TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 358, Number 9, September 2006, Pages 3781–3805
S 0002-9947(05)03767-0
Article electronically published on May 26, 2005

3-MANIFOLDS WITH PLANAR PRESENTATIONS AND THE WIDTH OF SATELLITE KNOTS

MARTIN SCHARLEMANN AND JENNIFER SCHULTENS

ABSTRACT. We consider compact 3-manifolds M having a submersion h to R in which each generic point inverse is a planar surface. The standard height function on a submanifold of S^3 is a motivating example. To (M, h) we associate a connectivity graph Γ. For $M \subset S^3$, Γ is a tree if and only if there is a Fox reimbedding of M which carries horizontal circles to a complete collection of complementary meridian circles. On the other hand, if the connectivity graph of $S^3 - M$ is a tree, then there is a level-preserving reimbedding of M so that $S^3 - M$ is a connected sum of handlebodies.

Corollary.
- *The width of a satellite knot is no less than the width of its pattern knot and so*
- $w(K_1\#K_2) \geq max(w(K_1), w(K_2))$.

The notion of thin position, introduced by D. Gabai [G], has been employed with great success in many geometric constructions. Yet the underlying notion of the width of a knot remains shrouded in mystery. Little is known about the width of specific knots, or how knot width behaves under connected sum. By stacking a copy of K_1 in thin position on top of a copy of K_2 in thin position, it is easily seen that $w(K_1\#K_2) \leq w(K_1) + w(K_2) - 2$. Here we establish a lower bound for the width of a knot sum: the width is bounded below by the maximum of the widths of its summands and therefore also by one-half the sum of the widths of its summands.

Knot width can be thought of as a kind of refinement of bridge number. Interest in how the width of a knot behaves under connected sum is inspired, in part, by the fact that bridge number behaves very well. Indeed for bridge number, $b(K_1\#K_2) = b(K_1) + b(K_2) - 1$; see the paper [S] by H. Schubert, or [Sch] for a much shorter proof. The shorter proof in [Sch] crystallized out of an investigation into whether or not thin position arguments clarify the behaviour of bridge number under connected sum. The answer to that question appears to be no: width seems to be a much more refined invariant than can be useful for the recovery of Schubert's result. In particular, the argument in [Sch] fails in settings where the swallow follow torus is too convoluted. One suspects that degeneration of width under connected sum of knots is possible, i.e., that there might be knots K_1, K_2, such that $w(K_1\#K_2) < w(K_1) + w(K_2) - 2$. The situation may be analogous to that of

Received by the editors September 28, 2003 and, in revised form, May 18, 2004.
2000 *Mathematics Subject Classification.* Primary 57M25, 57M27.
The authors thank RIMS Kyoto, where this work was begun, Professor Tsuyoshi Kobayashi for inviting us to RIMS, Yo'av Rieck for helpful conversations there, and the NSF for partial support via grants DMS 0203680 and DMS 0104039. The second author also thanks the MPIM-Bonn for support.

another knot invariant, tunnel number. For small knots (knots whose complements contain no essential closed surfaces), neither width nor tunnel number degenerate under connected sum; i.e., for small knots, width of knots satisfies $w(K_1 \# K_2) = w(K_1) + w(K_2) - 2$ and tunnel number satisfies $t(K_1 \# K_2) \geq t(K_1) + t(K_2)$. This is proven in [RS] and [MS], respectively. On the other hand, it is known that tunnel number can degenerate under connected sum, for knots that are not small. See for example [Mo]. Our results on knot width are in a spirit similar to that of [ScSc], establishing an upper bound for such possible degeneration. Explicitly:

Corollary 6.4. *For any two knots K_1, K_2,*

$$w(K_1 \# K_2) \geq max\{w(K_1), w(K_2)\} \geq \frac{1}{2}(w(K_1) + w(K_2)).$$

We obtain Corollary 6.4 by applying the following more general result to the swallow-follow companion tori that are associated to the connected sum of knots (see [L, p. 10], or the discussion in Section 6).

Corollary 6.3. *Suppose K' is a satellite knot with pattern K. Then $w(K') \geq w(K)$.*

Our approach to the latter result is to think of the companion solid torus as a simple example of a handlebody in S^3. We ask, in general, how a handlebody H in S^3 might be reimbedded so that its complement is also a handlebody, hoping in particular to find a reimbedding that preserves the natural projection to R (called *height*: $h : H \subset S^3 \subset R^4 \to R$). There is a theory of reimbeddings in S^3 going back to Fox [Fo], who showed that any connected $M \subset S^3$ can be reimbedded so that its complement is a union of handlebodies. What is new here is the concern about height $h : M \to R$.

In Section 2 we associate to an arbitrary compact $M \subset S^3$, a certain graph Γ, and show that Γ is a tree if and only if there is a collection of horizontal (with respect to height) circles in ∂M which constitute a complete collection of meridian circles after a reimbedding whose complement is a handlebody. This discussion is in some sense only a digression; the main argument begins with Section 3.

