AUTOMORPHISMS OF THE THREE-TORUS PRESERVING A GENUS-THREE HEEGAARD SPLITTING

JESSE JOHNSON

AUTOMORPHISMS OF THE THREE-TORUS PRESERVING A GENUS-THREE HEEGAARD SPLITTING

JESSE JOHNSON

The mapping class group of a Heegaard splitting is the group of connected components in the set of automorphisms of the ambient manifold that map the Heegaard surface onto itself. For the genus-three Heegaard splitting of the 3-torus, we find an eight element generating set for this group. Six of these generators induce generating elements of the mapping class group of the 3-torus and the remaining two are isotopy trivial in the 3-torus.

1. Introduction

Given a 3-manifold M and a Heegaard splitting (Σ, H_1, H_2) of M, consider the set $\mathrm{Aut}(M)$ of orientation preserving automorphisms $M \to M$. The set of connected components of $\mathrm{Aut}(M)$ forms a group $\mathrm{Mod}(M)$ called the mapping class group. We will define $\mathrm{Aut}(M, \Sigma)$ to be the subset of $\mathrm{Aut}(M)$ consisting of maps that send Σ onto itself. The set of connected components of $\mathrm{Aut}(M, \Sigma)$ again forms a group, which we will denote $\mathrm{Mod}(M, \Sigma)$.

The 3-torus $T^3 = S^1 \times S^1 \times S^1$ is known to have a unique (up to isotopy) genus-three Heegaard splitting (Σ, H_1, H_2) [Boileau and Otal 1990]. The mapping class group $\mathrm{Mod}(T^3)$ is isomorphic to $\mathrm{SL}(3, \mathbb{Z})$. However, in $\mathrm{Mod}(T^3, \Sigma)$ there are automorphisms that are nontrivial on Σ but are isotopy trivial in T^3. Thus $\mathrm{Mod}(T^3, \Sigma)$ is in some sense much larger than $\mathrm{Mod}(T^3)$. We will prove the following:

Theorem 1. *For (Σ, H_1, H_2) a genus-three Heegaard splitting of T^3, the group* $\mathrm{Mod}(T^3, \Sigma)$ *is generated by the automorphisms* $\alpha_{12}, \alpha_{21}, \alpha_{13}, \alpha_{31}, \alpha_{23}, \alpha_{32}, \sigma$ *and* τ *defined in Sections 2 and 3.*

This is the first nontrivial example of a genus-three Heegaard splitting of an irreducible Heegaard splitting for which a finite generating set can be explicitly described. Goeritz [1933] found a finite generating set for the mapping class group

of the genus-two Heegaard splitting of S^3. Scharlemann [Scharlemann 2004] recently published a new proof of this result after discovering that the two purported proofs of the higher-genus cases are fatally flawed. Akbas [2008] later found a finite presentation for this group and shortly afterwards, Cho [2008] presented a new proof that this presentation is correct.

For genus-one Heegaard splittings of lens spaces, the mapping class group is finite and easy to understand. For minimal Heegaard splittings of connect sums of $S^1 \times S^2$, the mapping class group of the Heegaard splitting is isomorphic to the mapping class group of a handlebody, and is thus understood. These two classes and the genus-two Heegaard splitting of S^3 are the only previously understood examples.

The problem of understanding mapping class groups of Heegaard splittings is equivalent to problems in algebra and geometry: Algebraically, the mapping class group of (M, Σ) is the intersection of the subgroups of $\text{Mod}(\Sigma)$ that extend into the two handlebodies or that take one handlebody to the other. These two subgroups are conjugate (by the gluing map) and each is finitely generated, but from the algebraic viewpoint there appears to be no general method for calculating their intersection.

Geometrically, for genus greater than two, the mapping class group is the group of automorphisms of the curve complex preserving two handlebody sets. Because the large scale geometry of the curve complex is understood, this point of view is useful for high distance Heegaard splittings. (See [Namazi 2004].) However, the local geometry of the curve complex is not well behaved, making this a difficult problem for low distance Heegaard splittings. By appealing to the topology of T^3, we can solve this problem for this one case. It seems that in general, solving this problem in the 3-manifold setting should be more reasonable than the equivalent problems in algebra and geometry.

2. The mapping class group

A *Heegaard splitting* for a 3-manifold M is a triple (Σ, H_1, H_2) where $H_1, H_2 \subset M$ are handlebodies (connected manifolds homeomorphic to closed regular neighborhoods of graphs in S^3) and Σ is a compact, connected, closed and orientable surface embedded in M such that $H_1 \cup H_2 = M$ and $\partial H_1 = \Sigma = \partial H_2 = H_1 \cap H_2$.

As defined above, $\text{Mod}(M, \Sigma)$ is the group of equivalence classes of orientation preserving automorphisms of M that take Σ onto itself. Two automorphisms are equivalent if there is an isotopy from one to the other by automorphisms of M that take Σ onto itself.

Each connected component of $\text{Aut}(M, \Sigma)$ is a subset of a connected component of $\text{Aut}(M)$, so the inclusion map in $\text{Aut}(M)$ determines a canonical homomorphism $i : \text{Mod}(M, \Sigma) \to \text{Mod}(M)$. In other words, i is induced by "forgetting" Σ and

considering each element of $\mathrm{Mod}(M, \Sigma)$ as an automorphism of M. For further discussion of this homomorphism, see [Johnson and Rubinstein 2006]. The kernel of i is the subgroup of $\mathrm{Mod}(M, \Sigma)$ consisting of automorphisms that are isotopy trivial on M but whose restrictions to Σ are not isotopy trivial. (If we include orientation reversing automorphisms of M, then there is one exception to this fact: any minimal-genus Heegaard splitting of a connect sum of $S^1 \times S^2$s will have an automorphism that interchanges the handlebodies, but restricts to the identity on the boundary.)

Consider $\mathbb{R}^3$ with axes labeled x_1, x_2, x_3. Let T_1, T_2, T_3 be isometries of $\mathbb{R}^3$ where T_i is translation by 1 unit along the axis x_i. We can think of T^3 as the quotient of $\mathbb{R}^3$ by the group generated by T_1, T_2 and T_3. Each automorphism of T^3 lifts to an automorphism of $\mathbb{R}^3$. Within the isotopy class for this automorphism, there is a representative that lifts to $\mathbb{R}^3$ such that the automorphism of $\mathbb{R}^3$ fixes the origin. It sends the vectors $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ to integral vectors v_1, v_2, v_3, respectively. Thus an automorphism of T^3 determines an element of the matrix group $\mathrm{GL}_3(\mathbb{Z})$.

The matrix determined by an automorphism is unique and an element of $\mathrm{GL}_3(\mathbb{Z})$ is represented by an automorphism of T^3 if and only if its determinant is one. Thus $\mathrm{Mod}(T^3)$ is isomorphic to the group $\mathrm{SL}_3(\mathbb{Z})$ of 3×3 integral matrices with determinant one. This group is generated by the six automorphisms that send x_i to $x_i + x_j$ for $i \neq j$. Let A_{ij} be this automorphism for each distinct pair i, $j \in \{1, 2, 3\}$.

We would like to construct a Heegaard splitting for T^3 that fits naturally into the picture of the 3-torus described above. Let $K_1 \subset T^3$ be the image in T^3 of the three edges in $\mathbb{R}^3$ from the origin to $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$, respectively. This K_1 is a graph with a single vertex and three edges. Let K_2 be the image in T^3 of the same three edges, translated by the vector $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. This is again a graph with one vertex and three edges.

