{x}||^2 \\ &= \min(\lambda_i)||\mathbf{x}||^2. \end{aligned}$$

Hence (since $\min(\lambda_i) > 0$), given an h^2 we only need to search for solutions where

$$||\mathbf{x}||^2 \leq \frac{h^2}{\min(\lambda_i)},$$

or in our case

$$||\mathbf{x}||^2 \leq 6h^2,$$

in particular there are only finitely many solutions in $\mathbb{Z}$ for each h. This is the best we can expect since elliptic curves cover themselves so in general if there is a single solution we would expect solutions for arbitrarily high h obtained by first mapping the genus 2 curve to the elliptic curve and then self-mapping the elliptic curve.

At this stage we have a reasonably practical (polynomial time in h) algorithm to find candidate coverings by a genus 2 curve. When applied, the solutions in Table 6.1 are found after moving back from $\mathbf{x} = (\alpha, \beta, \gamma)$ to the desired $\mathbf{q}$ variables.

Each row has a pair of solutions, obtained by the symmetry of (6.5) and (6.6) under $\mathbf{q} \mapsto -\mathbf{q}$. Notice that most of the solutions for $h = 4$ are simply double those from $h = 2$. This is a manifestation of the self-mapping of the elliptic curves. However, the final pair is distinctive.

h	$\mathbf{q} = (q_1, q_2, q_3, q_4, q_5)$	
2	(0 1 0 -1 2)	(0 -1 0 1 -2)
	(1 -2 0 1 -1)	(-1 2 0 -1 1)
	(1 -1 0 0 1)	(-1 1 0 0 -1)
3	No solutions	
4	(0 2 0 -2 4)	(0 -2 0 2 -4)
	(2 -4 0 2 -2)	(-2 4 0 -2 2)
	(2 -2 0 0 2)	(-2 2 0 0 -2)
	(1 3 -10 5 -6)	(-1 -3 10 -5 6)

Table 6.1: Candidate solutions to (6.5) and (6.6) for small h.

Moving back to the general theory, the question arises of how lucky we were to discover a positive definite quadratic form that only allowed finitely many solutions. Certainly the original equation (6.6) was not definite, and in principle the form of the solution (6.9) could be very different for other curves (perhaps even fewer or more degrees of freedom) so there's no *a priori* reason to expect a definite form to result.

In fact just splitting (6.5) into real and imaginary parts with $q_i \in \mathbb{R}$ always provides enough constraint to guarantee the quadratic part is positive definite:

Theorem 12. *Given a 2×2 period matrix $\hat{\tau}$ (in particular a symmetric matrix with positive definite imaginary part), the quadratic form*

$$q_3^2 + 4(q_1 q_5 - q_2 q_4)$$

is positive definite on the hyperplane defined by the solutions to (6.5).

Proof. Since the q_i are certainly real, we begin by splitting (6.5) into real and imaginary parts. Suppose

$$\hat{\tau}_{ij} = \alpha_{ij} + \mathrm{i}\beta_{ij},$$

and note that since $\hat{\tau}$ is symmetric with positive definite imaginary part we cannot have $\beta_{11} = 0$. Then (6.5) becomes

$$0 = q_1 + q_2\alpha_{11} + q_3\alpha_{12} + q_4\alpha_{22} - q_5(\alpha_{12}^2 - \beta_{12}^2 - \alpha_{11}\alpha_{22} + \beta_{11}\beta_{22}),$$

$$0 = q_2\beta_{11} + q_3\beta_{12} + q_4\beta_{22} - q_5(2\alpha_{12}\beta_{12} - \alpha_{11}\beta_{22} - \beta_{11}\alpha_{22}).$$

Since $\beta_{11} \neq 0$, this has solution

$$q_2 = -\frac{q_3\beta_{12} + q_4\beta_{22} - q_5(2\alpha_{12}\beta_{12} - \alpha_{11}\beta_{22} - \beta_{11}\alpha_{22})}{\beta_{11}},$$

$$q_1 = \frac{\alpha_{11}[q_3\beta_{12} + q_4\beta_{22} - q_5(2\alpha_{12}\beta_{12} - \alpha_{11}\beta_{22} - \beta_{11}\alpha_{22})]}{\beta_{11}}$$

$$+ q_3\alpha_{12} + q_4\alpha_{22} - q_5(\alpha_{12}^2 - \beta_{12}^2 - \alpha_{11}\alpha_{22} + \beta_{11}\beta_{22}).$$

When substituted back into (6.6) the resulting quadratic form has the following matrix

$$\begin{pmatrix} 1 & 2\frac{\beta_{12}}{\beta_{11}} & -2\frac{\beta_{11}\alpha_{12}+\alpha_{11}\beta_{12}}{\beta_{11}} \\ 2\frac{\beta_{12}}{\beta_{11}} & 4\frac{\beta_{22}}{\beta_{11}} & -4\frac{\alpha_{11}\beta_{22}-\alpha_{12}\beta_{12}}{\beta_{11}} \\ -2\frac{\beta_{11}\alpha_{12}+\alpha_{11}\beta_{12}}{\beta_{11}} & -4\frac{\alpha_{11}\beta_{22}-\alpha_{12}\beta_{12}}{\beta_{11}} & 4(\alpha_{12}^2 - \beta_{12}^2 + \beta_{11}\beta_{22} + \frac{\alpha_{11}^2\beta_{22}-2\alpha_{11}\alpha_{12}\beta_{12}}{\beta_{11}}). \end{pmatrix}$$

For this to be positive definite it is sufficient for successive determinants anchored at the top left to be positive, but the values of these are

$$\det(Q_{1,1}) = 1,$$

$$\det(Q_{1\ldots2,1\ldots2}) = 4\frac{\beta_{11}\beta_{22} - \beta_{12}^2}{\beta_{11}},$$

$$\det(Q_{1\ldots3,1\ldots3}) = 16\frac{(\beta_{11}\beta_{22} - \beta_{12}^2)^2}{\beta_{11}^2}.$$

The latter two are positive precisely because β is the imaginary part of a period matrix and hence itself positive definite. $\qquad\square$

A summary of the algorithm developed in this section is given in Algorithm 3.