Our goal is to reimbed a handlebody H (preserving height) so that the complement $M = S^3 - H$ is also a handlebody. What we in fact study carefully is the complement M, hoping that by reconstructing it appropriately, without changing h on M, we can turn M into a handlebody. One way to recognize that we are done is to observe that if H can be made to look like the neighborhood of a graph Λ and Λ lies in $S^2 \subset S^3$, then $S^3 - H$ is indeed a handlebody. We call such a graph Λ unknotted. In Section 3 we develop methods to construct and recognize unknotted graphs. In Section 4 we describe how, if the graph Γ associated to $M = S^3 - H$ is a tree, we can reconstruct M, without affecting height h, so that M becomes the complement of an unknotted graph, i.e. a handlebody. Such a reimbedding of H is called a Heegaard reimbedding. In Section 5 we observe that the only effect of this reconstruction of M on H is to alter it by braid moves; the corollaries on knot width then follow in Section 6.

1. Mathematical preliminaries

Throughout the paper, all manifolds will be orientable and, unless otherwise stated, compact. All embeddings will be locally flat. Since in dimension three there is little topological distinction between smooth manifolds and PL manifolds, and it will be convenient to use ideas and language from both smooth and PL topology, we will do so without apology, leaving it to the reader to make the appropriate translation if a specific structure (smooth or PL) is initially given on the manifold.

Definition 1.1. A *planar presentation* of a 3-manifold $(M, \partial M)$, $\partial M \neq \emptyset$ is a map $h : M \to R$ so that

(1) $Dh : T_M \to T_R$ is always surjective,
(2) $h|\partial M$ is a general position Morse function, and
(3) for t any regular value of $h|\partial M$, $h^{-1}(t)$ is a planar surface, denoted P^t.

The motivating source of examples is this: Consider $S^3 \subset R^4$ and let $p : R^4 \to R$ be a standard projection, so $p|S^3 : S^3 \to [-1, 1]$ has two critical points in S^3, typically called the north and south poles. Let $M \subset S^3$ be a compact submanifold that does not contain either pole. Then $h = p|M$ is a planar presentation of M.

Consider an index one (i.e. saddle) critical value t_σ of $h|\partial M$. The corresponding critical point is called an *upper saddle* (resp. *lower saddle*) if $\partial P^{t_\sigma - \epsilon}$ has one more (resp. one less) circle component than $\partial P^{t_\sigma + \epsilon}$. If the number of components in $P^{t_\sigma + \epsilon}$ and $P^{t_\sigma - \epsilon}$ is the same, we say the saddle is *nested*; otherwise the saddle is *unnested*. Here is an alternate description: an upper (resp. lower) saddle is nested if and only if the outward normal from M points up (resp. down) at the saddle point. (In particular, if $S \subset S^3$ is a surface, then a saddle singularity of S is a nested saddle for the component of $S^3 - S$ lying just above the singularity if and only if it is unnested for the component of $S^3 - S$ lying just below the saddle.) Similarly, a maximum (resp. minimum) of h on ∂M is called an external maximum (resp. minimum) of h on M if the outward pointing normal from M points up (resp. down) at the critical point. Other maxima and minima on ∂H will be called *internal* maxima and minima. See Figure 1. (This and other figures can be computer viewed in color at this paper's ArXiv site.)

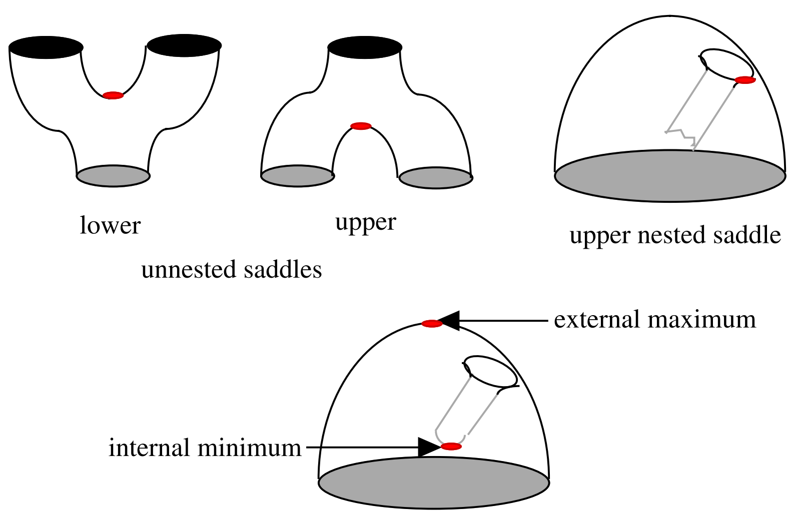

FIGURE 1.

2. The connectivity graph and Fox reimbedding

Throughout this section, (M, h) will be a planar presentation, and $s_1 < s_2 < \ldots < s_n$ will be the set of critical values at which $h|\partial M$ has an unnested saddle or an external maximum or minimum. The points $x_1, \ldots, x_n \in \partial M$ will be the corresponding critical points.

