Trimming

Daryl Cooper and Darren D. Long

CONTENTS

Both authors were supported in part by the NSF.

We consider the question of deciding whether or not certain subgroups of certain cocompact Kleinian groups are geometrically finite or infinite. This is reduced to a question about similarity interval exchange maps. The concept of trimming produces an practical algorithm to settle the latterquestion. This algorithm has been implemented on a computer and the results are perplexing.

In [Cooper et al. 1997], we described an algorithm to determine whether a certain kind of surface in a three-manifold is quasifuchsian or a virtual fibre. The surface in question is required to be transverse to the suspension flow in a closed hyperbolic three-manifold which fibres, but need not be embedded. One of the reasons that this question is interesting is that it is connected to showing that hyperbolic surface bundles have infinite virtual Betti number.

Although the algorithm of [Cooper et al. 1997] is practical in the sense that one can implement it on a variety of examples, such implementations were rather *ad hoc*. In this paper we improve the algorithm somewhat so that it may be more uniformly implemented on a computer. As a consequence we obtain a method which is sufficiently efficient and practical that one can deal with a broader range of examples.

The first step of [Cooper et al. 1997] is to reduce the question to one about similarity interval exchange (SIE) maps; we recall what is relevant of the definition below. To expedite the algorithm, we introduce a concept called *trimming*. This produces from a given geometric SIE, a sequence of SIE's which has the property that after a finite number of trims, either the SIE has a fixed point (from which it follows that the surface is quasi-

1058-6458/1998 $0.50 per page

Fuchsian), or else the sequence is eventually periodic (and the surface is a virtual fibre).

We conclude with examples to which this algorithm was applied; the results suggest that "most" simple immersed surfaces are quasifuchsian. However, any general understanding which would enable one to predict ahead of time which case occurs seems elusive.

1. TRIMMING

The precise definition of an SIE is contained in [Cooper et al. 1997]; roughly speaking it is a generalisation of the more well known notion of an interval exchange map [Masur 1982], except that one fixes a certain number λ and is allowed to stretch and expand intervals by powers of λ.

We may bypass most of the technical considerations here. To this end, we uniquely define a map $\theta : [0, a] \to [0, a]$ by the following data:

- A partition of the domain and range into n closed intervals $D_1, \dots, D_n$ and $R_1, \dots, R_n$.
- A bijection σ from the set of domain intervals to the set of range intervals.
- For each domain interval D_i, a sign $\varepsilon_i = \pm 1$.

This data determines θ by the requirement that θ map a domain interval D_i to the range interval σD_i by a linear bijection which is orientation preserving or reversing depending on whether $\varepsilon_i = \pm 1$.

A special subclass of this type of map arises in the following way. Let S be a closed surface equipped with a pair of oriented transverse singular affinely-measured foliations $\mathcal{F}$ and $\mathcal{G}$ (see [Hatcher and Oertel 1992] for definitions) with finitely many singularities of prong type. Let I be a closed interval lying inside a leaf of $\mathcal{G}$. (In particular this means that it is transverse to $\mathcal{F}$.) We require that both endpoints of I lie on separatrices of $\mathcal{F}$. (This is a technical requirement which is used later.)

Suppose that I is so long that every infinite segment of every leaf of $\mathcal{F}$ hits I. Given a point x on I there is a leaf λ of $\mathcal{F}$ containing x. Then start at x and move along λ in the direction given by the orientation until either λ returns to I at some point y, or else the leaf λ runs into a prong singularity before returning to hit I. In the latter case we say that x is a *break point* of θ. The break points of θ partition the domain into subintervals. The break points of θ^{-1} (obtained by moving backwards along leaves of $\mathcal{F}$) partition the range. If x is not a break point of θ, define $\theta(x) = y$. This defines a map θ on I minus the break points. There is one break point for each separatrix which is oriented to point towards a singularity. Exactly half the separatrices point towards a singularity and half point away from a singularity. We thus see that

$$\#\text{ of domain intervals} = 1 + \tfrac{1}{2}(\#\text{ of separatrices}).$$

We parametrise the interval I as $[0, a]$ using the transverse affine-measure. This means that in the universal cover of S the parametrization of every lift of I induces a measure on I which is a scalar multiple of the transverse measure on I given by the measured foliation $\tilde{\mathcal{F}}$. Then the map $\theta : I \to I$ constructed above is an *SIE*. In what follows we restrict attention to maps constructed in this way.