Algorithm 3 Find all $q_1, \ldots, q_5$ for a given genus 2 curve of a given covering order
---
Require: $h \in \mathbb{N}_+$, genus 2 period matrix $\hat{\tau}$.
Ensure: s is the set of all quintuples $(q_1, \ldots, q_5)$ satisfying (6.5) and (6.6).
  $s \leftarrow \emptyset$.
  $(\mathbf{q} \mapsto A\alpha) \leftarrow$ solution of (6.5) in terms of $\alpha \in \mathbb{Z}^k$.
  $\alpha^T Q\alpha = h^2 \leftarrow$ substitution into (6.6).
  $\lambda \leftarrow$ minimum eigenvalue of Q.
  for all α with $||\alpha|| \leq \frac{h}{\sqrt{\lambda}}$ **do**
    if $\alpha^T Q\alpha = h^2$ **then**
      $s \leftarrow s \cup \{A\alpha\}$
    end if
  end for

6.3   Solve for m_{ij}

Now that we have all possible values for q_i (given a covering order) we must refer back to the equations satisfied by m_{ij},

$$0 = \det \begin{pmatrix} 1 & 0 & m_{11} & m_{12} \\ 0 & 1 & m_{21} & m_{22} \\ \hat{\tau}_{11} & \hat{\tau}_{12} & m_{31} & m_{32} \\ \hat{\tau}_{12} & \hat{\tau}_{22} & m_{41} & m_{42} \end{pmatrix}, \tag{6.10}$$

where we now know each possible 2×2 determinant made up of m_{ij} and want the m_{ij} themselves.

Once again we are faced with solving a system of quadratic diophantine equations – dangerously close to the impossible problem of [7] discussed in the Klein chapter earlier. However, this situation is not hopeless. Another result of Murabayashi's paper [28] is that at least one solution exists for each set of q_i found. We also have extra freedoms and symmetries that allow us to restrict our search for m_{ij} to a finite, and even practical, subset of all 4×2 matrices.

6.3.1   Matrices equivalent to given M

In order for a 2×2 period matrix to cover an elliptic curve equation (6.10) must be solvable, and then the values of m_{ij} determine the elliptic curve. However the solution will not be unique, and certain obvious transformations can be applied to simplify the situation

1. We may perform a symplectic transformation on the genus 2 period matrix.

2. We may perform a symplectic transformation on the (implicit) genus 1 matrix.

The first of these is less interesting. It changes the admissible values for q_i since the linear equation for q_i, (6.5) depends on the particular values of the genus 2 period matrix $\hat{\tau}$. At this stage we are imagining we have an exhaustive list of possible solutions q_i so this doesn't really gain anything.

The second freedom is much more fruitful. A genus 1 symplectic transformation is simply a member $T \in SL(2, \mathbb{Z})$. Its action on M follows from the definition (6.1), and can be seen in (6.2) to be

$$M \mapsto MT,$$

Each q_i, being a 2×2 sub-determinant of M, gets mapped to itself:

$$\begin{vmatrix} m_{i1} & m_{i2} \\ m_{j1} & m_{j2} \end{vmatrix} \mapsto \left| \begin{pmatrix} m_{i1} & m_{i2} \\ m_{j1} & m_{j2} \end{pmatrix} T \right| = \begin{vmatrix} m_{i1} & m_{i2} \\ m_{j1} & m_{j2} \end{vmatrix}.$$

As a result, given a solution to the equations for q_i, each M falls into an $SL(2, \mathbb{Z})$ orbit. $SL(2, \mathbb{Z})$ may be generated by elementary matrices, which act act as integer-invertible column operations on M itself. We use this freedom to pick out a special representative from each orbit and prove that only finitely many such solutions exist.

6.3.2 Orbit representatives

We will now use the freedoms described above. By performing elementary integer-invertible column operations on M we can guarantee that rows 3 and 4 have the form

$$\begin{pmatrix} m_{31} & m_{32} \\ m_{41} & m_{42} \end{pmatrix} = \begin{pmatrix} x & 0 \\ y & z \end{pmatrix}. \tag{6.11}$$

Now, if the determinant of this matrix is nonzero (i.e. $q_1 \neq 0$) then neither x nor z can vanish ($xz = q_1$) and we can also demand $0 \leq y < |z|$. We will deal with $q_1 = 0$ later, as it splits into more complicated cases.

The important point to note here is that there are only finitely many possible values this submatrix can take for a given q_1, so we algebraically try to fill in the remaining 4 ms from each possibility.

Hence, here and afterwards we will adopt the following notation

- Letters from the end of the alphabet denote variables we know very little about.

- Letters from the beginning of the alphabet denote unknowns with only finitely many solutions – once every variable to be solved for is in this form, algorithmic loops can be written to exhaustively find all solutions.

In this manner we rewrite (6.11) as

$$\begin{pmatrix} m_{31} & m_{32} \\ m_{41} & m_{42} \end{pmatrix} = \begin{pmatrix} a & 0 \\ c & d \end{pmatrix}. \tag{6.12}$$

6.3.3   Continuing progress

Having found a candidate as above, the remaining equations are

$$m_{12}c - m_{11}d = q_2,$$

$$m_{22}c - m_{21}d - m_{12}a = q_3,$$

$$-m_{22}a = q_4,$$

$$m_{11}m_{22} - m_{12}m_{21} = q_5.$$

Once again we are reduced to a single quadratic equation (the one for q_5) and a set of linear equations in m_{ij}. In this case we will see that there are enough linear constraints that ad-hoc solution reduces the quadratic equation to a triviality.

The next obvious equation to attack is q_4, yielding just one value for m_{22} (call it e), and remaining equations

$$m_{12}c - m_{11}d = q_2,$$

$$ec - m_{21}d - m_{12}a = q_3,$$

$$m_{11}e - m_{12}m_{21} = q_5.$$

By the assumption $q_1 \neq 0$, a and d are nonzero. Hence the first two equations here are nondegenerate and can be solved to give a set of solutions $m_{11}(\alpha), m_{12}(\alpha), m_{21}(\alpha)$ in terms of a single integral parameter α. This reduces the final equation to

$$Q(\alpha) = q_5, \tag{6.13}$$

which we can simply solve over $\mathbb{R}$ and check whether the solutions are actually integers (n.b. if one is integral, the other will be).

Thus, at the conclusion, provided $q_1 \neq 0$ each choice for (6.12) extends consistently to 0, 1 or 2 solutions for the rest of M, depending on how many solutions the quadratic equation (6.13) has.