Lemma 2.1. *Suppose M_0 is a component of $M - \bigcup_{i=1}^n P^{s_i}$. Then for any generic height $\overline{t}$, $P_0 = M_0 \cap P^{\overline{t}}$ is connected (possibly empty).*

Proof. Choose any two points in P_0. Since M_0 is connected, there is an arc $\alpha \subset M_0$ that runs between them; a generic such arc will have its critical heights at different levels than ∂M does. Since $\alpha \subset M_0$, α is disjoint from $\{P^{s_i}\}$. So for some i, the height of α lies between s_i and s_{i+1}. Let $t_1, \ldots, t_m$ be the critical values (if any) of $h|\partial M_0$ between s_i and s_{i+1} and choose α to minimize the number of points in $T_\alpha = \{t_j \in h(\alpha)\}$. If T_α is empty, then α lies entirely in a region with no critical values, i.e. $M_0 \cap h^{-1}h(\alpha) \cong P_0 \times I$. Project α to P_0 and deduce that the ends of P_0 lie in the same component of P_0.

We now show that in fact T_α is always empty. For suppose t_j is the greatest value (if any) of T_α that is greater than $\overline{t}$ (or, symmetrically, the lowest value of T_α below $\overline{t}$). The same argument as above shows that each subarc of α that lies above P^{t_j} can be projected to lie in $P^{t_j+\epsilon}$ for any small ϵ. Since, by assumption, passing through the critical level t_j does not connect or disconnect any component of P^t; in fact such a subarc can then be pushed below t_j. Once this is done for every subarc of α above t_j, T_α is reduced by the removal of t_j, a contradiction.

We have thereby shown that any two points in P_0 can be connected by an arc in P_0, so P_0 is connected. □

Definition 2.2. The *connectivity graph* Γ of (M, h) is the graph whose vertices correspond to components of $M - \bigcup_{i=1}^n P^{s_i}$ and whose edges correspond to components of $\bigcup_{i=1}^n (P^{s_i} - x_i)$. An edge corresponding to a component P_0 of $P^{s_i} - x_i$ has its ends at the vertices that correspond to the components of $M - \bigcup_{i=1}^n P^{s_i}$ that lie just above and below P_0.

It is an old theorem of Fox [Fo] that any compact connected 3-dimensional submanifold M of S^3 can be reimbedded in S^3 so that the closure of $S^3 - M$ is a union of handlebodies. (This theorem has recently been updated to include other non-Haken 3-manifolds [Th].) As described above, let $p : S^3 \to R$ be the standard height function and let $M \subset S^3$ be a 3-manifold in general position with respect to p. One can refine Fox's question and ask if M can be reimbedded in S^3 so that the complement is a collection H of handlebodies and, furthermore, each horizontal circle in ∂M (that is, each component of each generic ∂P^t) bounds a disk in H. Put another way, the question is whether a Fox reimbedding of M can be found so that in the complementary handlebodies a complete collection of meridian disks is horizontal with respect to the original height function on M.

A first observation is that we may as well assume M does not contain the poles. For if M contains the north pole, say, let t be the highest critical value of $h = p|M$ on ∂M, necessarily the image of a maximum on ∂M. Alter M by simply removing the ball $h^{-1}(t - \epsilon, \infty)$. The result does not contain the north pole and (after a tiny isotopy) is homeomorphic to M via a homeomorphism that preserves the height function h on ∂M. So, after this initial reimbedding, we may think of the pair (M, h) as a planar presentation of M.

Then the answer is straightforward:

Proposition 2.3. *There is a collection of handlebodies H so that $M \cup_\partial H \equiv S^3$. Moreover, there is a complete collection of meridian disks for H whose boundaries are all horizontal (with respect to h) in ∂M if and only if the connectivity graph Γ of M is a tree.*

Proof. The first claim is the central theorem of [Fo].

The second claim follows from the central theorem of [Sc1]. This says that a collection of 0-framed curves $C \subset \partial M$ contains a complete collection of meridians for some complementary handlebody H if and only if it has this property: Any properly embedded surface S in M whose boundary is disjoint from C separates M.

If Γ is not a tree, then some component P_0 of some P^t is non-separating, and clearly such a component can be made disjoint from any finite collection of horizontal circles in M. If M could be imbedded in S^3 so that the complement consisted of handlebodies H in which a complete collection of meridian boundaries were horizontal with respect to h, then P_0 could be capped off in H by adding disks to ∂P_0. The result would be a non-separating closed surface in S^3, and this of course is impossible.

Conversely, suppose C is a finite collection of horizontal circles in ∂M chosen so large that any horizontal circle in ∂M is parallel to an element of C in ∂M. Suppose S is a generic non-separating properly embedded surface in M with boundary disjoint from C. Let α be a generic simple closed curve in M which intersects S in an odd number of points. Choose such an S to minimize $|S \cap (\bigcup_{i=1}^n P^{s_i})|$, where, as above, $\{s_i\}$ is the set of heights of the unnested saddles and of the external minima and maxima of M.

The first observation is that in fact $S \cap (\bigcup_{i=1}^n P^{s_i}) = \emptyset$. For otherwise, choose an innermost circle c of intersection of S with a component P_0 of $\bigcup_{i=1}^n P^{s_i}$. Here "innermost" means that c cuts off from P_0 a subplanar surface P_- whose boundary consists of c and a collection of boundary circles of P_0. Then replacing a vertical collar of c in S with two parallel horizontal copies of P_- gives a surface which has fewer components of intersection with $\bigcup_{i=1}^n P^{s_i}$ but which still contains a non-separating component, since the number of intersections with α is increased by $2 \cdot |\alpha \cap P_-|$ and so remains odd. Since the boundary of P_- is horizontal, generically it is disjoint from C.