Let H_1 be the set of points in T^3 whose distance to K_1 (in the Euclidean metric on T^3) is greater than or equal to their distance to K_2. Let H_2 be the set of points closer to K_2 and let Σ be the set of points equidistant to K_1 and K_2. Each of H_1, H_2 is the closure of a regular neighborhood of K_1, K_2, respectively so each set is a handlebody. Moreover, Σ is the boundary of each handlebody, so (Σ, H_1, H_2) is a (genus-three) Heegaard splitting for T^3, shown on the right. (Gluing opposite faces of the cube by translations forms T^3. The image in T^3 of the surface shown is the Heegaard splitting Σ. It is drawn here as a smooth surface, though the way it's defined it is piecewise linear.) Boileau and Otal [1990] showed that two Heegaard splittings of T^3 are isotopic if they have the same genus, so (Σ, H_1, H_2) is the *standard* (up to isotopy) genus-three Heegaard splitting of T^3.
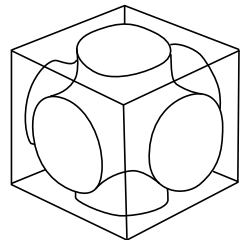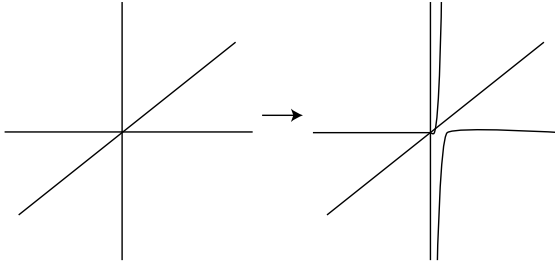
Figure 1. The move α_{ij} replaces the edge along the x_i axis with an edge that follows the x_j axis, then the x_i axis.

Orient each edge of K_1 so that it points in the increasing direction along its axis. Choose distinct i, j, $k \in \{1, 2, 3\}$. Applying the transformation A_{ij} to T^3 sends the edges of K_1 along the x_j and x_k axes onto themselves and sends the edge along the x_i axis to the diagonal of a square in the x_i-x_j plane. Isotope this diagonal across the square into a neighborhood of the original spine so that it first passes along the edge in the x_j axis and then along the edge in the x_i axis. The resulting spine is shown in Figure 1.

There is a unique element α_{ij} of $\mathrm{Mod}(T^3, \Sigma)$ that maps to A_{ij} in $\mathrm{Mod}(T^3)$ and sends K_1 onto the new spine defined above. If we had chosen to slide the diagonal edge in the opposite direction across the square, we would have gotten a different element of $\mathrm{Mod}(T^3, \Sigma)$. Because each α_{ij} maps to A_{ij}, the images in i of $\{\alpha_{ij}\}$ generate $\mathrm{Mod}(T^3)$. In order to extend this set of elements to a generating set for $\mathrm{Mod}(T^3, \Sigma)$, we must understand the kernel of i.

3. The kernel

In this section, we will define elements σ and τ of $\mathrm{Mod}(T^3, \Sigma)$ that are nontrivial in $\mathrm{Mod}(T^3, \Sigma)$, but isotopy trivial in $\mathrm{Mod}(T^3)$. First consider the translation of $\mathbb{R}^3$ by the vector $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$. This descends to an automorphism σ of T^3 that sends the handlebody H_1 onto H_2 and sends H_2 onto H_1. The translation of $\mathbb{R}^3$ is isotopic to the identity by a family of translations, inducing an isotopy of T^3 taking σ to the identity. Thus σ is in the kernel of i. Note that there are many automorphisms that are isotopic to the identity and swap the two handlebodies. For our purposes, we could have picked any of these, but σ happens to be the easiest to define. (Note that σ has order two in $\mathrm{Mod}(T^3, \Sigma)$.)

The second automorphism in the kernel that we will define is called a torus twist and can be defined on Heegaard splittings in a large family of manifolds. Let $D_1 \subset H_1$ and $D_2 \subset H_2$ be properly embedded, essential disks such that the intersection of ∂D_1 and ∂D_2 is exactly two points. We will assume that the algebraic intersection number of ∂D_1 and ∂D_2 is zero, though this is not necessary in
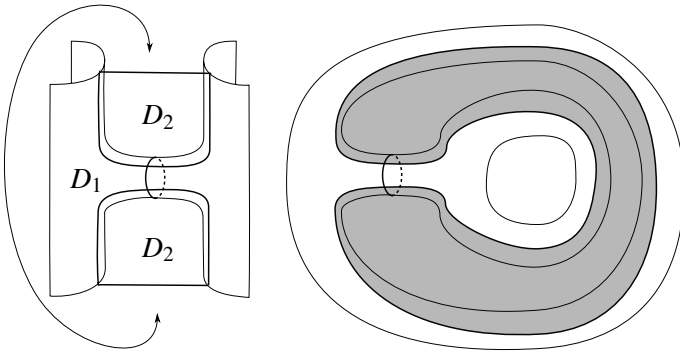
Figure 2. A neighborhood of the two disks is a solid torus.

general. In [Johnson and Rubinstein 2006] showed that D_1 and D_2 determine an element of the kernel of i as follows:

Let N be the closure of a regular neighborhood of $D_1 \cup D_2$. Because there are two points of intersection, N is a solid torus, as in Figure 2. (On the left, the disk D_2 is shown cut in half, at the top and bottom of the figure.) Because the orientations at the two intersections are opposite, the surface $\Sigma \cap N$ is a four punctured sphere whose boundary consists of four simple closed curves in ∂N.

The loops $\Sigma \cap \partial N$ are parallel longitudes in ∂N. There is a one-parameter family of automorphisms of the torus, beginning and ending at the identity, such that each automorphism takes each longitude onto itself and such that the image of any point in the longitude under this family goes around the loop once. This family defines an homeomorphism from $T^2 \times [0, 1]$ onto itself that is the identity on its boundary. A *torus twist* is a homeomorphism of the 3-torus consisting of this map on a regular neighborhood of ∂N and the identity everywhere else. Because N is a solid torus, such an automorphism is isotopic to the identity. This is, in some sense, a 3-dimensional version of a Dehn twist. A torus twist around the ∂N constructed above induces an automorphism of Σ consisting of Dehn twists along loops parallel to $\Sigma \cap \partial N$.

A torus twist is defined by the solid torus N constructed from a pair of disks D_1, D_2. If τ is a torus twist and ϕ is any automorphism of (M, Σ) then $\phi^{-1}\tau\phi$ is a torus twist around $\phi(N)$, which is defined by the disks $\phi(D_1), \phi(D_2)$. We will choose a specific pair of disks and show that the kernel of i is generated by torus twists that are conjugates of this fixed torus twist.

For each $i \in \{1, 2, 3\}$, the image in T^3 of the plane $\{(x_1, x_2, x_3) \mid x_i = 0\}$ intersects H_2 in a properly embedded, essential disk D_2^i and the image of the plane $\{(x_1, x_2, x_3) \mid x_i = \frac{1}{2}\}$ intersects H_1 in a disk D_1^i. For each i, j the disks D_1^i and D_2^j are disjoint when $i = j$ and intersect in two points when $i \neq j$. When $i \neq j$, a regular neighborhood of $D_1^i \cup D_2^j$ is a solid torus parallel to the x_k axis ($k \neq i, j$)
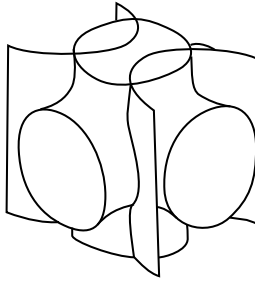
Figure 3. The torus defining an automorphism τ_{ij}.

and we can twist in the positive direction along this axis. Thus for each choice of distinct i and j, D_1^i and D_2^j define a torus twist τ_{ij}. One of these is shown in Figure 3.

For $i \neq j$, the boundary of a regular neighborhood of $D_1^i \cup D_2^j$ intersects Σ in loops parallel to the loops of intersection with the boundary of a neighborhood of $D_1^j \cup D_2^i$. However, twisting in the positive direction around the first solid torus induces oppositely oriented Dehn twists along these loops than twisting in the positive direction along the second solid torus. Thus $\tau_{ij} = \tau_{ji}^{-1}$ on Σ. For our last generator, define $\tau = \tau_{12}$.