6.3.4   When $q_1 = 0$

Unfortunately this case fragments rather quickly. Looking at (6.12) again we now know that $ad = 0$ giving four mutually exclusive possibilities, three of which are interesting

1. $a = c = d = 0$. This implies all q_i and hence m_{ij} vanish. We will call cases like this where the only solution is $M = 0$ degenerate, and ignore them.

2. $a = 0$, at least one of $c, d \neq 0$. This can only occur if $q_4 = 0$

3. $a \neq 0, d = 0, c = 0$. This can only occur if $q_2 = 0$.

4. $a \neq 0, d = 0, c \neq 0$. No further restrictions on q_i vanishing.

These cases are mutually exclusive relative to a known a, c and d but a given set of q_i allows for more than one solution in those variables. Thus when designing an algorithm to find m_{ij} given q_i we must usually consider more than one situation in the list. Specifically, one of the following cases will hold

- $q_1 = 0; q_2, q_4 \neq 0$. A solution could only come from 4 above.

- $q_1, q_4 = 0; q_2 \neq 0$. A solution could only come from 2 or 4 above.

- $q_1, q_2 = 0; q_4 \neq 0$. A solution could come from either 3 or 4 above.

- $q_1, q_2, q_4 = 0$. A solution could come from any of 2, 3 or 4 above.

In subsequent sections we will deal with exactly what solutions can be obtained from cases 2, 3 and 4 in the numbered list. An outline of the algorithmic application for these cases is provided in Algorithm 4; it essentially duplicates the information in the previous list, though in more procedural terms.

6.3.5 The case $a = 0$ (implies $q_1 = 0, q_4 = 0$).

In the first case further column operations can reduce the bottom submatrix (reusing the variable-name c) to

$$\begin{pmatrix} 0 & 0 \\ c & 0 \end{pmatrix},$$

and the interesting remaining equations are

$$m_{12}c = q_2,$$

$$m_{22}c = q_3,$$

$$m_{11}m_{22} - m_{12}m_{21} = q_5.$$

Algorithm 4 Construct viable solutions in m_{ij} from q_i

Require: Integers $q_1, \ldots, q_5$ not all zero.

Ensure: s is a set of matrices M, each satisfying (6.10), and with one from each orbit of all solutions under $SL(2, \mathbb{Z})$.

  if $q_1 \neq 0$ **then**

    $s \leftarrow \{$solutions from $\begin{pmatrix} a & 0 \\ c & d \end{pmatrix}$ according to 6.3.3$\}$

  else

    $s \leftarrow \{$solutions from $\begin{pmatrix} a & 0 \\ c & 0 \end{pmatrix}$ according to 6.3.7$\}$

    if $q_4 = 0$ **then**

      $s \leftarrow s \cup \{$solutions from $\begin{pmatrix} 0 & 0 \\ c & 0 \end{pmatrix}$ according to 6.3.5$\}$

    end if

    if $q_2 = 0$ **then**

      $s \leftarrow s \cup \{$solutions from $\begin{pmatrix} a & 0 \\ 0 & 0 \end{pmatrix}$ according to 6.3.6$\}$

    end if

  end if

The first two permit just finitely many solutions in m_{12} and m_{22}. So we may write (reusing a and b)

$$M = \begin{pmatrix} m_{11} & a \\ m_{21} & b \\ 0 & 0 \\ c & 0 \end{pmatrix}.$$

Now, one of q_2 and q_3 are nonzero (otherwise the solution is again degenerate) so one of a and b are nonzero. Suppose without loss of generality $q_2 \neq 0$. Then $a \neq 0$ and further column operations can guarantee $0 \leq m_{11} < |a|$. Calling this d in line with our finite conventions, the final equation becomes

$$db - am_{21} = q_5,$$

which uniquely determines m_{21} (and may have no solution over $\mathbb{Z}$).

6.3.6   The case $c, d = 0$ (implies $q_1, q_2 = 0$).

The bottom submatrix becomes

$$\begin{pmatrix} a & 0 \\ 0 & 0 \end{pmatrix}.$$

Essentially the same argument as above applies to give only finitely many possible orbits.

6.3.7   The case $a, c \neq 0; d = 0$ (only guarantees $q_1 = 0$).

The bottom submatrix now becomes

$$\begin{pmatrix} a & 0 \\ c & 0 \end{pmatrix},$$

with corresponding equations

$$m_{12}c = q_2,$$

$$m_{22}c - m_{12}a = q_3,$$

$$m_{22}a = q_4,$$

$$m_{11}m_{22} - m_{12}m_{21} = q_5.$$

If $q_2 = q_4 = 0$ then these equations are again degenerate. If both q_2 and q_4 are nonzero then these equations have immediate solution (both m_{12} and m_{22} also have finitely many possibilities; column operations then give (say) m_{11} finitely many options and then there is just one variable to be solved for in the q_5 equation).

So without loss of generality assume $q_2 \neq 0, q_4 = 0$. Then $m_{22} = 0$ and m_{12} is also limited to finitely many values. The overall matrix becomes

$$M = \begin{pmatrix} u & b \\ v & 0 \\ a & 0 \\ c & 0 \end{pmatrix},$$

with interesting equation

$$-bv = q_5,$$

which can be solved immediately for v. One final column operation guarantees $0 \leq u < |b|$ and gives finitely many possible orbits.

6.4   Martens algorithm

Martens algorithm, described in [26], simplifies the next step of converting discovered solutions in m_{ij} into elliptic curves covered by the genus 2 curve.

Following Martens, for each solution m_{ij} which we write as the matrix M, we know that

$$\hat{\Pi}\lambda = M\Pi \tag{6.14}$$

for some 2×1 matrix λ and period matrix Π. As input to the entire algorithm we know $\hat{\Pi}$ explicitly so our task is simply to compute Π.

Following Martens' calculation closely we are first provided with a factorisation of M as

$$M = TNS = T \begin{pmatrix} 0 & x \\ 0 & 0 \\ -1 & 0 \\ 0 & -\epsilon \end{pmatrix} S,$$

where T is a symplectic integer matrix (of dimension 4); S is a unimodular matrix (of dimension 2); ϵ is either 0 or 1 and $x \in \mathbb{Z}$.