So S lies in a component of $M - \bigcup_{i=1}^n P^{s_i}$ whose closure we denote M_0. Let $h(M_0) = [s_i, s_{i+1}]$, so M_0 lies in a slice of S^3 homeomorphic to $S^2 \times [s_i, s_{i+1}]$. So as not to be distracted by other parts of M, let Q be a 2-sphere and momentarily think of M_0 as lying in $Q \times [s_i, s_{i+1}]$. Since every horizontal cross-section of M_0 is connected, at any generic height a cross-section of $Q - M_0$ is a collection of disks. In particular, the boundary components of S can be capped off in $Q \times [s_i, s_{i+1}]$ to give a closed surface $S_+ \subset Q \times [s_i, s_{i+1}]$.

Now consider how the arcs $\alpha \cap M_0$ lie in $Q \times [s_i, s_{i+1}]$. Any arc with both ends in $Q \times (s_i)$ or both ends in $Q \times (s_{i+1})$ can be entirely homotoped in $Q \times [s_i, s_{i+1}]$ into that end and so be made disjoint from S_+. It follows that such an arc intersects S an even number of times. Since α intersects S an odd number of times, it follows that there are an odd number of arcs of $\alpha \cap M_0$ that run from the top of M_0 to the bottom. Then, returning again to $M \subset S^3$ there must be an odd number of arcs of $\alpha - M_0$ that run from the top of M_0 to the bottom of M_0 in $M - M_0$. In particular, there is at least one such arc, so one can construct a closed curve in M that intersects the bottom of M_0 in a single point p. Hence removing the edge in Γ corresponding to the component of $P^{s_i} - x_i$ in which p lies does not disconnect Γ. Since we can remove an edge and not disconnect Γ, Γ is not a tree. $\square$

3. Unknotted graph complements

Definition 3.1. For N a compact 3-manifold and Λ a finite graph, a *proper embedding* $\Lambda \subset N$ is an embedding so that $\partial N \cap \Lambda$ consists of a collection of valence one vertices of Λ. These vertices are denoted $\partial\Lambda$. The other vertices, some of which may also have valence one, are called *interior vertices*.

In case $N = B^3, S^3$ or $S^2 \times I$, the pair $(N - \eta(\Lambda), \partial N - \eta(\Lambda))$ will be denoted (N_Λ, P_Λ) and will be called a *graph complement with planar part* P_Λ. Graphs Λ and Λ' are equivalent if there is a homeomorphism $(N, \eta(\Lambda)) \cong (N, \eta(\Lambda'))$. In particular, if Λ' is any graph obtained from Λ by sliding and isotoping edges rel $\partial\Lambda$, then Λ and Λ' are equivalent graphs.

Two graph complements (N_Λ, P_Λ) and $(N_{\Lambda'}, P_{\Lambda'})$ will be called equivalent if they are pairwise homeomorphic. In particular, equivalent graphs have equivalent graph complements.

An interval $J \subset R$ is *proper* if it intersects any compact subset of R in a compact set. Equivalently, it is proper if and only if $J = R$ or J has one of the forms $[a, b], [a, \infty), (-\infty, b]$.

Lemma 3.2. *Suppose (M, h) is a planar presentation of a compact manifold M, $J \subset R$ is a proper interval with endpoint(s) generic for h, and suppose that in a component M^J of $h^{-1}(J)$ all saddles are nested. Then M^J is homeomorphic to a graph complement with planar part $h^{-1}(\partial J) \cap M^J$.*

Proof. Since M^J is connected and contains no unnested saddles, each generic horizontal cross-section is a connected planar surface, by Lemma 2.1. Since J is proper, M^J is compact, so we may as well assume J is compact (say $J = [0, 1]$), although we do not know that $h(M^J) = J$. We will describe the graph Λ for which M^J is the complement; the details of the homeomorphism then follow from standard Morse theory. See Figure 2.

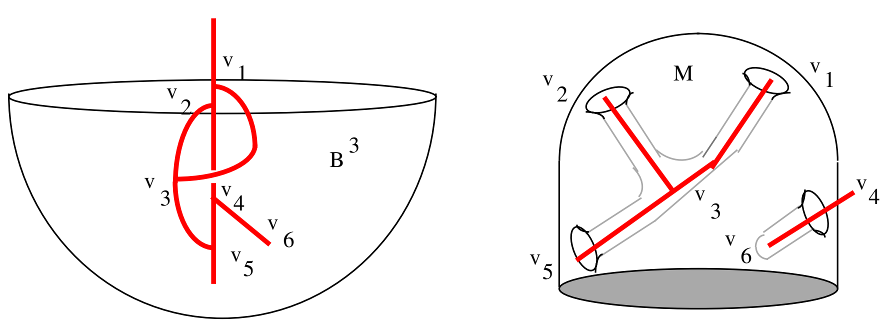

Figure 2. M as graph complement (in B^3)

Suppose first that $J = h(M^J)$, so each component of $h^{-1}(\partial J) \cap M^J$ is a non-empty connected planar surface. We will describe $\Lambda \subset S^2 \times I$. Each (circle) boundary component of $h^{-1}(\{1\}) \cap M^J$ can be capped off with a disk to give a two sphere; dually, $h^{-1}(\{1\}) \cap M^J$ can be thought of as obtained from $S^2 \times \{1\}$ by removing some vertices. These will be vertices in $\partial\Lambda$. As t descends from 1 through generic values of t, each boundary component of P^t can be capped off by a disk to give a sphere S^t. This gives an embedding $P^t \subset S^t$; dually P^t can be obtained from the sphere S^t by removing a neighborhood of the center of each disk. As t varies, these points form vertical edges in Λ incident to those vertices of $\partial\Lambda$ that lie at height 1.