4. Torus twists

Let $G \subset \mathrm{Mod}(T^3, \Sigma)$ be the subgroup generated by $\{\alpha_{ij}, \sigma, \tau\}$. In order to show that $G = \mathrm{Mod}(T^3, \Sigma)$, we will show first that the kernel of i is generated by a certain class of torus twists and second that these torus twists are all conjugate to τ by elements of G. We will begin by defining some relations on elements of G. Define $r_{ij} = \alpha_{ij} \alpha_{ji}^{-1} \alpha_{ij} \tau_{jk}$. The reader can check the following:

Lemma 2. *Let $i, j, k \in \{1, 2, 3\}$ be distinct. Then the automorphism r_{ij} takes the image in T^3 of the x_i axis onto the image of the j axis and takes the x_j axis onto the negative i axis.*

Because $\tau = \tau_{12} = \tau_{21}^{-1}$, the rotations $r_{31} = r_{13}^{-1}$ and $r_{32} = r_{23}^{-1}$ are in G. Moreover, we have that $\tau_{13} = r_{23} \tau_{12} r_{23}^{-1}$ so τ_{13} is in G, as is r_{12} and therefore τ_{23}. Thus the torus twists $\{t_{ij}\}$ are all contained in G.

Let S be the image in T^3 of the plane $\{(x_1, x_2, x_3) \mid x_3 = 0\}$ in $\mathbb{R}^3$. Let

$$\phi \in \mathrm{Mod}(T^3, \Sigma)$$

be a torus twist defined by disks $D_1 \subset H_1$ and $D_2 \subset H_2$. Let T be the boundary of a regular neighborhood of $D_1 \cup D_2$, that is, the torus that defines the twist. Assume S and T are transverse.

Lemma 3. *If T is a torus defining a torus twist on Σ and $T \cap S$ is connected and essential in T then ϕ is conjugate to τ by elements of G (and thus ϕ is in G).*

Proof. Because the intersection $T \cap S$ is a single loop and this loop is essential in T, the complement $T \setminus S$ is an annulus $A \subset T$. The intersection of T with Σ consists of four loops, each of which intersects A in a properly embedded, essential arc. These arcs cut the twice punctured, genus-two surface $\Sigma \setminus S$ into planar pieces. Because the pieces are planar, the arcs must cut $\Sigma \setminus S$ into two annuli.

The intersection $H_1 \cap S$ is a punctured torus $S' \subset S$. The intersection of T with S' is a pair of properly embedded arcs (with endpoints in $T \cap \Sigma$) that cut S' into a disk and an annulus. There is a spine for S' such that one edge of the spine intersects T in two points and the other edge is disjoint from T. Let γ_1 be the edge that intersects T and γ_2 the edge disjoint from T.

Let γ_1' be the edge of the spine K_1 for H_1 that is the image of the x_1 axis of $\mathbb{R}^3$. Let γ_2' be the edge of K coming from the x_2 axis. These two edges form a spine for S. By applying the generators α_{12}, α_{21}, τ_{13} and τ_{23}, one can send γ_1', γ_2' onto any spine for S'. In particular, one can send these edges onto γ_1 and γ_2. This automorphism sends the torus defined by D_1^1 and D_2^2 onto T so conjugating τ by this automorphism produces ϕ. $\square$

5. Weak reducing disks

Let $D_1 \subset H_1$ and $D_2 \subset H_2$ be properly embedded, essential disks in the standard Heegaard splitting (Σ, H_1, H_2) of T^3.

Definition 4. The disks D_1, D_2 are a *weak reducing pair* if the boundary of D_1 can be isotoped disjoint from ∂D_2.

Each of the pairs of disks D_1^i, D_2^i constructed above is a weak reducing pair.

Lemma 5. *A weak reducing pair of disks determines a unique isotopy class of incompressible tori in T^3.*

Proof. First note that (Σ, H_1, H_2) is irreducible because cutting Σ along a reducing sphere would produce a genus-two Heegaard splitting. This is impossible because $\pi_1(T^3)$ has rank three. Thus the boundaries of D_1 and D_2 cannot be isotopic.

We will see momentarily that both D_1 and D_2 must be nonseparating disks in H_1, H_2, respectively, but for now note that if D_1 is separating, then it cuts H_1 into a genus-two handlebody and a genus-one handlebody. The boundary of D_2 must be contained in the boundary of the genus-two handlebody so a (nonseparating) meridian disk for the genus-one handlebody will be disjoint from D_2. Thus if D_1 is separating then it uniquely determines a nonseparating disk in H_1 disjoint from ∂D_2 and we can replace D_1 with this disk. Likewise, if D_2 is separating then

we can replace D_2 with the unique nonseparating disk in H_2 determined by D_2. Thus we can assume D_1 and D_2 are nonseparating.

Scharlemann and Thompson [1994] showed that given a pair $D_1 \subset H_1$, $D_2 \subset H_2$ of disjoint essential disks, one can form a generalized Heegaard splitting, consisting of four compression bodies G_1, G_2, G_3, G_4 such that $G_1 \cap G_2 = \partial_+ G_1 = \partial_+ G_2$, $G_2 \cap G_3 = \partial_- G_2 = \partial_- G_3$ and $G_3 \cap G_4 = \partial_+ G_3 = \partial_+ G_4$. The thin surface

$$S = \partial_- G_2 = \partial_- G_3$$

is the result of compressing Σ across D_1 and D_2. Since ∂D_1 and ∂D_2 are nonseparating in the genus-three surface Σ, the surface S consists of one or two tori.

By [Scharlemann and Thompson 1994], if S is compressible then the generalized Heegaard splitting can be weakly reduced further, in which case one of the thin levels will contain a sphere. This implies that the original Heegaard splitting is reducible. Because the genus-three Heegaard splitting of T^3 is irreducible, S must consist of incompressible tori.

Because every incompressible torus in T^3 is nonseparating, S must consist of two parallel incompressible tori whose union is separating, though each is nonseparating on its own. There was no choice involved in the construction of S, and the two components of S are isotopic, so D_1 and D_2 define a unique isotopy class of incompressible tori (shown in Figure 4). $\square$
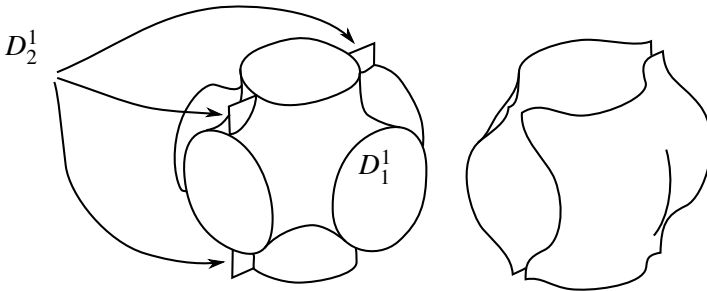

Figure 4. Compressing Σ along a weak reducing pair produces a pair of incompressible tori.

Lemma 6. *If D_1 and D_2 are a weak reducing pair then each of ∂D_1 and ∂D_2 is nonseparating in Σ. If D_2' is an essential, properly embedded disk in H_2 disjoint from ∂D_1 then D_2' is isotopic to D_2. Similarly, if D_1' is essential and properly embedded in H_1 and $\partial D_1'$ is disjoint from ∂D_2 then D_1' is isotopic to D_1.*

In other words, a disk in H_1 or H_2 is half of at most one weak reducing pair. (Some disks are not half of any weak reducing pair.)