Substituting this back into (6.14) we obtain

$$\hat{\Pi}\lambda = TNS\Pi,$$

$$\hat{\Pi}'\lambda = N\Pi',$$

$$\begin{pmatrix} \mathbf{1} \\ \hat{\tau}' \end{pmatrix} \lambda' = N \begin{pmatrix} 1 \\ \tau' \end{pmatrix}.$$

In each case going back and forward between the primed and unprimed version is simple, being multiplication by an explicitly calculable matrix.

Expanding these equations yields

$$\begin{pmatrix} \lambda_1' \\ \lambda_2' \\ \hat{\tau}_{11}'\lambda_1' + \hat{\tau}_{12}'\lambda_2' \\ \hat{\tau}_{21}'\lambda_1' + \hat{\tau}_{22}'\lambda_2' \end{pmatrix} = \begin{pmatrix} x\tau' \\ 0 \\ -1 \\ -\epsilon\tau' \end{pmatrix}.$$

Upon eliminating λ_i' we get

$$\hat{\tau}_{11}' x \tau' = -1,$$

$$\tau'(x\hat{\tau}_{12}' + \epsilon) = 0.$$

In passing we note that this means $x, \epsilon \neq 0$ and that in $\hat{\tau}'$ we have produced a neat form of the genus 2 period matrix that makes explicit its cover of this curve ($\hat{\tau}_{12}'$ being rational leads to splitting of the Θ-functions on the genus 2 curve in terms of the Θ-functions on the elliptic curve).

But the main point is that

$$\tau' = -\frac{1}{x\hat{\tau}_{11}'},$$

that is, we have found an algorithm to produce the elliptic period matrix τ' from $\hat{\tau}$ and M.

6.5 Equivalent genus 1 curves

We now have a set of genus 1 period matrices covered by the initial genus 2 curve. However we have no guarantee that the curves given by these period matrices are distinct. As a simple change of symplectic basis on the genus 1 curve acts as the modular group, we can use the well-known fact that the modular group has the following fundamental domain in the upper half-plane (i.e. the space of period matrices).

$$D = \left\{ z : |z| > 1, -\frac{1}{2} < \operatorname{Re} z \leq \frac{1}{2} \right\} \cup \left\{ z : |z| = 1, 0 \leq \operatorname{Re} z \leq \frac{1}{2} \right\},$$

as in Figure 6.2. Abstractly, we simply put every discovered period matrix into this domain and remove duplicates.

The one remaining challenge is finding an algorithm to do this that is implementable and guaranteed to terminate. Algorithms seem to exist for parts of this procedure (see for example [24]) even in higher dimensions. However, none of the algorithms I have encountered are suitable for actual implementation. Typically they rely on finding the optimal transformation to apply at any given step by minimizing a quantity over all possible modular transformations.

Therefore, we will develop the following algorithm which *will* terminate. I would be entirely unsurprised to learn it has been done before.

Figure 6.2: Fundamental domain of the modular group on the upper half-plane.

Given a τ_0, assumed without loss of generality to satisfy $-\frac{1}{2} < \operatorname{Re} \tau_0 \leq \frac{1}{2}$, we want an algorithm to map it into D. Form the sequence defined by iterating the function f which inverts τ_n and brings its real part back to near 0

$$\tau_{n+1} = f(\tau_n) := k - \frac{1}{\tau_n},$$

where $k \in \mathbb{Z}$ is chosen so that $-\frac{1}{2} < \operatorname{Re} \tau_{n+1} \leq \frac{1}{2}$. We will show that τ_n enters the fundamental domain after a finite number of steps. Now, τ_n is in the fundamental domain when $|\tau_n| \geq 1$, so we simply need to prove that $|\tau_n| \geq 1$ eventually and we will have an algorithm.

The following, rather technical theorem, gives the required result

Theorem 13. *There exists an n such that* $\operatorname{Im} \tau_n \geq \frac{5}{12}$, *and hence* $\tau_{n+1} \in F$.

Proof. Consider how the iteration affects the imaginary parts of the sequence. Let $\tau_n = \alpha_n + \mathrm{i}\beta_n$, then

$$
\begin{aligned}
\beta_{n+1} &= \operatorname{Im} f(\alpha_n + \mathrm{i}\beta_n) \\
&= \operatorname{Im}\left(\frac{-1}{\alpha_n + \mathrm{i}\beta_n}\right) \\
&= \frac{\beta_n}{\alpha_n^2 + \beta_n^2} \\
&\geq \frac{\beta_n}{\frac{1}{4} + \beta_n^2}.
\end{aligned}
$$

If we let $g(z) = \frac{z}{\frac{1}{4}+z^2}$ then this can be restated as $(\operatorname{Im} \circ f)(\tau_n) \geq (g \circ \operatorname{Im})(\tau_n)$. Now consider the associated real sequence

$$a_n = g^n(\beta_0) = g^n(\operatorname{Im} \tau_0).$$

Figure 6.3: f sends points with imaginary part above $\frac{5}{12}$ to D in one step.

Because $g(z)$ is an increasing function when $z \leq \frac{1}{2}$ we can use the following induction argument to show that provided $\operatorname{Im} \tau_n \leq \frac{1}{2}$, the inequality $\operatorname{Im} \tau_n \geq a_n$ is satisfied:

$$
\begin{aligned}
\operatorname{Im} \tau_{n+1} &= \operatorname{Im} f(\tau_n) \\
&\geq g(\operatorname{Im} \tau_n) \\
&\geq g(a_n) \\
&= a_{n+1}.
\end{aligned}
$$

But given $\epsilon > 0$, a_n eventually exceeds $\frac{1}{2} - \epsilon$ (if not, a_n would be bounded above and increasing, hence convergent and could only converge to a fixed point of g, i.e. $\frac{\sqrt{3}}{2}$, a contradiction). But this means that $\operatorname{Im} \tau_n$ eventually exceeds $\frac{1}{2} - \epsilon$ too.

Now take $\epsilon = \frac{1}{12}$ (say) and refer to Figure 6.3. Since there is an N such that $\operatorname{Im} \tau_N \geq \frac{1}{2} - \frac{1}{12} = \frac{5}{12}$, τ_N is above the dotted line and in one of the lower shaded regions. But on the very next iteration of $\tau_N \mapsto \frac{-1}{\tau_N}$ these map to the upper shaded regions which are sent to the fundamental domain by f itself so $\tau_{N+1} \in D$ ☐

A summary of the resulting procedure for arbitrary τ_0 is given by Algorithm 5.