A generalization of the idea of break point will be used later. We say that a point x on I is a *potential end point* if moving either forwards or backwards along the leaf λ starting at x one gets to a singularity before hitting any point on the left hand subinterval $[0, x]$ of I. Thus break points of both θ and θ^{-1} are potential end points of I.

Suppose that one now shrinks the interval I by decreasing a a small amount to obtain a subinterval J in such a way that leaf rays always return to J. Then the SIE θ_J obtained for J is said to be *induced* from θ by restriction to the subinterval J. In fact the map θ_J is determined solely by θ and does not depend on which foliation $\mathcal{F}$ produces θ. (A referee pointed out that the passage θ to θ_J is a special case of Rauzy induction; see [Veech 1982] or [Masur 1982].)

We will call an SIE *finite* if there is a periodic point, and *infinite* if there is a linear map

$$\xi : [0, a] \to [0, \xi(a)]$$

where $\xi(a) < a$ and for which θ induces an SIE $\bar{\theta}$ on $[0, \xi(a)]$ which is conjugate via ξ to θ. In [Cooper et al. 1997], we showed:

Theorem 1.1. *If θ is an SIE coming from an immersed surface S then θ is finite or infinite depending on whether the surface S is geometrically finite or infinite.*

We will now describe a process called *trimming* which produces from θ a new SIE $\bar{\theta}$ defined on a subinterval $[0, \bar{a}]$ of $[0, a]$. This new SIE is the SIE induced on the subinterval $[0, \bar{a}]$ by θ. The partitions of the domain and range of θ are

$$0 = d_0 < d_1 < d_2 < \cdots < d_n = a$$

and

$$0 = r_0 < r_1 < r_2 < \cdots < r_n = a.$$

The inverse of an SIE is also an SIE, so that after replacing θ by its inverse if necessary, we may assume that $d_{n-1} \le r_{n-1}$. If $\sigma(D_n) = R_n$ then θ is finite so we may suppose that $\sigma(D_n) = R_m$ with $m < n$. We define $\bar{a} = r_{n-1}$, and the new SIE $\bar{\theta}$ as follows.

Case 1: $d_{n-1} < r_{n-1}$ (see Figure 1, left). The partition $\bar{D}_1, \ldots, \bar{D}_n$ of the new domain $[0, \bar{a}]$ consists of the intersections of the domain intervals D_i with $[0, \bar{a}]$. Thus for $i < n$ we have $\bar{D}_i = D_i$ and $\bar{D}_n = [d_{n-1}, \bar{a}]$. The partition of the new range consists of the intervals $R_1, \ldots, R_m^1, R_m^2, \ldots, R_{n-1}$ which are the same as intervals in the range of θ except that the last interval has been ommitted and the m-th interval R_m has been subdivided into two subintervals R_m^1, R_m^2 where $R_m^1 = \theta(\bar{D}_n)$ and R_m^2 is the (closure of) the rest of R_m. The induced SIE $\bar{\theta}$ equals θ everywhere except on $D_p \equiv \sigma^{-1}(R_n)$ on which it equals $\theta \circ \theta$.

Case 2: $d_{n-1} = r_{n-1}$ (see Figure 1, right). In this case, the induced SIE has one fewer interval in the domain and range partitions. This may be viewed as a degenerate example of the previous case where the new domain interval $\bar{D}_n$ and new range interval $\bar{R}_m^1$ have zero length. The partition $\bar{D}_1, \ldots, \bar{D}_{n-1}$ of the new domain $[0, \bar{a}]$ consists of the all the domain intervals D_i except the last one. Similarly, the partition $\bar{R}_1, \ldots, \bar{R}_{n-1}$ of the new range consists of the all the range intervals R_i except the last one. The induced SIE $\bar{\theta}$ equals θ everywhere except on $\sigma^{-1}(R_n)$ on which it equals $\theta \circ \theta$.