Proof. Let D_1 and D_2 be a weak reducing pair for (Σ, H_1, H_2). First assume that each of D_1 and D_2 is nonseparating. As in the proof of Lemma 5, compressing Σ along D_1 and D_2 produces a surface S consisting of two parallel incompressible tori. If we first compress S across D_1 then we get a surface isotopic to the boundary of $H_2 \cup N(D_1)$, where $N(D_1)$ is the closure of a regular neighborhood of D_1.

Compressing $\partial(H_2 \cup N(D_1))$ across D_2 produces a pair of parallel tori that cut T^3 into two pieces homeomorphic to $T^2 \times [0, 1]$. Thus $H_2 \cup N(D_1)$ is homeomorphic to attaching a one-handle to $T^2 \times [0, 1]$ with ends on opposite boundary tori. Any compressing disk for ∂H_2 that is disjoint from D_1 will be a compressing disk for $\partial(H_2 \cup N(D_1))$. Any compressing disk for $\partial(H_2 \cup N(D_1))$ is isotopic to a meridian of the one handle so there is a unique compressing disk for H_2 disjoint from D_1. In other words a compressing disk D_2' will be isotopic to D_2. The same argument with D_1 and D_2 reversed shows that any disk in H_1 disjoint from D_2 is isotopic to D_1.

If D_2 is separating and D_1 is nonseparating then, as noted in the proof of Lemma 5, there is a nonseparating disk D_2' that is disjoint from both D_1 and D_2. Because D_1 and D_2' are disjoint and each is nonseparating, we have just shown that D_2 must be isotopic to D_2'. Thus if D_1 is nonseparating then D_2 must be nonseparating.

Finally, if D_1 is separating then there is a nonseparating disk D_1' disjoint from D_1 and D_2. Because D_1' and D_2 are disjoint and D_1' is nonseparating, D_2 must be nonseparating. Because D_1' and D_2 are disjoint and each is nonseparating, D_1 must be isotopic to D_1', so D_1 is nonseparating. $\qquad\square$

6. Graphics and disks

The key to the proof of Theorem 1 is the following lemma:

Lemma 7. *Let D_1, D_2 and D_1',D_2' be weak reducing pairs. Then the incompressible tori determined by the two pairs are isotopic in T^3 if and only if there is a sequence of torus twists on (Σ, H_1, H_2) taking D_1', D_2' onto D_1, D_2 such that each torus twist is along a torus that cuts Σ into planar surfaces and intersects the incompressible torus defined by D_1, D_2 in a single loop.*

This will be proved in Section 8. The proof uses Morse functions and stable functions on 3-manifolds, which we will describe in this and the next sections.

Let $f : T^3 \to \mathbb{R}$ be a Morse function on T^3 with one index zero critical point at level 0, three index one then three index two critical points (at distinct levels), then one index three critical point at level 1. Let $\pi : T^3 \to S^1$ be a fiber bundle map such that $\pi^{-1}(s)$ is a torus for each $s \in S^1$. We can think of π as a circle valued Morse function on T^3. As was shown in [Kobayashi and Saeki 2000], after an arbitrarily small isotopy the product of two Morse functions on a 3-manifold

M is a stable function from M to $\mathbb{R}^2$. The same argument in this situation implies that after an arbitrarily small isotopy, the product of f and π is a stable function $f \times \pi : T^3 \to [0, 1] \times S^1$.

Assume f and π have been isotoped so that $F = f \times \pi$ is stable. For each $t \in [0, 1]$, define f_t to be the level surface $f^{-1}(t)$. Define π_t to be the restriction of π to f_t. As was shown in [Kobayashi and Saeki 2000], the discriminant set $\mathcal{J}$ of F is a one dimensional submanifold in T^3, whose image $F(\mathcal{J})$ is a finite graph in $[0, 1] \times S^1$, called the *graphic* of f. We will say that F is *generic* if F is stable and for any $s \in S^1$, there is at most one vertex of $F(\mathcal{J})$ in the line $[0, 1] \times \{s\}$.

Because f is Morse, f_t is a surface for all but finitely many values of t. Because F is stable, π_t is a Morse function on f_t for all but finitely many of the regular values of t. If F is generic then at the values of t where π_t fails to be Morse, π_t will have either exactly two critical points at the same level or exactly one degenerate critical point. (If F is stable but not generic, there will be a value of t where there are more than one pairs of critical points at the same level.) Let $\ell_1 \in [0, 1]$ be the level of the highest index one critical point in f and let ℓ_2 be the level of the lowest index two critical points. (By assumption, $\ell_1 < \ell_2$.)

Assume some level set f_t is equal to Σ (as a subset of T^3) for some $t \in (\ell_1, \ell_2)$. Then the surfaces $\{f_t \mid t \in (\ell_1, \ell_2)\}$ determine an isotopy of Σ, inducing a canonical (up to isotopy) identification $c_t : \Sigma \to f_t$ for each $t \in (\ell_1, \ell_2)$.

Lemma 8. *Assume f has the property that for any regular value $t < \ell_1$, the surface f_t is not compressible into $f^{-1}([t, 1])$ and for any regular $t > \ell_2$, the surface f_t is not compressible into $f^{-1}([0, t])$. If $f \times \pi : T^3 \to [0, 1] \times S^1$ is generic then there is a unique (up to isotopy) weak reducing pair D_1, D_2 such that for some $t \in (\ell_1, \ell_2)$, c_t sends ∂D_1 and ∂D_2 to loops isotopic into regular level sets of π_t.*

We should emphasize that there are essentially two parts to the statement: first that for some t, the level sets of π_t determine a weak reducing pair of disks and second that if π_t and $\pi_{t'}$ determine weak reducing pairs, then they determine the same weak reducing pair.

Proof. Let $C_1 \subset [0, 1]$ be the set of points $t \in [0, 1]$ such that an essential level loop of π_t bounds a disk in $f^{-1}([0, t])$. Similarly, let C_2 be the set of points where an essential level loop of π_t bounds a disk in $f^{-1}([t, 1])$. We first will show that $C_1 \cap C_2$ is a nonempty, connected, open interval $(a, b) \subset [0, 1]$, implying that the level loops of π_t for $t \in (a, b)$ contain the boundaries of a weak reducing pair.

For any $\varepsilon > 0$, the set $f([\ell_1 + \varepsilon, \ell_2 - \varepsilon]) \subset M$ is foliated by level surfaces of f and its complement is a pair of handlebodies. Thus there is a sweep-out $f' : M \to [0, 1]$ that agrees with f on the interval $[\ell_1 + \varepsilon, \ell_2 - \varepsilon]$. (Each level set of f' is isotopic to Σ.) Define the sets $C_1', C_2' \subset [0, 1]$ as the values of t for which the level set

of f' intersects a level set of π in a loop bounding a compressing disk for the surface $f'^{-1}(t)$.

Bachman and Schleimer [2005, Claim 6.7] showed that for a sweep-out f' of a surface bundle, the set C_1' is of the form $[0, b)$ for some b and C_2' is of the form $(a, 0]$ for some a. Thus if $a < b$ then $C_1' \cap C_2'$ is an open interval (a, b). Because f agrees with f' on $[\ell_1 + \varepsilon, \ell_2 - \varepsilon]$ for every $\varepsilon > 0$, the intersection of $C_1 \cap C_2$ with (ℓ_1, ℓ_2) is a (possibly empty) open interval.

If for some t and s there are regular loops of $\pi_t(s)$ that are essential in f_t and trivial in $\pi^{-1}(s)$ then t is in either C_1 or C_2. Cooper and Scharlemann [1999] showed that for a torus bundle M with bundle map π, if there is a surface $S \subset M$ such that $\pi|_S$ is Morse and every regular level loop is either trivial in both surfaces or essential in both surfaces then S is an essential torus. If $\pi|_S$ is near Morse (that is, at a crossing of the graphic) and all regular level loops are trivial or essential in both surfaces then S is a strongly irreducible, genus-two Heegaard surface.