Algorithm 5 Perform symplectic transform on $\tau \in H$ to map it into the fundamental domain $D = \{z \in H : |z| \geq 1, -\frac{1}{2} < |\operatorname{Re}(z)| \leq \frac{1}{2}\}$.

Ensure: $\tau \in D$.
Require: $\tau \in H = \{z \in \mathbb{C} | \operatorname{Im} z > 0\}$.
  $\tau \leftarrow \tau - a$ where $a = \max\{x \in \mathbb{Z} : \operatorname{Re} \tau - x > -\frac{1}{2}\}$
  while $|\tau| < 1$ **do**
    $\tau \leftarrow -\frac{1}{\tau}$
    $\tau \leftarrow \tau - a$ where $a = \max\{x \in \mathbb{Z} : \operatorname{Re} \tau - x > -\frac{1}{2}\}$
  end while

6.6   Generic algorithm

At its most basic, our algorithm now consists of 4 steps.

1. Use Algorithm 3 to solve the equations in q_i giving each of the subdeterminants of M. A finite set of possible q_i results for each covering order.

2. Using Algorithm 4, for each solution q_i above find a representative of each of the finitely many classes of M with those subdeterminants.

3. For each class of Ms discovered, calculate the period matrix of the elliptic curve corresponding to this solution using Martens' algorithm (described in [26]).

4. Use Algorithm 5 to map each elliptic period matrix to the canonical one and eliminate duplicates.

6.7   Numeric extension to algorithm

If a period matrix is given over a field extension of $\mathbb{Q}$ with degree more than 2 then the algorithm presented in 6.2 extracts progressively more information and reduces the complexity of the inevitable exhaustive search for valid $\mathbf{q}$. However, in practice we often do not know the period matrix $\hat{\tau}$ symbolically and only have numeric approximations to its entries; in these cases the algorithm as presented is inapplicable. But notice that Theorem 12 actually relies only on splitting τ into real and imaginary parts. This can be accomplished just as well in the numeric case.

The linear equations (6.5) can then be solved without regard for integrality and substituted into (6.6) to obtain a real quadratic equation

$$\begin{pmatrix} q_3 & q_4 & q_5 \end{pmatrix} A \begin{pmatrix} q_3 & q_4 & q_5 \end{pmatrix}^T = h^2$$

which we wish to solve in integers q_3, q_4, q_5, h. Theorem 12 applies and puts bounds on q_i in terms of h so an exhaustive search can be performed, to within a certain tolerance.

With (6.6) solved the final step is to check that the q_1 and q_2 required to satisfy (6.5) are also integral (to a certain degree of certainty). All other stages of the algorithm deal with these integers, q_i, as found and follow through without modification.

6.7.1 Numeric example

Here we present a simple example of the conversion above. The period matrix

$$\hat{\tau} \cong \begin{pmatrix} 0.00865693 + 0.34649672i & 0.45003704 + 0.12098331i \\ 0.45003704 + 0.12098331i & 0.63692830 + 0.22355689i \end{pmatrix}$$

occurred in calculations, with reasons to believe it should cover a curve of genus 1 with $h = 2$ in (6.6). We will find out which curve is covered and put $\hat{\tau}$ into a form making this manifest.

In this case the linear constraints on q_i reduce (examining purely real and imaginary components) to the matrix equation

$$\begin{pmatrix} 1 & 0.00865693 & 0.45003704 & 0.63692830 & -0.25984426 \\ 0 & 0.34649672 & 0.12098331 & 0.2235569 & 0.11373494 \end{pmatrix} \mathbf{q} \cong 0,$$

which has solution

$$q_1 \cong -0.44701437 q_3 - 0.63134291 q_4 + 0.26268583 q_5,$$

$$q_2 \cong -0.34916150 q_3 - 0.64519193 q_4 - 0.32824247 q_5.$$

The form associated to the quadratic constraints (6.6) becomes (with $\mathbf{x} := (q_3, q_4, q_5)$)

$$\mathbf{x}^T A \mathbf{x} := \mathbf{x}^T \begin{pmatrix} 1 & 0.69832299 & -0.89402874 \\ 0.69832299 & 2.58076771 & -0.60620089 \\ -0.89402874 & -0.60620089 & 1.05074332 \end{pmatrix} \mathbf{x} \cong 4,$$

which has a minimum eigenvalue of about 0.127, so

$$|\mathbf{x}^T A \mathbf{x}| \geq 0.127 ||\mathbf{x}||^2,$$

and we only have to search for solutions with

$$||\mathbf{x}||^2 \leq \frac{4}{0.127} \leq 32.$$

Performing this search with very loose bounds (only asking that $|\mathbf{x}^T A \mathbf{x} - 4| \leq 0.01$)

gives 3 possible (essentially distinct) solutions in q_3, q_4, q_5. Of these only

$$\begin{pmatrix} q_1 & q_2 & q_3 & q_4 & q_5 \end{pmatrix} \cong \begin{pmatrix} -1.00000000 & -1.99999986 & 2 & 1 & 2 \end{pmatrix}$$

also has q_1 and q_2 within 0.01 of an integer. We have now essentially confirmed that the original surface covered one with genus 1, and all that remains is to put $\hat{\tau}$ into a form that makes the cover manifest.

To do this we must find the appropriate matrices M. The algorithm presented in Section 6.3 applies directly to any integers $\mathbf{q}$ no matter how they were obtained and gives two classes for M

$$\left\{ \begin{pmatrix} -2 & -2 \\ 0 & -1 \\ 1 & 0 \\ 0 & -1 \end{pmatrix}, \begin{pmatrix} -2 & 0 \\ -2 & -1 \\ 1 & 0 \\ 0 & -1 \end{pmatrix} \right\}.$$

Choosing the first of these and applying Martens algorithm as in Section 6.4 we obtain the Poincaré decomposition

$$\begin{pmatrix} -1 & 0 & -2 & 0 \\ 0 & 1 & 0 & -1 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} -2 & -2 \\ -0 & -1 \\ 1 & 0 \\ 0 & -1 \end{pmatrix} = \begin{pmatrix} 0 & 2 \\ 0 & 0 \\ -1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Applying this symplectic transformation to $\hat{\tau}$ yields

$$\hat{\tau}' \cong \begin{pmatrix} 0.34303902 + 0.16638383\mathrm{i} & -0.49999999 + 0.00000003\mathrm{i} \\ -0.49999999 + 0.00000003\mathrm{i} & 0.24710783 + 0.43466967\mathrm{i} \end{pmatrix},$$

which has the expected $-\frac{1}{h} = -\frac{1}{2}$ terms off-diagonal. It follows that the elliptic curve covered has period matrix

$$\tau' = -\frac{1}{x\hat{\tau}'_{11}} \cong -1.179968913 + 0.5723189827\mathrm{i},$$

which has canonical form

$$\tau'' \cong 0.4999999978 + 1.59004955\mathrm{i}$$

via the symplectic transformation given by $z \mapsto \frac{-1}{z+1}$.