Now consider what happens as t descends through a critical point of $h|\partial M$. Each such critical point corresponds to an interior vertex of Λ. In particular, edges descend from those vertices that correspond to internal maxima (and descend *to* the vertices that correspond to internal minima). At lower saddles two edges descend into the corresponding vertex and one edge descends from it whereas at upper saddles one edge descends into the corresponding vertex and two edges descend from it. There are no external maxima or minima, for these would necessarily start (or end) a different planar surface, which could never be connected to M^J since all saddles are nested.

The argument is little changed if $h(M_J) \neq J$. Say $1 \notin h(M^J)$; then there is an external maximum on M^J at height $t_{max} \in J$ and at a generic height just below it the ball $M^J \cap [t_{max} - \epsilon, 1]$ can be thought of as the complement of a radius of B^3, and so as a graph complement in B^3. The rest of the construction proceeds as above, although now viewed as a construction in a collar $\partial B^3 \times I$. Ultimately M^J is thereby described as a graph complement in B^3 (or in S^3 if also $0 \notin h(M^J)$). $\square$

The case when the graph Λ is planar will be particularly important. Let $S^1 \times I$ denote the standard vertical cylinder in $S^2 \times I$.

Definition 3.3. A properly imbedded graph Λ in $N = S^3$ (resp. B^3 or $S^2 \times I$) is *unknotted* if it lies in $S^2 \subset S^3$ (resp. $B^2 \subset B^3$ or $S^1 \times I \subset S^2 \times I$). The pair (N_Λ, P_Λ) is then called an *unknotted graph complement* with *planar part* P_Λ. (Note that the number of components of P_Λ determines whether the ambient manifold is S^3, B^3 or $S^2 \times I$.)

More generally, any graph which is equivalent to an unknotted graph will be called an unknotted graph.

Unknotted graphs are in some sense unique:

Proposition 3.4. *Suppose Λ and Λ' are unknotted graphs in $N = S^3, B^3$, or $S^2 \times I$. Suppose that $\partial\Lambda = \partial\Lambda' \subset \partial N$.*

Then there is a homeomorphism of pairs $(N, \eta(\Lambda)) \cong (N, \eta(\Lambda'))$ which is the identity on $\eta(\partial\Lambda) = \eta(\partial\Lambda')$ if and only if there is a correspondence between the components of Λ and the components of Λ' with two properties:

- *The Euler characteristic of a component of Λ is the same as the Euler characteristic of the corresponding component of Λ' and*
- *the boundary points of each component of Λ are also the boundary points of the corresponding component of Λ'.*

Proof. The existence of such a homeomorphism clearly implies that the partitions and the corresponding Euler characteristics are the same. The difficulty is in proving the other direction. We consider the case $N = S^2 \times I$, for it is representative (and in fact the most difficult).

It will be convenient to number the p components of Λ (and the corresponding components of Λ') in some order $\Lambda_i, i = 1, \dots, p$, and then order the points $\partial\Lambda \cap (S^2 \times \{1\}) = \{w_j\}$ and $\partial\Lambda \cap (S^2 \times \{0\}) = \{v_k\}$ in some subordinate order, i.e. so that in the ordering all the boundary points of any earlier component of Λ come before all the boundary points of any later component.

Since Λ is an unknotted graph we can assume (up to homeomorphism of the pair $(N, \eta(\Lambda))$ rel $\eta(\partial(\Lambda)))$ that $\Lambda \subset S^1 \times I \subset S^2 \times I$. In a small neighborhood of Λ collapse a forest that is maximal in Λ among those not incident to $\partial\Lambda$. Then each component is the cone on its boundary vertices, wedged with some circles. (In particular, any component with no boundary is just a wedge of circles.) Each circle (even those that are essential in the vertical cylinder $S^1 \times I$) can be moved (rel the cone point) in $S^2 \times I$ until it bounds a tiny disk in $S^1 \times I$ whose interior is disjoint from Λ. For the purposes of the following argument, these tiny circles can be ignored, since the assumption on Euler characteristic means there will be as many tiny circles on a component of Λ' as there are on the corresponding component of Λ (namely $1 - \chi$). Hence, with no loss of generality, we may assume that Λ (and Λ') consist entirely of collections of cones on (corresponding) subsets of vertices. (In particular, we henceforth ignore any components with no boundary.) See Figure 3.

Since Λ contains no circles there is a spanning arc of the cylinder $S^1 \times I$ that is disjoint from Λ. After an isotopy in $S^1 \times I$, we may as well assume the arc is vertical and then break up a neighborhood of this vertical arc into a sequence of p vertical strips $\alpha_i \times I \subset S^1 \times I, i = 1, \dots, p$, where each $\alpha_i \cap \alpha_{i+1}, i = 1, \dots, p-1$, is a single end point of each.