Because Σ is not a genus-two surface, every level surface f_t' must have a level loop that is essential in f_t' but trivial in the appropriate level surface of π. Thus $C_1' \cap C_2'$ is a nonempty open interval. The set $C_1 \cap C_2$ is the intersection of $C_1' \cap C_2'$ with $[\ell_1, \ell_2]$ (because f and f' agree on this set) so $C_1 \cap C_2$ is empty if and only if $C_1 \cap (\ell_1, \ell_2)$ or $C_2 \cap (\ell_1, \ell_2)$ is empty.

Assume for contradiction C_1 is disjoint from (ℓ_1, ℓ_2). Let t be a value between ℓ_1 and the last crossing before ℓ_1. The surface f_t has genus two and $\pi|_{f_t}$ is Morse so there is a regular level $s \in S^1$ such that a loop in $\pi_t^{-1}(s)$ is essential in f_t and trivial in $\pi^{-1}(s)$. Thus some essential loop $\gamma \subset f_t^{-1}(s)$ bounds a disk in either $f^{-1}([0, t])$ or $f^{-1}([t, 1])$.

If γ bounds a disk in $f^{-1}([0, t])$ then attaching a one-handle at the critical point at time ℓ_1 does not affect this disk. So γ must bound a disk in $f^{-1}([t, 1])$. This contradicts the assumption that for any regular value $t < \ell_1$, f_t is not compressible into $f^{-1}([t, 1])$. Thus C_1 must intersect (ℓ_1, ℓ_2). A similar argument implies C_2 must intersect (ℓ_1, ℓ_2). Thus $C_1 \cap C_2$ is a nonempty, connected, open interval (a, b).

For $t \in (a, b)$, at least one loop in the pair of pants decomposition bounds a disk in H_1 and at least one bounds a disk in H_2. Because any two loops in the pants decomposition are disjoint, Lemma 6 implies that there is exactly one loop bounding a disk in H_1 and exactly one bounding a disk in H_2.

At each crossing in $(a, b) \times S^1$, f_t passes through a near-Morse function and the induced pants decomposition changes in one of two ways (see [Hatcher and Thurston 1980]): (1) One loop in the pants decomposition may be replaced by a new loop that intersects the original loop in one or two points or (2) two loops that are contained in a twice punctured torus component are simultaneously removed and replaced. After the change, there are still loops bounding disks in opposite handlebodies. The second type of move cannot change the two loops that bound

a weak reducing pair because two such loops are never contained in a twice punctured torus component. Because the loops bounding the disks cannot both change simultaneously, Lemma 6 implies that neither can change. Thus the weak reducing pair is uniquely determined.   □

7. Stable functions in dimension three

We have seen how the graphic defined by a Morse function f for a genus-three Heegaard splitting for T^3 and a torus bundle map π for T^3 determine a weak reducing pair of disks for the Heegaard splitting. An automorphism ϕ of the Heegaard splitting will take this weak reducing pair to a new weak reducing pair, determined by the graphic of the original sweep-out f and a new bundle map $\pi \circ \phi$. If ϕ is isotopy trivial then the isotopy determines a family of bundle maps $\{\pi^t\}_{t \in [0,1]}$ such that $\pi^0 = \pi$ and $\pi^1 = \pi \circ \phi$.

Each bundle map π^t determines a graphic with f. At each value of t where the graphic is generic, π^t determines a pair of weak reducing disks. Thus the family $\{\pi^t\}$ determines a sequence of weak reducing disks for (Σ, H_1, H_2). If we choose this family carefully, we can understand the sequence of weak reducing disks well enough to find a sequence of automorphisms of (Σ, H_1, H_2) taking the original weak reducing pair to each consecutive pair in the sequence. In this section we will describe how the graphic can change during the isotopy of π and in the next section we will show how this corresponds to a sequence of weak reducing disks which suggest a sequence of automorphisms of (Σ, H_1, H_2).

In order to understand how the graphic changes, we will consider an isotopy of f rather than an isotopy of π. Because ϕ is isotopic to the identity, there is a continuous family $\{\phi_t : T^3 \to T^3\}$ such that ϕ_0 is the identity and $\phi_1 = \phi$. The family of bundle maps defined above is given by $\pi^t = \pi \circ \phi_t$ and the graphic at time t is determined by $f \times (\pi \circ \phi_t)$. If we compose the stable function with ϕ_t^{-1}, we find that the graphic is also given by the map $(f \circ \phi_t^{-1}) \times \pi$.

Define $f^t = f \circ \phi_t^{-1}$. Because f^t is a Morse function for each t, there is an open neighborhood $N_t \subset C^\infty(T^3, \mathbb{R})$ (with the Whitney C^∞ topology [Golubitsky and Guillemin 1973]) such that each function in N is isotopic to f^t. Moreover, N_t can be chosen to be convex in $C^\infty(T^3, \mathbb{R})$. Because the set $\{f^t\} \subset C^\infty(T^3, \mathbb{R})$ is compact, it is covered by a finite subset of the convex open neighborhoods $\{N_t\}$. Because each neighborhood is convex and consists of Morse functions isotopic to f, the path f^t can be replaced with a piecewise linear path consisting of arcs connecting consecutive functions $g^0, \dots, g^n \in \{f^t\}$ such that each $g^i \in N_i \cap N_{i+1}$ determines a generic graphic with π and each arc determines an isotopy of g^i onto g^{i+1}. Thus in order to understand how the graphic changes as f^t changes, we can restrict our attention to straight arcs in $C^\infty(T^3, \mathbb{R})$.

Each intermediate function in the arc from g^i to g^{i+1} is of the form $ag^i + bg^{i+1}$ where $a, b > 0$ and $a + b = 1$. By scaling this function as in [Johnson 2009], we can make it of the form $\cos(s)g^i + \sin(s)g^{i+1}$. Thus we are interested in the graphic defined by the stable function $(\cos(s)g^i + \sin(s)g^{i+1}) \times \pi$. This is the projection of the map $g^i \times g^{i+1} \times \pi : T^3 \to \mathbb{R}^1 \times S^1$ onto the annulus $L \times S^1$ where $L \in \mathbb{R}^2$ is the line through the origin with slope $\cos(s)/\sin(s)$. Thus we can understand the effect of the isotopy on the graphic by looking at projections of the map $g^i \times g^{i+1} \times \pi$.

We can choose the sequence $\{g^i\}$ such that each of these maps g^i, g^{i+1}, π, $g^i \times g^{i+1}$, $g^i \times \pi$ and $g^{i+1} \times \pi$ is stable, so each is contained in an open ball of isotopic maps in its respective vector space. The projections of $\mathbb{R}^2 \times S^1$ into the appropriate subspaces define continuous maps from $C^\infty(T^3, \mathbb{R}^3 \times S^1)$ into the spaces containing the above maps and the preimage of each open neighborhood is open in $C^\infty(T^3, \mathbb{R}^2 \times S^1)$. Their intersection is an open set containing $g^i \times g^{i+1} \times \pi$.

Mather [1970] showed that stable functions between three dimensional manifolds are dense in the Whitney C^∞ topology, so this open neighborhood contains a stable function from T^3 into $\mathbb{R}^2 \times S^1$. This stable function projects to maps isotopic to g^i, g^{i+1}, π, $g^i \times g^{i+1}$, $g^i \times \pi$ and $g^{i+1} \times \pi$ so if we replace g^i, g^{i+1} and π with these isotopic maps, $g^i \times g^{i+1} \times \pi$ will be stable and the rest of the maps will be isotopic to the original maps.