An identical argument applied to the second M gives us the elliptic period

$$\tau'' = 0.49421567 + 0.86933933i.$$

6.8 Summary

In this chapter we have developed an algorithm which could, in principle, be used to enumerate all curves covered by an arbitrarily chosen genus 2 curve and give an explicit relation between their period matrices. Each stage of the process is computationally feasible, and could quite easily be integrated into some kind of "black box" function for use when the methods of derivation are not a primary concern.

There are two main open issues with this approach. First, it does not currently integrate a means of bounding the degree of the cover, h. Thus in principle the search will never terminate. However, often we have external reasons to either bound the maximum possible h or only take interest in lower values of h.

The second issue is the repeated appearance of the same elliptic curves due to the fact that self-covers of elliptic curves exist to arbitrarily high degree. That is, given any projection $\pi : \hat{\Sigma} \to \Sigma$ of degree h, any self-cover $\sigma : \Sigma \to \Sigma$ will induce infinitely many "spurious" results

$$\pi' = \sigma \circ \pi$$
$$\pi'' = \sigma^2 \circ pi$$
$$\vdots$$

However, any cover derived in this manner (even via different self-covers) will ultimately be derived from some primitive genus-2 over genus-1 cover. Further work could involve reducing any cover found to this simplest possible version, probably via a search of all divisors of the degree of the discovered map.

7

FURTHER WORK

The work in this thesis is by no means a complete solution to the issue of finding an optimal homology basis for arbitrary curves and making use of it. Many interesting questions remain unanswered, and large steps could be made towards a richer, more intuitive interface for the calculations enabled by this work.

`CyclePainter` could benefit from an overhaul of its algorithms. Work was started based on the existing Maple libraries for dealing with Riemann surfaces but there's actually comparatively little use made of those facilities. Decoupling the program from Maple would at the very least ease installation by not relying on the rather fragile interface to the Maple external call libraries. Given Maple's rather lax attitude to encapsulation and global state, such a move could also improve robustness.

Similarly, the algorithms used by `extcurves` could be overhauled for large performance gains. I believe converting the intersection calculation code to use principles more similar to those found in `CyclePainter` (work out whether an intersection is real by counting cuts crossed rather than by analytic continuation) would yield orders of magnitude speed increase. This work wasn't undertaken before now because `extcurves` is the classical "fast enough". However, such a performance improvement would open up larger problems to analysis and probably new methods too. For example it would be feasible to iterate computationally through many attempted homologies looking for one with desired properties.

Another potential direction for `CyclePainter` and `extcurves`, and one with rather more mathematical interest, would be extending their representations to the hyperbolic models of Riemann surfaces. Some thought would be needed on graphically representing paths that cross the boundaries of the disc intuitively, particularly while actually being constructed. But the problems are probably not insurmountable. On the `extcurves` side intersection numbers would be easier to calculate but there is more interest in computing arbitrary period matrices.

The most interesting work to be done following on from the results on Klein's and Bring's curve is probably extending those techniques to general curves. Currently, much use is made of symmetry. It would be very interesting to discover either methods that work without such requirements.

In a slightly more limited scope, while answering questions on those two surfaces shared many techniques, each has its unique challenges and it would be very interesting to discover a fully unified argument that worked even on this limited sample. Such a unification could point towards more generally applicable versions of the arguments.

Looking at depth rather than breadth, the current mapping from hyperbolic to algebraic representations is rather spartan. It's fully pinned down only at the vertices of each polygon and slightly hazily specified along the interesting geodesics of the surfaces. Discovering fully explicit formulae mapping from every point of the hyperbolic models to a corresponding algebraic model would be valuable for future interrelations of results.

The work on genus two curves covering elliptic ones could usefully be extended in multiple directions. Currently we can answer the question of which curves are covered, but not give an explicit formula for the covering map.

In another direction, the work is rather limited in only dealing with genus at most two. While there are combinatorial difficulties in answering the general question naively (as noted, it leads to solving a quadratic diophantine equation in many variables – a hard problem), in this case we potentially have more information derived from geometric sources and that may open up more avenues that reduce the problem to one that's practically solvable.

The above problems are interesting for their own sake, however potentially the most useful extension and the original goal of this project was deciding on the optimal basis for any arbitrarily presented curve. Inroads have been made here, but there's still a very large reliance on human effort.

Gilman goes some of the way towards specifying what might be possible in [17]. However her results take as given a covering map from one surface to another, and only give what I would call an aesthetically pleasing result (in the context of simple homology bases) if the degree of that map is prime.

The basis proven to exist there may not even be the optimal choice, and computing which basis is most symmetric (whatever that means) is still very much an open question and far too reliant on ingenuity in each individual case at the moment.

BIBLIOGRAPHY

[1] J. V. Armitage and W. F. Eberlein, *Elliptic functions*, London Mathematical Society, 2006.

[2] Babelon, Bernard, and Talon, *Introduction of classical integrable systems*, Cambridge University Press, 2003.

[3] H. W. Braden and V. Z. Enolski, *Remarks on the complex geometry of the 3-monopole* (2006), available at `math-ph:0601040`.

[4] H W Braden and T P Northover, *Klein's curve*, Journal of Physics A: Mathematical and Theoretical **43** (2010), no. 43.

[5] Thomas Breuer, *Characters and automorphism groups of compact Riemann surfaces*, Cambridge University Press, 2000.

[6] E. S. Bring, Quart. J. Math. **6** (1864).