Theorem 1.2. *If θ is an SIE coming from an immersed surface S, then after a finite number $N(\theta)$ of trims, one obtains an induced SIE $\bar{\theta}$ such that either $\bar{\theta}$ has a fixed point, or else after rescaling the domain of $\bar{\theta}$ to equal the domain of θ then $\theta = \bar{\theta}$.*

Proof. First suppose that θ is infinite. Thus there is a pseudo-Anosov φ of the surface S which preserves a pair of transverse foliations $\mathcal{F}$ and $\mathcal{G}$. By iterating the map φ we may suppose that φ fixes all the singularities and separatrices of both foliations.

By choice the interval I is contained in some leaf of $\mathcal{G}$ and if we also arrange (as we may) that the lefthand endpoint is a singularity of $\mathcal{G}$, we may assume that φ maps I into a subinterval of itself.

We need to show that only finitely many trims are required to get from I to $\varphi(I)$. Let $x_1, x_2, \ldots$ be the sequence of right hand end points of intervals obtained by trimming I. We claim that these are potential end points of I. The proof is by induction. The first trim of I produces x_1 which is an endpoint of either a domain or a range subinterval. Thus x_1 is a break point of I hence a potential end point. The argument works for further trims with the following observation: x is a potential end point of I if and only if it is a break point of the subinterval $[0, x]$ with the induced SIE.

This is a decreasing sequence of points on I and we wish to show that they converge to 0. Suppose that they converge to $\varepsilon > 0$. There are finitely many separatrices, therefore there is some separatrix λ which intersects I in infinitely many of these potential end points. There is a uniform lower bound to the distance measured along λ between successive hits with I. Hence the sequence of potential end points which lie on λ are not contained in any compact subset of λ. This implies that λ contains no point of $[0, \varepsilon]$, because if y were such