Mather's classification of singularities of stable functions between three dimensional manifolds implies that at each point p in T^3, some neighborhood N of p can be parametrized and some open ball in $\mathbb{R}^2 \times S^1$ can be parametrized such that $f|_N$ has one of the following forms:

$$f(x, y, z) = (x, y, z),$$
$$f(x, y, z) = (x^2, y, z),$$
$$f(x, y, z) = (xy + x^3, y, z),$$
$$f(x, y, z) = (xy + x^2z + x^4, y, z).$$

The first type of point on the list is a regular point and at such a point, the discriminant map from $T(T^3)$ to $T(\mathbb{R}^2 \times S^1)$ is one-to-one. At the last three types of points, the discriminant map has a one dimensional kernel, so these points are in the discriminant set. The discriminant set intersects each such neighborhood in an open disk, so the discriminant set of $g_0 \times g_1 \times \pi$ is a compact 2-dimensional submanifold in T^3.

The image of the discriminant set in $\mathbb{R}^2 \times S^1$ is an immersed 2-manifold with "cusps". The cusps in the immersion of the 2-manifold form edges consisting of points with neighborhoods of the third type, and these edges come together at points with neighborhoods of the last type.
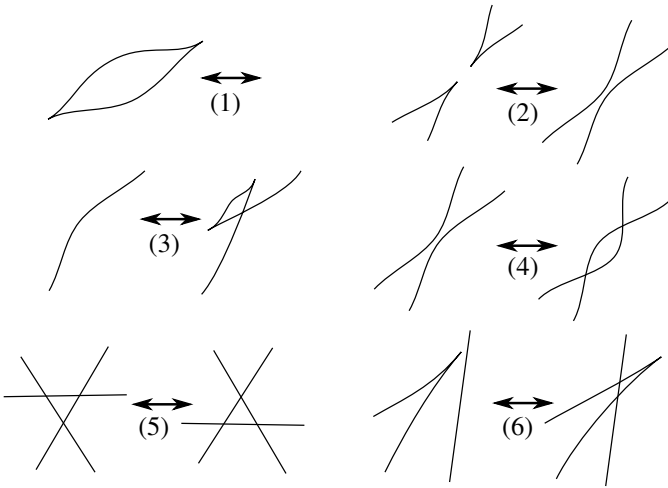
Figure 5. Any two graphics are related by isotopies and a sequence of these six moves.

Lemma 9. *Given a Morse function* $f : T^3 \to \mathbb{R}$, *and isotopic torus bundle maps* π^0, π^1 *such that* $f \times \pi^0$ *and* $f \times \pi^1$ *are stable then the maps* $f \times \pi^0$ *and* $f \times \pi^1$ *are related by a sequence of isotopies and moves of the types shown in Figure 5.*

In fact, the theorem is true for any generic path of maps from a 3-manifold to a 2-manifold, but we do not need to prove it in such generality, so we will stick to the situation described above.

Proof. Because π^0 and π^1 are isotopic, there is a smooth family of stable functions $\{\pi^t \mid t \in [0, 1]\}$ from π^0 to π^1. As noted above, we can approximate the family of stable functions $f \times \pi^t = f \times (\pi \circ \phi_t)$ by considering the isomorphic stable functions $(f \circ \phi_t^{-1}) \times \pi$, then approximating the path $\{f \times \phi_t^{-1}\}$ by a piecewise linear path of straight arcs between a sequence of functions $\{g^i \in C^\infty(T^3, \mathbb{R})\}$. The stable function defined by π and a function in the arc from g^i to g^{i+1} can be recovered as a projection of $g^i \times g^{i+1} \times \pi$ onto $L \times S^1$ for some line L in $\mathbb{R}^2$. Thus in order to understand how the graphic changes with t, we must consider projections of $g^i \times g^{i+1} \times \pi$. Without loss of generality, we will consider projections of $g^0 \times g^1 \times \pi$.

Let $S \subset T^3$ be the discriminant set of the map $F = g^0 \times g^1 \times \pi : T^3 \to \mathbb{R}^2 \times S^1$. As noted above, S is a compact, closed surface in T^3. The image of S in $\mathbb{R}^2 \times S^1$ is a surface with "cusps". On the complement of the cusps, there is a well defined map from $T_p S$ to $T_{F(p)}(\mathbb{R}^2 \times S^1)$ where p is a noncusp point in S. If the plane tangent to $F(S)$ at $F(p)$ is not parallel to the plane $\mathbb{R}^2 \times \{\pi(p)\}$ then its intersection with this plane determines a slope in $\mathbb{R}^2$. For each $p \in S$, let $s(p)$ be this slope. The

function s is defined on the complement in S of the points where $F(S)$ is parallel to the plane $\mathbb{R}^2 \times \{\pi(p)\}$.

By perturbing $F(S)$ slightly (by isotoping g^0, g^1 and π), we can ensure that the function s on S is a Morse function defined on the complement in S of a finite number of points. Locally, we can identify a patch of $F(S)$ with the graph of a function γ from $\mathbb{R}^2$ (in variables y, z) to $\mathbb{R}$ as follows: Let $\gamma' : \mathbb{R} \to \mathbb{R}$ be a smooth function. For each a, b, define $\gamma(a, b) = \gamma'(a) + \int_0^b s \, dy$, where the integral is taken along the arc from $(a, 0)$ to (a, b). The slope of the intersection with the plane $\mathbb{R}^2 \times \{\pi(p)\}$ is precisely $d\gamma/dy$ so for some γ', $F(S)$ is the graph of γ.

Let p_t be the orthogonal projection of $\mathbb{R}^2 \times S^1$ onto the annulus $L \times S^1$ where $L \subset \mathbb{R}^2$ is the line through the origin with slope $\cos(t)/\sin(t)$. The composition of F with p_t is a function from M to $\mathbb{R}^2$ and the discriminant set of $F \circ p_t$ is the projection of the closure in S of the subset $s^{-1}(t)$. In order to understand how the graphic changes with t, we must understand how $s^{-1}(t)$ maps to the graphic.

Because s is a Morse function, $s^{-1}(t)$ will be a collection of closed loops in S. By recovering S from s as the graph of a function γ (defined up to choice of γ'), one can check that the graphic defined by $F \circ p_t$ is related to the level sets as follows:

If the image of $s^{-1}(t)$ in $F(S)$ is transverse to the plane $\mathbb{R}^2 \times \{\pi(p)\}$ at a point $p \in S$ then in the projection, p maps into the interior of an edge in the graphic. If the level set is tangent to the plane at $F(p)$ then p maps to a cusp in the graphic. Each cusp point of S in the closure of $s^{-1}(t)$ maps to a cusp in the graphic. A noncusp point p where s is not defined (because $F(S)$ is tangent to the plane $\mathbb{R}^2 \times \{\pi(p)\}$) maps to the interior of an edge. (Such a point is in the closure of $s^{-1}(t)$ for every t.)

When t passes through a critical level of s, there are two cases to consider: If the critical point has index zero (or two) then right after (right before) the critical point, a component of $s^{-1}(t)$ will be tangent to the plane $\mathbb{R}^2 \times \{\pi(p)\}$ in exactly two points. Its image in $F \circ p_t$ will be an eye as in move (1). Before (after) the critical point, there is no such component so when t passes through the critical level, the eye is created (removed) as in move (1). Similar reasoning shows that the graphic changes by move (2) if the critical point has index one.

When t passes through a level where tangencies to the plane $\mathbb{R}^2 \times \pi(p)$ are created or removed, the tangencies are created or removed in pairs, and in the graphic this corresponds to move (3). At the points where $F(S)$ is tangent to the plane, the graphic does not change at all.

Generically, there will be no critical point of s in the (one dimensional) set of cusp points in S. Thus the graphic only changes along the set of cusps when a level set of s becomes tangent to an arc of cusp points. In this case, two cusps in the graphic are created or eliminated as in move (3).