Now push the first component Λ_1 of Λ into a vertical cylinder C parallel to $S^1 \times I$ and, exploiting the fact that Λ_1 is just a cone on its end points, do this so that the vertices in $\partial\Lambda_1$ appear in their correct order in a vertical strip in C. Now move this vertical strip (and so Λ_1) to $S^1 \times I$ by moving the strip to $\alpha_1 \times I$. Similarly place the second component Λ_2 in the second strip $\alpha_2 \times I$ and continue through all of Λ. Call the resulting graph $\Lambda^{c(anonical)} \subset S^1 \times I$ and observe that the process we have described gives a homeomorphism of pairs $g : (N, \eta(\Lambda)) \to (N, \eta(\Lambda^c))$. Finally, observe that the process is so canonical that if we had done the same process on Λ' we would have obtained a homeomorphism of pairs $g' : (N, \eta(\Lambda')) \to (N, \eta(\Lambda^c))$ that preserves the orderings. In particular g and g' could be taken to be the same on $\eta(\partial\Lambda)) = \eta(\partial\Lambda'))$. Then $g^{-1}g$ is the required homeomorphism of pairs. □

If Λ is unknotted, then M_Λ has a particularly simple structure:

Lemma 3.5. *An unknotted graph complement is a connected sum of handlebodies.*

Note: Here we regard a 3-ball as a handlebody of genus 0.

Proof. The case in which the ambient manifold is B^3 is representative. We have $\Lambda \subset B^2 \subset B^3$. In a small neighborhood of Λ collapse a forest that is maximal among all forests not incident to $\partial\Lambda$. Then each component is the cone on its boundary vertices, wedged with some circles. Each circle can be pushed out of B^2 (rel its wedge point) and made to bound a tiny disk. Removing such a circle from Λ has the effect in the graph complement of removing a 1-handle, dual to the tiny

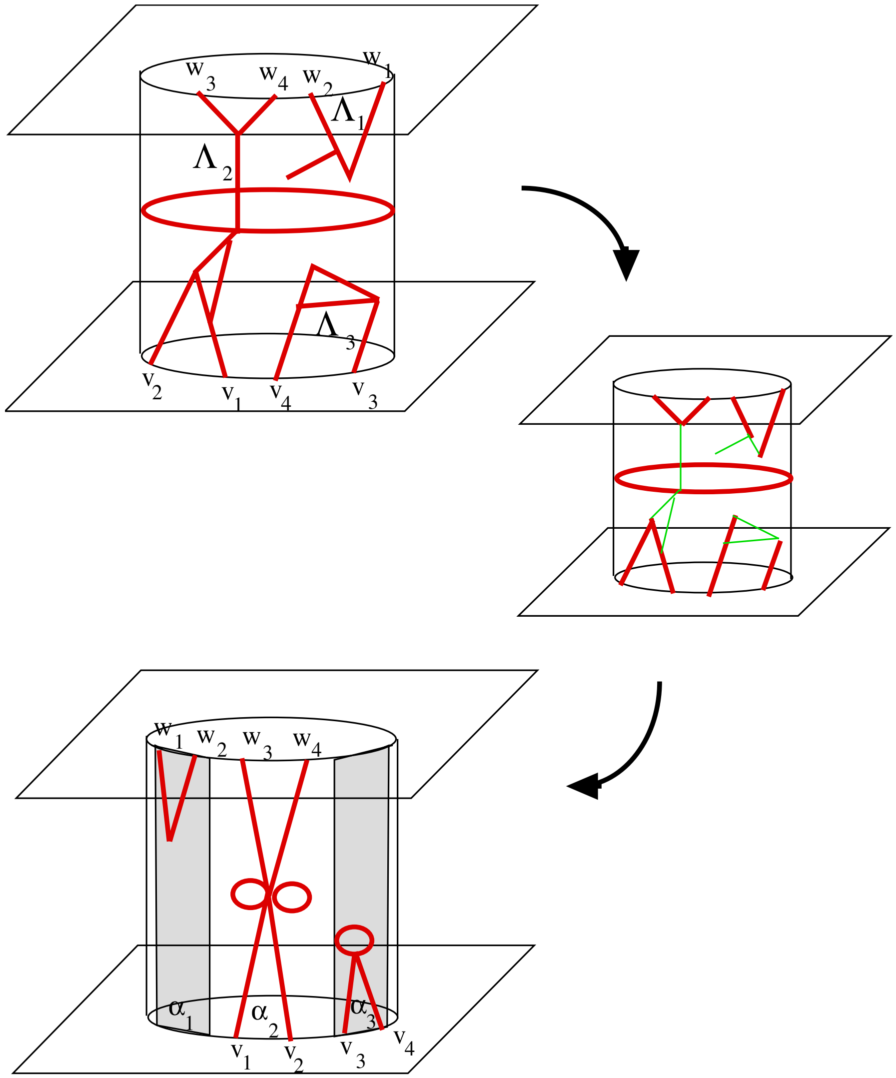

FIGURE 3. Reimbedding Λ canonically

disk. In particular, with no loss, we can assume that no such circles arise and so each component of Λ is a cone on its boundary vertices.