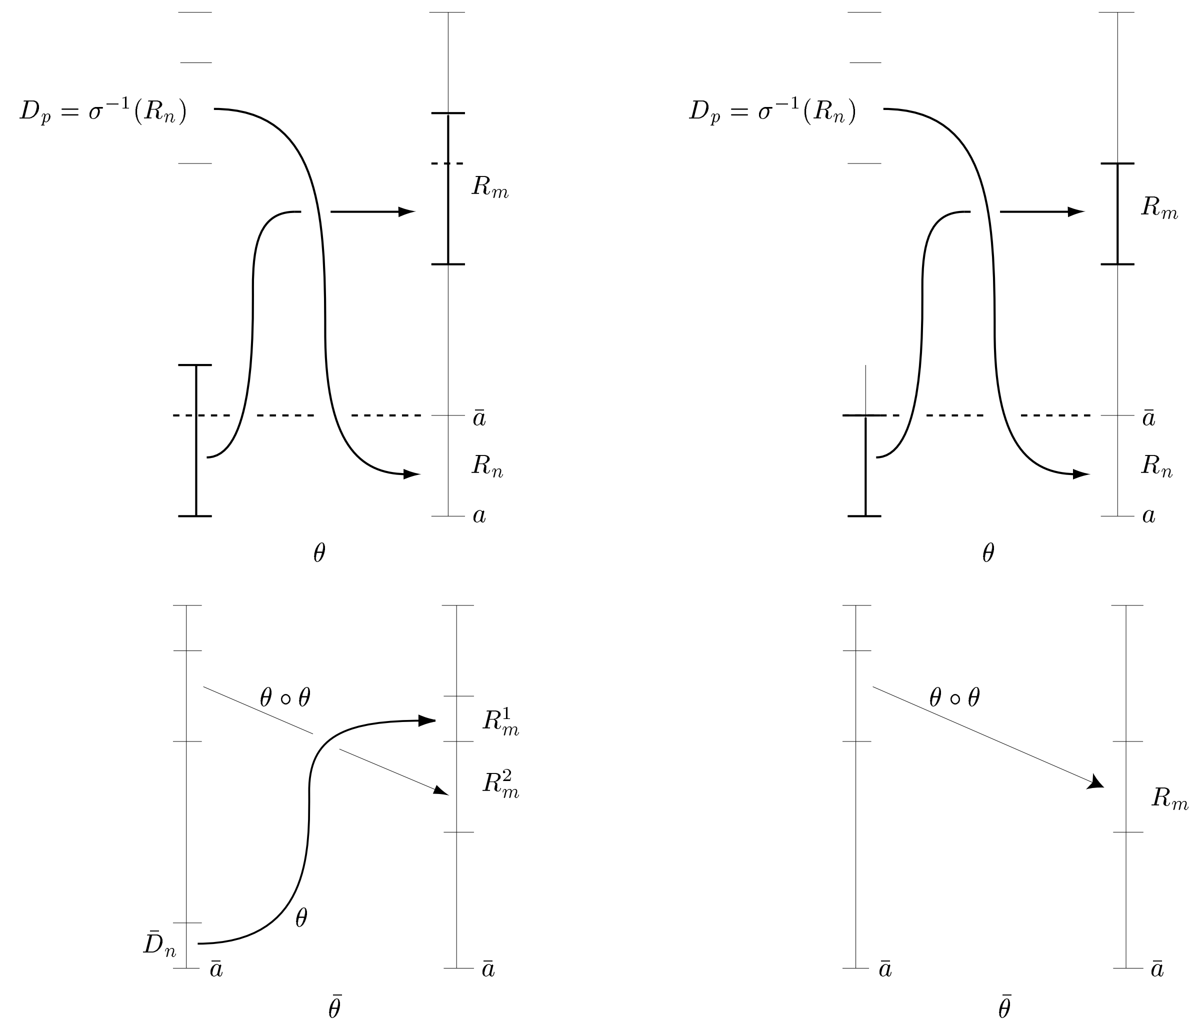

FIGURE 1. Trimming. Left: Case 1. Right: Case 2.

a point, then all the potential end points which lie on λ must be between the singularity and y.

Now every infinite segment of every leaf of $\mathcal{F}$ is dense in S because this foliation is invariant under a pseudo-Anosov. Thus λ hits $[0, \varepsilon]$ infinitely often, for each $\varepsilon > 0$. This is a contradiction.

Notice that we have arranged that the righthand endpoint of I is a potential end point and that it follows from the definition that the φ image of a potential end point is a potential end point. Thus the righthand endpoint of $\varphi(I)$ is a potential end point. Moreover, the argument of the above paragraph shows that the only accumulation point possible for potential end points in I is the left hand endpoint. Since trimming removes one potential end point at a time, we deduce that after a finite number of trims we get from I to $\varphi(I)$.

Now suppose that θ is finite. We need to show that after finitely many trims the induced SIE $\bar{\theta}$ has a fixed point. The analysis again proceeds by looking at the accumulation points of the potential end points. However, this time each separatrix λ spirals in towards some closed leaf of $\mathcal{F}$ (see [Cooper et al. 1994]). Thus the accumulation points are a subset of the intersections of the closed leaves of $\mathcal{F}$ with I.

A closed leaf γ of $\mathcal{F}$ in general meets I in finitely many points $z_1, \ldots, z_n$ which we label in increasing order along I. We claim that the only point in this set which is an accumulation point of potential end points is z_1. This is because if λ meets I at a point w near z_i with $i > 1$, then as λ spirals away from γ towards the singularity, it meets I at a point near z_1 and thus w is not a potential end point. It is also clear from this description that the intersections of λ sufficiently close to z_1 are all potential end points, and thus z_1 is an accumulation point of potential end points.

Thus trimming shortens I so that the right hand endpoint y of I converges to a point z_1 on a closed leaf γ. After finitely many trims, the only point of intersection of γ with I is z_1 therefore the induced SIE on I fixes z_1. □

This theorem gives rise to the following algorithm. Given an SIE, check if it has a fixed point. If not, then trim it. Rescale the new SIE and check if it equals the original SIE. If not, check the new SIE for a fixed point. If it does not have one, then trim it. Rescale the new SIE and check if it equals the original SIE. Repeat.

Checking whether an SIE has a fixed point is tantamount to looking at each of the finitely many linear maps which make up the SIE, and for each linear map with dilatation not equal to 1 calculating the unique fixed point by solving a linear equation. One must then check if this fixed point lies in the interval on which this linear map equals the SIE. This involves comparing the magnitude of the fixed point with the endpoints of the interval. All these numbers lie in a number field over $\mathbb{Q}$, so these checks are algorithmic. See [Cooper et al. 1997].