A stable map from a 3-manifold to $\mathbb{R}^2$ can fail to be stable for two reasons: there may be critical points that do not have neighborhoods of the necessary forms, or the

image of the discriminant set may be nongeneric, that is, have triple points, tangencies or double points at cusps. We have shown that there is a family of functions from $f \times \pi_0$ to $f \times \pi_1$ such that there are finitely many intermediary functions with unstable neighborhoods, and these correspond to the first three moves. In between these functions, the graphic changes by some homotopy of the image of the discriminant set. By perturbing the family $\{f_t\}$ slightly we can ensure that the homotopy is generic, consisting of a finite sequence of moves (4), (5) and (6). This completes the proof.   □

8. Graphics and isotopies

Proof of Lemma 7. A torus twist determines an automorphism that is isotopy trivial in T^3 so it takes the torus determined by a weak reducing pair of disks to an isotopic torus. If two weak reducing pairs are related by a sequence of torus twists then by induction the induced tori are isotopic. The majority of the proof of Lemma 7 will be devoted to the converse of this.

Let $f : T^3 \to [0, 1]$ be a Morse function with one index zero critical point, three index one critical points followed by three index two critical points, then one index three critical point. Let ℓ_1 be the level of the last index one critical point and ℓ_2 the level of the first index two critical point. Any level set $f_t = f^{-1}(t)$ for $t \in (\ell_1, \ell_2)$ is a Heegaard surface for T^3 and is thus isotopic to Σ for the standard (and unique) Heegaard splitting (Σ, H_1, H_2). Moreover, there is a continuous family $\{c_t : \Sigma \to f_t \mid t \in (\ell_1, \ell_2)\}$ identifying each level surface with Σ.

Assume f has the property that for $t < \ell_1$, f_t is not compressible into $f^{-1}([t, 1])$. This will be the case whenever the meridian disk defined by the last index one critical point is not part of a weak reducing pair, so such an f exists. Similarly, assume that for $t > \ell_2$, f_t is not compressible into $f^{-1}([0, t])$.

For a weak reducing pair D_1, D_2, there is a bundle map $\pi^0 : T^3 \to S^1$ such that the weak reducing pair determined by $f \times \pi^0$ and the maps c_t is isotopic to D_1 and D_2. Likewise, there is a bundle map π^1 such that $f \times \pi^1$ determines a weak reducing pair isotopic to the second pair D'_1, D'_2. By assumption the weak reducing pairs define isotopic incompressible tori, so π^0 and π^1 are isotopic. Thus Lemma 9 implies that $f \times \pi^0$ and $f \times \pi^1$ are related by a sequence of the moves shown in Figure 5 and isotopies of the graphic.

Because we chose f so that f_t is incompressible in one direction when $t < \ell_1$ or $t > \ell_2$, π^r determines a unique weak reducing pair of disks for every value of r for which the graphic for $f \times \pi^r$ is generic. In particular, Lemma 8 implies that for every generic s, there is a unique maximal interval (a_r, b_r) and a unique weak reducing pair such that for $t \in (a_r, b_r)$, the boundaries of the weak reducing pair are level sets of π^r restricted to f_t. The essential level sets of π^r restricted to f_t

change only when t passes through a value where there is a crossing in $\{t\} \times S^1$. Thus the circles $\{a_r\} \times S^1$ and $\{b_r\} \times S^1$ pass through crossings in the graphic.

The map $f \times \pi^r$ can fail to be generic for two reasons: the map may fail to be stable (that is, when it undergoes one of the moves in Figure 5) or there may be two crossings of the graphic at the same value of t. Generically, each of these will happen for only finitely many values of r.

If $f \times \pi^r$ fails to be stable because there is a point in M with a nonstable neighborhood (that is, when the graphic changes by one of the first three types of moves), the graphic changes within a subset $I \times S^1$ for some interval $I \subset [0, 1]$. For moves (1) and (2), there are no crossings in this band, so I is either disjoint from (a_r, b_r) or a proper subset of (a_r, b_r). For move (3), there is a single crossing in $I \times S^1$ but this crossing cannot be an endpoint of (a_r, b_r) because at least one of the arcs involved corresponds to a pair of pants with a trivial boundary loop in f_t. Again I is disjoint from or properly contained in (a_r, b_r). The crossings created (removed) by moves (4) and (6) cannot be endpoints of (a_r, b_r) because before (after) the move, these crossings don't exist.

For moves (1), (2), (3), (4) and (6), the move is disjoint from $\{t\} \times S^1$ for some $t \in (a_r, b_r)$ so the level loops of π^r restricted to f_t bounding a weak reducing pair do not change. The weak reducing pair with boundaries in level sets of π_r restricted to f_t does not change as r passes through r_0 so the weak reducing pair induced by π^r does not change.

The last two cases to deal with are cases in which more than one crossings pass through the same vertical arc of the graphic. This will only change the induced weak reducing disks if before and after the move, these crossings sit in the vertical loops $a_r \times S^1$ and $b_r \times S^1$. When the move occurs, this region shrinks down to a single arc, then expands back to an annulus, but with a new pair of loops bounding weak reducing disks as in Figure 6.

To understand how this can happen, we must consider the following situation: for some t and r, two level sets of π^r restricted to f_t bound a weak reducing pair. Just to the left of t there is a crossing in the graphic that eliminates one of the loops in the pair, and just to the right of t there is a crossing that eliminates the other loop in the pair.
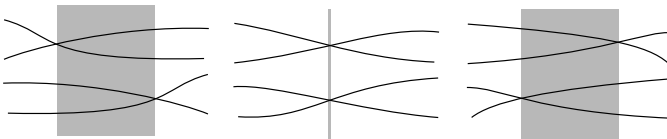
Figure 6. The induced weak reducing disks change when the interval $[a_s, b_s]$ shrinks to a point and the graphic fails to be generic.
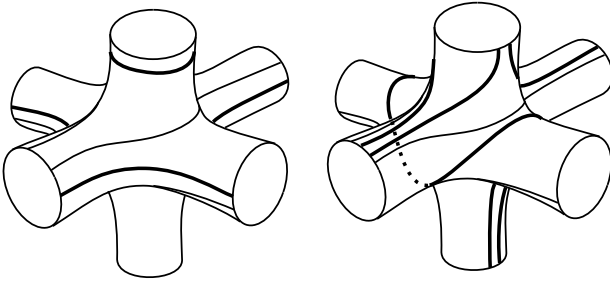
Figure 7. The pairs of weak reducing disks before and after the graphic becomes nongeneric.

Note that no component of the complement of $\Sigma \setminus (\partial D_1 \cup \partial D_2)$ is a pair of pants. This implies that the critical points just above and just below one of the loops in the weak reducing pair are distinct from the critical points just above and below the other loop. Thus the two crossings must involve four distinct edges of the graphic, so move (5) cannot change the pair of weak reducing disks.

The only case in which four distinct edges of the graphic can be involved in the crossing is when $f \times \pi^r$ is stable but not generic because there are two crossings at the same value of t. Each boundary loop of the weak reducing pair sits in a four punctured sphere bounded by essential level loops and each crossing corresponds to replacing the boundary of the weak reducing disk with a new loop in the same four punctured sphere. The new loops form a new weak reducing pair because after the move, the graphic is generic again. Thus each boundary loop of the weak reducing pair is replaced by a loop bounding a disk in the opposite handlebody.

Given a pair of weak reducing disks and bounding level sets of some π^r as on the left side of Figure 7, we can find a new pair of weak reducing disks as on the right. A four punctured sphere in Σ that contains a loop bounding a disk in H_1 contains at most one loop bounding a disk in H_2 and vice versa. Thus the new pair of weak reducing disks shown on the right of the figure is the pair defined by the graphic on the right.

One can check that the original weak reducing pair shown on the left is taken to the new pair shown on the right by a torus twist along a torus T that intersects S in a single loop and cuts Σ into two four punctured spheres.