[7] J. L. Britton, *Integer solutions of systems of quadratic equations*, Math. Proc. Camb. Phil. Soc. **86** (1979), 385–389.

[8] Peter Buser and Robert Silhol, *Geodesics, periods, and equations of real hyperelliptic curves*, Duke Math. J. **108** (2001), no. 2, 211–250.

[9] Maurice Craig, *A sextic diophantine equation*, Aust. Math. Soc. Gazette **29** (2002), 27–29.

[10] Antonella d'Avanzo, *On charge 3 cyclic monopoles*, Ph.D. Thesis, 2010.

[11] Bernard Deconinck and Mark van Hoeij, *Computing Riemann matrices of algebraic curves*, Mathematics of Computation **73** (2004), no. 247, 1417–1442.

[12] Rainer Dietmann, *Small solutions of quadratic Diophantine equations*, Proc. London Math. Soc. **86** (2003), 545–582.

[13] R. H. Dye, *A plane sextic curve of genus 4 with A_5 for collineation group*, J. London Math. Soc. **52** (1995), 97–110.

[14] Greg Egan, *Klein's quartic equation*, `http://www.gregegan.net/SCIENCE/KleinQuartic/KleinQuarticEq.html`, 2006.

[15] H. M. Farkas and I. Kra, *Riemann surfaces*, Springer-Verlag, 1980.

[16] John D. Fay, *Theta functions on Riemann surfaces*, Springer, 1973.

[17] Jane Gilman, *A matrix representation for automorphisms of compact Riemann surfaces*, Linear Algebra and its Applications **17** (1977), 139–147.

[18] ———, *Canonical symplectic representations for prime order conjugacy classes of the mapping-class group*, Journal of Algebra **318** (2007), 430–455.

[19] P. A. Griffiths, *Linearizing flows and a cohomological interpretation of Lax equations*, American Journal of Mathematics **107** (1985), 1445–1484.

[20] Phillip Griffiths and Joseph Harris, *Principles of algebraic geometry*, Wiley-Blackwell, 1994.

[21] Morris W. Hirsch, *Differential topology*, Springer, 1976.

[22] H. Karcher and M. Weber, *The geometry of Klein's Riemann surface*, Vol. Math. Sci. Res. Inst. Publ. 35, Cambridge University Press, 1998.

[23] Frances Kirwan, *Complex algebraic curves*, London Mathematical Society, 1992.

[24] Helmut Klingen, *Introductory lectures on Siegel modular forms*, Cambridge University Press, 1990.

[25] John M. Lee, *Introduction to smooth manifolds*, Springer, 2003.

[26] Henrik H. Martens, *A footnote to the Poincaré complete reducibility theorem*, Publications Matemàtiques **36** (1992), 111–129.

[27] Rick Miranda, *Algebraic curves and riemann surfaces*, American Mathematical Society, 1991.

[28] Naoki Murabayashi, *The moduli space of curves of genus two covering elliptic curves*, Manuscripta Mathematica **84** (1994), no. 2, 125–133.

[29] Harry E Rauch and J Lewittes, *The Riemann surface of Klein with 168 automorphisms*, Problems in analysis. a symposium in honor of Salomon Bochner, 1970, pp. 297–308.

[30] Gonzalo Riera and Rubí E. Rodríguez, *The period matrix of Bring's curve*, Pacific Journal of Mathematics **154** (1992), 179–199.

[31] C. L. Tretkoff and M. D. Tretkoff, *Combinatorial group theory, Riemann surfaces and differential equations*, Contemporary Mathematics **33** (1984), 467–517.

[32] Egbert Vrieskorn and Horst Knörrer, *Plane algebraic curves*, Birkhäuser, 1986.

A

EXTCURVES API

A.1   Accessing the library

The `extcurves` routines are compiled into the library `extcurves.mla`, so it is necessary to tell Maple where this is in order to access it. The basic commands are

```
> march('open', "./extcurves.mla");
> with(extcurves):
```

A.2   Design principles

There are two fundamental datatypes introduced by `extcurves`. The first is mainly for convenience; we package the conventional first three arguments used by `algcurves` routines (a polynomial describing the surface, and the two variables in that polynomial) into a single `Record` structure.

For example

```
> curve := Record('f'=t^3*s+s^3+t, 'x'=t, 'y'=s):
```

Note the single quotes around the left hand side of each equation. This is necessary if any of `f`, `x`, `y` have definitions in Maple and prevents substitution too early, so the following code would work

```
> f := x^3 * y + y^3 + x;
> curve := Record('f'=f, 'x'=x, 'y'=y);
```

The second, and more important, datatype introduced is called an `extpath`. It describes a path on a Riemann surface as a sequence of straight line segments in the `x` plane together with a specification of the sheet at the initial point in this projection (so strictly, it's only meaningful relative to a given curve object above). This system is flexible enough to give any path up to homotopy, and simple enough to be amenable to computation.

When a homology basis is needed, it is simply a list of `extpath` objects.

A.3   Creating an extpath

There are multiple functions to create `extpath` objects. In the descriptions there is always an implicit `curve` object that projections and sheets are defined relative to.

When the x projection is referred to, we mean the projection of `curve:-f` onto `curve:-x`. Similarly the sheet will be defined by the value of `curve:-y` on `curve:-f` at `curve:-x`

`extpath_from_description(curve, base, basey, descr, how)`

curve	curve Record.
base	Initial point in x plane.
basey	Initial sheet at `base`.
descr	List of of branch points and number of circuits to make around each. Each entry consists of a list [`branch`, `N`] which means go around `branch` `N` times anticlockwise.
how	Either a table where the entry for `branch` is the radius of the circle to use around that point, or a complete list of branch points (from which appropriate radii will be derived automatically).
description:	Creates an extpath from an algebraic description of the path in terms of the branch points circuited. `how` is necessary to ensure that the circular paths around each branch don't interfere with nearby ones.

`extpath_from_homology(curve, wanted, hom)`

curve	curve Record
wanted	Token representing the desired homology cycle. For example `a[1]`.
hom	(Optional) Homology description in extended algcurves format
description:	Extracts an extpath from a homology description conforming to the output of algcurves' `homology` function. By default calls `homology` itself to get this information, but could work with a modified version if needed.