This algorithm can be implemented on a computer. We give a few comments that will improve the efficiency. The calculation of the initial SIE can take some time. This is because if the initial interval I is short, a leaf of $\mathcal{F}$ may wander around the surface for a long time between hits with I. Thus it is a good idea to start with a long interval I. In fact the initial interval can be replaced by a strategically chosen finite collection of initial intervals. The initial SIE is then defined on this finite set of intervals. One then trims until either a fixed point is found, or else the trimming produces a subinterval of the first interval. Only after this point does one start checking for an infinite SIE.

The check for a fixed point is slow compared to doing a trim. Thus one does many trims between checks for a fixed point. This is justified, for if an SIE has a fixed point, then doing a trim produces a new SIE with a fixed point. Similarly, rescaling an SIE to check for the infinite situation is also slow. However if one first checks to see if the permutation associated to the SIE's are equal, and only rescales in this situation, then much time is saved.

The basic process of doing one trim is a number of computer operations which depends only on the magnitudes of the integers involved in describing the elements in the algebraic number field which determine the SIE. It does not depend on the topological complexity of the surface. In particular, it is especially easy for quadratic fields and we have implemented this algorithm on a micro-computer for surfaces of genus up to a hundred in this case.

2. SOME EXAMPLES

We briefly describe two families of examples based on the simple construction of [Cooper et al. 1994], to which we refer for more details. We recall briefly that the polygon of Figure 2 with the identifications shown, defines a genus two surface. If we regard each square in the picture as having unit side length, we may then define a branched covering of this surface to a torus. Then we may define maps $T_{\mathcal{M}}$ covering $\left(\begin{smallmatrix}1&2\\0&1\end{smallmatrix}\right)$ and $T_{\mathcal{L}}$ covering $\left(\begin{smallmatrix}1&0\\-2&1\end{smallmatrix}\right)$ on this torus.

In this way, we obtain a family θ_k of pseudo-Anosov maps which carry the curve C shown below to the curve $\theta_1(C) = \theta_k(C)$ given by

$$\theta_k = T_{\mathcal{L}}^{-1} T_{\mathcal{M}}^{k},$$

which covers the product

$$\begin{pmatrix}1 & 0\\ 2 & 1\end{pmatrix}\begin{pmatrix}1 & 2k\\ 0 & 1\end{pmatrix} = \begin{pmatrix}1 & 2k\\ 2 & 1+4k\end{pmatrix}.$$

Our convention is that $\theta_1 = \theta$. The slopes of the invariant foliations in these examples is given by $1 \pm \sqrt{k^2+k}/k$.

The surfaces in these examples are constructed by taking several copies of this basic surface and crossjoining various levels along annuli in a pattern described by a picture. (See Figure 3.) We recall that a *crossjoin* involves removing a pair of annuli from the surface and regluing the resulting boundary components in a certain way so as to give an immersion of the surface into each of the mapping tori $M(\theta_k)$ which is transverse to the foliation by lines coming from the mondromy. Such an immersion is incompressible; see [Mangum 1998] or [Cooper et al. 1994].

The annuli carry labels to indicate which curve is involved; usually C or $\theta(C)$. However, for reasons to do with the way the trimming program handles triangulations, it's necessary to use curves which are parallel to these; for example the curve labelled $\theta(C)^{||}$ is the obvious curve in the flat structure parallel to $\theta(C)$. See Figure 2.

To encode the series of examples with this surface we need to describe the surface; this is done by specifying curves and the levels which they join. If a *pair* of curves is specified, one will be the θ image of the other and this crossjoin runs around the circle direction. (In the context of these simple examples we only use the annulus joining C to $\theta_k(C) = \theta(C)$.) To explain our notation consider the surface presented schematically in Figure 3.