We have found a sequence of stable functions such that each determines a pair of weak reducing disks for Σ, the first and last stable functions determine D_1, D_2 and D_1', D_2', respectively, and the weak reducing disks determined by consecutive stable functions are either isotopic or related by a torus twist along a torus meeting the criteria for Lemma 7. This defines a sequence of torus twists taking D_1, D_2 onto D_1', D_2', completing the proof. $\square$

9. Proof of Theorem 1

Recall that G is the subgroup of $\mathrm{Mod}(T^3, \Sigma)$ generated by the elements α_{ij}, σ and τ defined above. To prove Theorem 1, we must show that G is equal to $\mathrm{Mod}(T^3, \Sigma)$.

Proof of Theorem 1. Let ϕ be an element of $\mathrm{Mod}(T^3, \Sigma)$. If ϕ interchanges the handlebodies H_1 and H_2, then composing ϕ with σ produces an automorphism that sends H_1 to itself. Since σ is an element of G, $\sigma\phi$ will be in G if and only if ϕ is in G. Thus by replacing ϕ with $\sigma\phi$ if necessary, we will assume that ϕ preserves each handlebody.

The image $i(\phi)$ of ϕ in $\mathrm{Mod}(T^3)$ is a composition of the images $\{i(\alpha_{ij})\}$. Thus composing ϕ with a sequence of these maps produces an element of $\mathrm{Mod}(T^3, \Sigma)$ that is isotopy trivial in T^3. The maps $\{\alpha_{ij}\}$ are in G so ϕ is in G if and only if this composition is in G. We can thus assume ϕ is in the kernel of i.

Let T be the incompressible torus determined by D_1^1 and D_2^1. The disks $\phi(D_1^1)$ and $\phi(D_2^1)$ form a weak reducing pair for Σ and determine an incompressible torus T' isotopic to $\phi(T)$. Because ϕ is isotopic to the identity on T^3, T' is in fact isotopic to T. Because D_1^1, D_2^1 and $\phi(D_1^1)$, $\phi(D_2^1)$ determine isotopic tori in T^3, Lemma 7 implies that there is a sequence of torus twists, along tori that intersect T in a single loop and cut Σ into planar pieces, taking D_1^1 to $\phi(D_1^1)$ and D_2^1 to $\phi(D_2^1)$.

Any torus twist along a torus that intersects T in a single loop is a conjugate of τ by elements of G, by Lemma 3. Composing ϕ by this conjugate of τ produces an element of $\mathrm{Mod}(T^3, \Sigma)$ that is isotopy trivial on T^3 and preserves D_1^1 and D_2^1. This conjugate is in G if and only if ϕ is in G, so we can replace ϕ with an element of $\mathrm{Mod}(T^3, \Sigma)$ that is in the kernel of i and takes D_1^1 and D_2^1 onto themselves.

The complement in Σ of $\partial D_1^1 \cup \partial D_2^1$ is a pair of twice punctured tori. Let S be one of these twice punctured tori. The automorphisms α_{23}, α_{32}, τ_{12} and τ_{13} take S onto itself and their restrictions to S generate the mapping class group of the twice punctured torus. Thus there is an element g of G such that $g \circ \phi$ restricts to the identity on S.

Each of the loops ∂D_1^2 and ∂D_1^3 intersects ∂D_2^1 in two points, so each loop intersects S in a single properly embedded arc. Because each loop bounds a disk in H_1, it must intersect the twice punctured torus $\Sigma \setminus S$ in a single arc parallel to the arc in S. Because $g \circ \phi$ is the identity on S, it fixes $\partial D_1^2 \cap S$ and $\partial D_1^3 \cap S$, and therefore also fixes the parallel arcs in $\Sigma \setminus S$. This implies that $g \circ \phi$ is the identity on $\Sigma \setminus S$ as well as S. Thus $g \circ \phi$ is the identity on all of Σ so $\phi = g^{-1} \in G$, completing the proof. $\qquad\square$

References

[Akbas 2008] E. Akbas, "A presentation for the automorphisms of the 3-sphere that preserve a genus two Heegaard splitting", *Pacific J. Math.* **236**:2 (2008), 201–222. MR 2009d:57029 Zbl 1157.57002

[Bachman and Schleimer 2005] D. Bachman and S. Schleimer, "Surface bundles versus Heegaard splittings", *Comm. Anal. Geom.* **13**:5 (2005), 903–928. MR 2006m:57027 Zbl 1138.57026

[Boileau and Otal 1990] M. Boileau and J.-P. Otal, "Sur les scindements de Heegaard du tore T^3", *J. Differential Geom.* **32**:1 (1990), 209–233. MR 91i:57006 Zbl 0754.53012

[Cho 2008] S. Cho, "Homeomorphisms of the 3-sphere that preserve a Heegaard splitting of genus two", *Proc. Amer. Math. Soc.* **136**:3 (2008), 1113–1123. MR 2009c:57029 Zbl 1149.57025

[Cooper and Scharlemann 1999] D. Cooper and M. Scharlemann, "The structure of a solvmanifold's Heegaard splittings", *Turkish J. Math.* **23**:1 (1999), 1–18. MR 2000h:57034 Zbl 0948.57015

[Goeritz 1933] L. Goeritz, "Die Abbildungen der Brezelfläche und der Volbrezel vom Geschlecht 2", *Abh. Math. Sem. Univ. Hamburg* **9** (1933), 244–259. Zbl 0007.08102

[Golubitsky and Guillemin 1973] M. Golubitsky and V. Guillemin, *Stable mappings and their singularities*, Graduate Texts in Math. **14**, Springer, New York, 1973. MR 49 #6269 Zbl 0294.58004

[Hatcher and Thurston 1980] A. Hatcher and W. Thurston, "A presentation for the mapping class group of a closed orientable surface", *Topology* **19**:3 (1980), 221–237. MR 81k:57008 Zbl 0447.57005

[Johnson 2009] J. Johnson, "Stable functions and common stabilizations of Heegaard splittings", *Trans. Amer. Math. Soc.* **361**:7 (2009), 3747–3765. MR 2010b:57023 Zbl 1168.57013

[Johnson and Rubinstein 2006] J. Johnson and J. H. Rubinstein, "Mapping class groups of Heegaard splittings", preprint, 2006. arXiv math.GT/0701119

[Kobayashi and Saeki 2000] T. Kobayashi and O. Saeki, "The Rubinstein–Scharlemann graphic of a 3-manifold as the discriminant set of a stable map", *Pacific J. Math.* **195**:1 (2000), 101–156. MR 1781617 (2001i:57026)

[Mather 1970] J. N. Mather, "Stability of C^∞ mappings: V, Transversality", *Advances in Math.* **4**:3 (1970), 301–336. MR 43 #1215c Zbl 0207.54303

[Namazi 2004] H. Namazi, "Big handlebody distance implies finite mapping class group", preprint, 2004. arXiv math.GT/0406551

[Scharlemann 2004] M. Scharlemann, "Automorphisms of the 3-sphere that preserve a genus two Heegaard splitting", *Bol. Soc. Mat. Mexicana* (3) **10**:Special Issue (2004), 503–514. MR 2007c:57020 Zbl 1095.57017

[Scharlemann and Thompson 1994] M. Scharlemann and A. Thompson, "Thin position for 3-manifolds", pp. 231–238 in *Geometric topology* (Haifa, 1992), edited by C. Gordon et al., Contemp. Math. **164**, Amer. Math. Soc., Providence, RI, 1994. MR 95e:57032 Zbl 0818.57013

JESSE JOHNSON
DEPARTMENT OF MATHEMATICS
OKLAHOMA STATE UNIVERSITY
STILLWATER OK 74078
UNITED STATES
jjohnson@math.okstate.edu
http://www.math.okstate.edu/~jjohnson/

PACIFIC JOURNAL OF MATHEMATICS

http://www.pjmath.org