`extpath_from_painter(curve, points, close_y)`

curve	curve Record giving surface
points	List of points
close_y	Specification of sheet at first point in `points`
description:	Creates an extpath consisting of straight-line segments between given points. This is essentially the same as `extpath_from_points`

except that it makes sure your choice of sheet is consistent with the curve specified. It chooses the sheet closest to the value given.

return value: extpath again

`extpath_from_parametrised(curve, paths, y0)`

curve	curve Record
paths	List of expressions in t, each specifying a straight line between $t = 0$ and $t = 1$
y0	Sheet at beginning of path
description:	Creates an extpath from the given information, choosing the sheet closest to guven y0. A slightly different form of `extpath_from_painter`.

`extpath_from_points(points, sheet)`

points	List consisting of points in x plane
sheet	Value for y at the first point in the list
description:	Low level function that simply creates an extpath with straight-line segments between each point in `points`. Since it has no knowledge of the curve you must be wary of floating point errors in specifying the sheet. Consider using the more forgiving `extpath_from_painter`.
return value:	extpath

There are also some utility functions for working with extpaths.

`all_extpaths_from_homology(curve)`

curve	curve Record
description:	Returns a list of extpaths converted from the output of algcurves' `homology` function. This list is suitable for use where a homology basis is needed.

`draw_paths(path)`

path	extpath
description:	Plots the segments of `path` on a graph. Useful for quick visualisation.

`extpath_is_closed(curve, path)`

curve	curve Record
path	extpath
description:	Checks whether a given extpath forms a closed loop on the given surface. Essentially checks that the projection to x is closed and that analytically continuing the y value along this path gives the same sheet at the end. Very vulnerable to floating point errors so beware.

`points_from_extpath(path)`

path	extpath
description:	Essentially the reverse of `extpath_from_points`. Returns a list of the points involved in the given path.

A.4 Automatic PIC handling

There is also an API for writing `.pic` files back.

`read_pic(filename)`

filename	String with name of file to read
description:	Reads a `.pic` file and constructs the necessary expaths.
return value:	sequence: first is curve Record, second is list of extpaths and third is list of names as Maple symbols

---

`write_pic(filename, curve, hom, base=0, sheets_base=1 + i, names=null)`

filename	Name of file to write as `.pic`
curve	curve Record
hom	List of extpaths
base	Monodromy base to use
sheets_base	Base for sheets definition
names	List of Maple symbols to name each extpath. Default will be "hom[i]"
description:	Outputs a `.pic` file that can be read in by CyclePainter.

A.5 Drawing an extpath

CyclePainter is a Java program that interfaces with Maple and can make inputting these paths far less painful. That said, it does have bugs so be careful.

Effectively you create paths with the buttons and entry boxes on the left, and then drag them down to the buttons just below that to activate. When there clicking on the canvas will start (or continue) a path. Right clicking finishes a path. Nodes can be dragged with the mouse, or deleted by selecting them (a single click) and pressing "Delete" on the keyboard. There is no ability to insert a node yet.

The sheet for the initial point can be selected just by the drop-down menu on the left.

Caveats

- Adding a path of the same name as an existing one overwrites it with no confirmation question.

- There may well be many more problems.

It produces a file that can be read directly into Maple, the idiom for converting from this to an extpath based homology basis called `hom` is

```
> curve, hom, names := read_pic("drawn.pic"):
```

It can also produce graphics usable by metapost for drawing the paths.

A.6 Using extpaths

The   fundamental   function   acting   on   extpaths   is   `isect`.

```
isect(curve, path1, path2)
```

`curve`	curve Record
`path1`	First extpath
`path2`	Second extpath
description:	Calculates the intersection number of `path1` and `path2`.

Using this the following can be defined.

```
find_homology_transform(curve, hom1, hom2)
```

curve	curve Record
hom1	List of extpaths giving a homology basis (not necessarily canonical)
hom2	List of extpaths giving a homology basis (not necessarily canonical)
description:	Returns a matrix giving `hom2` in terms of `hom1`. If the paths in `hom1` are γ_i and in `hom2` δ_i then the return value M_{ij} satisfies $\delta_i = M_{ij}\gamma_j$. If both `hom1` and `hom2` are canonical bases then M will be symplectic.
return value:	Matrix

`from_algcurves_homology(curve, hom)`

curve	curve Record
hom	Homology basis
description:	Calculates the transformation which takes the homology basis algcurves chooses to that given by `hom`. A thin wrapper around putting the algcurves homology in the `hom1` slot of `find_homology_transform`.
return value:	Matrix

`periodmatrix(curve, hom)`

curve	curve Record
hom	Homology basis
description:	Calculates the period matrix by transforming Maple's version with appropriate symplectic transform. Thus susceptible to the vagaries of Maple's choice of differential – not only arbitrary but not stable between executions.
return value:	Matrix

A.7 Experimental components

Experimental should probably be read as "Broken", but they exist and can still be useful occasionally.

`transform_extpath(initCurve, path, trans, finCurve := initCurve)`

initCurve	curve Record that `path` lives on

path	extpath
trans	function of two variables returning a list containing two elements. If (x, y) lies on `initCurve` then $trans(x, y) = [w, z]$ and (w, z) is expected to lie on `finCurve`.
finCurve	(Optional) curve Record; image of `trans`.
description:	Not a very reliable function, but can save time even when it fails. There are three essential cases:

- `trans` acts as a simple linear transformation on the x plane (includes common case of changing sheets, i.e. fixing x and just acting on y). The routine should work flawlessly here.

- `trans` is a Mobius transformation on the x plane. Fairly reliable. The straight lines forming extpaths get sent to lines or circles. Probably possible to design input which would fool it though.

- `trans` is more general. Highly unreliable. Straight lines often make tiny loops around branch points in the image; these are rarely detected by the algorithm. Even if `extpath_is_closed` succeeds on the result further tests should be undertaken.

For cases where it fails, the function modifies a global variable `tps`. This is a list of functions. Each function $[0, 1] \to \mathbb{C}$ and represents the image of one segment forming the `path`.

Plotting these with `plots[complexplot]` allows the path to be reconstructed manually.

return value:	extpath: the push-forward of `path` by `trans` from `initCurve` to `finCurve`.

---

reverse_extpath(path)

path	*Closed* extpath.
description:	Reverses the direction `path` is traversed. Assumes path is closed so that a parameter specifying the curve is not needed.
return value:	extpath: `path` in the opposite direction