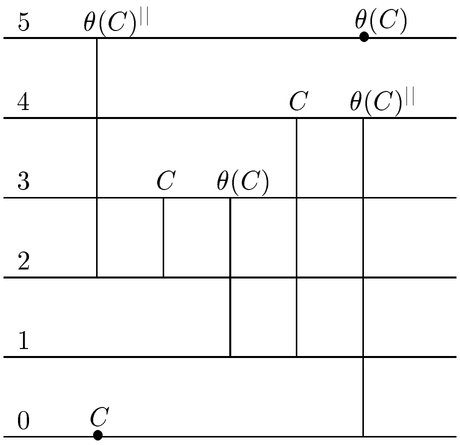

FIGURE 3. Schematic for cross-joined surfaces.

Each horizontal line represents a copy of the surface, in this case there are six copies, labelled $0, \ldots, 5$. Each vertical line represents a crossjoin along an annulus, where the curve in question is given by the label on the line. The symbol $\bullet$ is reserved for the crossjoin around the circle direction. The surface above will be coded thus:

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,4)$, $(2,5)$
C	$(1,4)$, $(2,3)$
$\theta(C)$	$(1,3)$
$(C, \theta(C))$	$(0,5)$

Using the program, we examined all values of $1 \le k \le 64$; such surfaces are geometrically infinite precisely for $k = 4, 30, 34, 60, 64$.

Various examples of this nature, together with the values of k when they are geometrically finite, are recorded below.

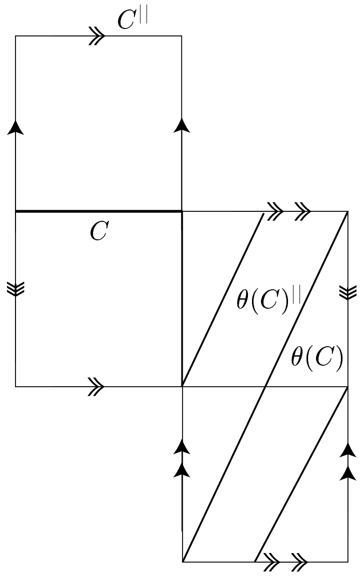

FIGURE 2. The fibre surface.

Example 2: Number of sheets = 7.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,5)$
$C^{\vert\vert}$	$(0,1)$
C	$(1,5)$, $(2,4)$, $(3,6)$
$\theta(C)$	$(1,4)$, $(2,3)$
$(C,\theta(C))$	$(0,6)$

For $1 \leq k \leq 45$, the surfaces are geometrically infinite precisely when k is even.

Example 3: Number of sheets = 9.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,7)$
C	$(1,7)$, $(2,6)$, $(3,5)$, $(4,8)$
$\theta(C)$	$(1,6)$, $(2,5)$, $(3,4)$
$(C,\theta(C))$	$(0,8)$

For $1 \leq k \leq 25$, the surfaces are geometrically infinite precisely when k is even.

Example 4: Number of sheets = 9.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,7)$
$C^{\vert\vert}$	$(0,1)$
C	$(1,7)$, $(2,6)$, $(3,4)$, $(5,8)$
$\theta(C)$	$(1,6)$, $(2,3)$, $(4,5)$
$(C,\theta(C))$	$(0,8)$

For $1 \leq k \leq 80$, the surfaces are geometrically infinite precisely when k is a multiple of 12.

Example 5: Number of sheets = 9.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,7)$
$C^{\vert\vert}$	$(0,2)$
C	$(1,7)$, $(2,6)$, $(3,4)$, $(5,8)$
$\theta(C)$	$(1,6)$, $(2,3)$, $(4,5)$
$(C,\theta(C))$	$(0,8)$

For $1 \leq k \leq 30$, the surfaces are geometrically infinite precisely when k is a multiple of 4.

Example 6: Number of sheets = 9.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,7)$
$C^{\vert\vert}$	$(0,3)$
C	$(1,7)$, $(2,6)$, $(3,4)$, $(5,8)$
$\theta(C)$	$(1,6)$, $(2,3)$, $(4,5)$
$(C,\theta(C))$	$(0,8)$

For $1 \leq k \leq 90$, these surfaces are all geometrically finite.

Example 7: Number of sheets = 11.

Curve(s)	Crossjoin
$\theta(C)^{\vert\vert}$	$(0,9)$
C	$(1,9)$, $(2,8)$, $(3,7)$, $(4,6)$, $(5,10)$
$\theta(C)$	$(1,8)$, $(2,7)$, $(3,6)$, $(4,5)$
$(C,\theta(C))$	$(0,10)$

For $1 \leq k \leq 18$, the surfaces are geometrically infinite precisely when k is even.