The proof is then by induction on $|\partial\Lambda|$. If $\partial\Lambda = \emptyset$, then Λ is a collection of isolated vertices, so its complement is a connected sum of balls. If any component of Λ has a single boundary vertex, then that component is just an arc with one end on ∂B^3; removing it from Λ has no effect on the complement in B^3. So without loss assume each component of Λ is the cone on two or more points in ∂B^2. A path in Λ between two such points divides the disk B^2 into two disks. An outermost such path will cut off a disk D from B^2 whose interior is disjoint from Λ. The disk D can be used to ∂-reduce M_Λ and the effect on M_Λ is the same as if we had removed one of the edges of Λ incident to D. The proof then follows by induction. $\square$

We now describe a few situations that guarantee that a graph is unknotted in $S^2 \times I$. We will be taking the standard height function on $S^2 \times I$, namely projection to I. A vertex v in a properly embedded graph $\Lambda \subset S \times I$ is a *Y-vertex* if two or more edges are incident to v from above and a *λ-vertex* if two or more edges are incident to v from below. (A vertex may be both a λ-vertex and a Y-vertex, or neither.)

Example 3.6. Suppose $\Lambda \subset S^2 \times I$ is a properly embedded graph so that

(1) the edges in Λ are all monotonic with respect to the projection $S^2 \times I \to I$,
(2) there are no Y-vertices.

Then Λ is an unknotted graph.

Proof. We first simplify Λ up to graph equivalence. By a small edge-slide arrange that each vertex is incident to at most two edges below; any vertex that is incident to a single edge above and a single edge below can be ignored. If an interior vertex is incident to two edges below, and none above, then add a small vertical edge above. After these initial maneuvers, each interior vertex of Λ has valence zero, one or three; in the last case, the vertex is a λ-vertex.

Pick a circle C in $S^2 \times \{1\}$ that contains all the vertices of $\partial\Lambda$ that lie in $S^2 \times \{1\}$. As $t \in [0, 1]$ descends, the monotonicity of edges means that, until another vertex of Λ is encountered, the cross-section $\Lambda \cap (S^2 \times \{t\})$ is a collection of points moving by isotopy in S^2. Extend the isotopy to all of S^2 to get a continuously varying circle $C_t \subset S^2 \times \{t\}$ that contains all of $\Lambda \cap (S^2 \times \{t\})$. When a valence one vertex (or an isolated vertex) is encountered, it can be easily added to or deleted from C_t, as appropriate, depending on whether the edge incident to the vertex is incident from below or from above.

So we only need to worry about λ-vertices. As t passes through the level of such a vertex (which we have arranged to lie in C_t), a single point in C_t simply splits in two and we may incorporate the arc between the two points as part of C. Continue the process down to $t = 0$. Now, in a standard argument, the continuously varying family of circles C_t bounds a continuously varying family of disks in S^2 and so there is a height-preserving isotopy of C_t to the standard $S^1 \times I$. □

The fact that, in the proof, the original circle C was ours to choose immediately leads to these additional examples:

Example 3.7. Suppose $\Lambda \subset S^2 \times I$ is a properly embedded graph and there is a generic height $t \in I$ so that

(1) the edges in Λ are all monotonic with respect to the projection $S^2 \times I \to I$,
(2) there are no λ-vertices above t,
(3) there are no Y-vertices below t.

Then Λ is an unknotted graph.

Proof. Apply the argument of Example 3.6 separately to $S^2 \times [0, t]$ and (upside down) to $S^2 \times [t, 1]$, starting with a circle in $S^2 \times \{t\}$ that contains all points in $\Lambda \cap (S^2 \times \{t\})$. See Figure 4. □

More generally

Example 3.8. Suppose $\Lambda \subset S^2 \times I$ is a properly embedded graph so that

(1) there is a generic level sphere $S^2 \times \{t\}$ for which Λ intersects $S^2 \times [t, 1]$ in an unknotted graph,

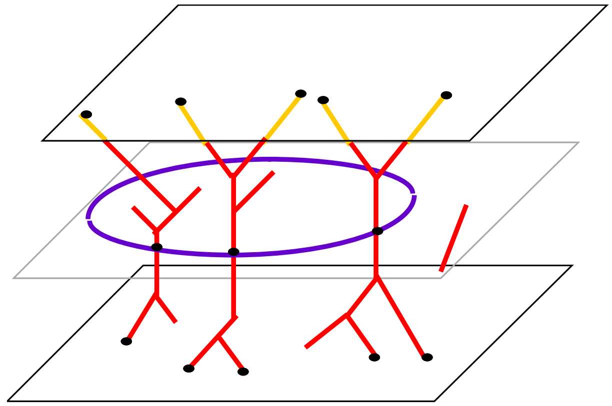

FIGURE 4.

(2) the edges in $\Lambda \cap [0,t]$ are all monotonic with respect to the projection to $[0,t]$,
(3) There are no Y-vertices in $\Lambda \cap [0,t]$.

Then Λ is an unknotted graph.

Proof. Apply the argument of Example 3.6 separately to $S^2 \times [0,t]$ starting with the circle in $S^2 \times \{t\}$ which is the base of the vertical cylinder in $S^2 \times [t,1]$ on which $\Lambda \cap (S^2 \times [t,1])$ lies. □

The next two examples simply reinterpret earlier examples in light of Lemma 3.2.