Conjecture. The sequence behaviour illustrated by each example persists as k tends to infinity.

For example, the conjecture implies that in Example 4 above, the surface described is geometrically infinite if and only if k is a multiple of 12.

Assuming the truth of the conjecture just posed, we may explain these examples as follows. Varying k amounts to composing the original monodromy θ with powers of some Dehn twists on the one-dimensional submanifold $\mathcal{M}$ in the fibre. In turn, this means that we can view each surface as an immersed incompressible surface S in the hyperbolic manifold with cusps obtained by deleting $\mathcal{M}$ from the bundle $M(\theta)$. The mapping torus $M(\theta_k)$ is obtained by doing certain Dehn surgeries on this manifold with cusps.

The surface S is geometrically finite in the hyperbolic structure on $M(\theta) - \mathcal{M}$, but has accidental parabolics. The behaviour above (modulo the conjecture) shows that the geometrically finite representation of S is a limit of both geometrically infinite and quasifuchsian surface groups.

Permutation-Type Surfaces

Using the base surface and monodromies as described above we can give a different construction. In this case surfaces are all crossjoined along an alternating sequence beginning $\theta(C)^{||}$, $\theta(C)$, ..., together with a single annulus $C, \theta(C)$ around the fibre direction. In addition, another series of annuli is then added all along curves $C^{||}$ according to the following scheme. In each level remove an annular neighbourhood of the curve $C^{||}$. One then writes down a permutation to describe how these annuli are reglued. The permutation is described by mapping (rather than cycle) notation. The example in Figure 4 is coded as $(2, 3, 1, 0)$.

Geometrically infinite surfaces occur in this family, but seem rarer. In fact for $1 \le k \le 30$ there are precisely three geometrically infinite four sheeted surfaces: All have $k = 1$ and are (2,0,1,3), (2,0,3,1), (2,1,0,3).

Using random methods (luck), we found three geometrically infinite six-sheeted surfaces for $k = 2$: (1,0,3,2,5,4), (1,2,3,0,5,4), (1,3,0,2,5,4).

REFERENCES

[Cooper et al. 1994] D. Cooper, D. D. Long, and A. W. Reid, "Bundles and finite foliations", *Invent. Math.* **118**:2 (1994), 255–283.

[Cooper et al. 1997] D. Cooper, D. D. Long, and A. W. Reid, "Finite foliations and similarity interval exchange maps", *Topology* **36**:1 (1997), 209–227.

[Hatcher and Oertel 1992] A. Hatcher and U. Oertel, "Affine lamination spaces for surfaces", *Pacific J. Math.* **154**:1 (1992), 87–101.

[Mangum 1998] B. S. Mangum, "Incompressible surfaces and pseudo-Anosov flows", *Topology Appl.* **87**:1 (1998), 29–51.

[Masur 1982] H. Masur, "Interval exchange transformations and measured foliations", *Ann. of Math.* (2) **115**:1 (1982), 169–200.

[Veech 1982] W. A. Veech, "Gauss measures for transformations on the space of interval exchange maps", *Ann. of Math.* (2) **115**:1 (1982), 201–242.

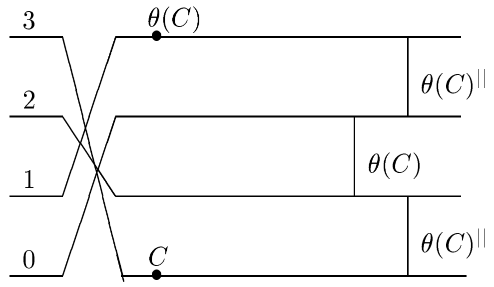

FIGURE 4. Permutation-type cross-joined surface.

Daryl Cooper, Department of Mathematics, University of California, Santa Barbara, CA 93106, United States (cooper@math.ucsb.edu)

Darren D. Long, Department of Mathematics, University of California, Santa Barbara, CA 93106, United States (long@math.ucsb.edu)

Received July 28, 1997; accepted in revised form February 19, 1998