Example 3.9. Suppose (M,h) is a planar presentation of a manifold M and for an interval J, M^J is a component of $h^{-1}(J)$. Suppose all saddles in M^J are nested and that all lower saddles occur at higher levels than all the upper saddles do. Then the pair $(M^J, P^{\partial J})$ is an unknotted graph complement.

Suppose (M,h) is a planar presentation of a manifold M and $M^{[a,c]}$ is a component of $h^{-1}([a,c])$. We will use the following notation: for J a subinterval of $[a,c]$ let $M^J = M^{[a,c]} \cap h^{-1}(J)$ and for $t \in [a,c]$ let $Q^t = M^{[a,c]} \cap P^t$.

Example 3.10. Suppose (M,h) is a planar presentation of a manifold M and $M^{[a,c]}$ is a component of $h^{-1}([a,c])$.

Suppose that for some $b \in [a,c]$

(1) Q^b is connected,
(2) the pair $(M^{[b,c]}, Q^b \cup Q^c)$ is an unknotted graph complement,
(3) all saddles in $M^{[a,b]}$ are nested upper saddles.

Then $M^{[a,c]}$ is an unknotted graph complement.

It would be useful to know that if $\Lambda_1 \subset B_1$, $\Lambda_2 \subset B_2$ are unknotted graphs in 3-balls B_i, and we are given some identification of $\partial\Lambda_1$ with $\partial\Lambda_2$, that we could find some way to attach ∂B_1 to ∂B_2 consistent with that identification so that the resulting graph is unknotted. Ultimately we will succeed (see Lemma 4.1) but first we observe that the most obvious way to try to prove this fact is doomed to fail. Specifically, it may be impossible to match up the boundary of a disk in B_1

containing Λ_1 to the boundary of a disk in B_2 containing Λ_2 in a manner that preserves the identification of $\partial\Lambda_1$ with $\partial\Lambda_2$.

To see that this is impossible, take the following simple example: let each Λ_i be three copies of a cone on three points, so that $\partial\Lambda_i$ is nine points. Identify $\partial\Lambda_1$ with $\partial\Lambda_2$ so that the resulting graph is the complete bipartite graph $K_{3,3}$. If one could identify the boundary of a disk containing Λ_1 with the boundary of a disk containing Λ_2 in a way consistent with the identification of $\partial\Lambda_1$ with $\partial\Lambda_2$, we would have found an embedding of $K_{3,3}$ into the 2-sphere, which is famously impossible.

Yet there is a way to attach ∂B_1 to ∂B_2 so that $\Lambda_1 \cup_\partial \Lambda_2$ is an unknotted embedding of $K_{3,3}$ in S^3; the argument above merely shows that, in order to demonstrate that such an embedding is unknotted, edges will need to be slid over edges, inevitably across the sphere ∂B_i. In other words, the demonstration that there is an unknotted embedding of $K_{3,3}$ is inevitably a bit harder than one might at first expect.

It will be extremely useful to demonstrate that *any* bipartite graph has an unknotted embedding in S^3, via a construction much as above. That is the goal of the following lemma. Recall that a bipartite graph with vertex sets A and B is a graph so that each edge has one end among the vertices of A and the other end among the vertices of B. We will show that any bipartite graph can be imbedded in a very controlled way into a cube so that the embedded graph is unknotted: that is, *after some edge slides* the graph can be made to lie in a plane. Some details of its structure will be crucial in the discussion of braid equivalence in Section 4.

Lemma 3.11. *Let Λ be a finite bipartite graph, with vertex sets A and B. Then there is an embedding of Λ in the cube $I \times [-1,1] \times I$ so that:*

1. $A = \{(i/|A|, -1, 0), i = 0, \dots, |A|-1\}$.
2. $B = \{(j/|B|, 1, 0), j = 0, \dots, |B|-1\}$.
3. *Each edge in Λ is monotonic with respect to the y-coordinate. That is, each edge projects to $[-1,1]$ with no critical points.*
4. *The edges may be isotoped and slid over each other (perhaps destroying the bipartite structure) in the cube, so that afterwards the resulting graph lies entirely in the face $I \times [-1,1] \times \{0\}$.*

Moreover, given a specific edge e in Λ, such an embedding can be found so that $e = \{0\} \times [-1,1] \times \{0\}$ and e never moves during the isotopy.

Remark. Note that the last numbered condition implies that Λ is an unknotted graph, since this property is unchanged by edge slides. (Technically, Λ is unknotted only in a larger cube, for the given cube contains e in a face and so does not contain Λ as a proper subgraph.)

Proof. We will assume Λ is connected; if not, the following argument can be carried out in each component separately.

Place the designated edge e as described. Denote its ends by $a_0 = (0,-1,0)$ and $b_0 = (0,1,0)$. The Λ-distance between two vertices in Λ will mean the number of edges in the shortest path between them. With no loss, order the indices of the remaining vertices $a_i, 1 \le i \le |A|-1$, of A subordinate to their Λ-distance from a_0, i.e. so that, for any pair of indices i_1 and i_2, if a_{i_1} is closer in Λ to a_0 than a_{i_2} is, then $i_1 < i_2$. (We do not care how vertices are ordered among those that are Λ-equidistant from a_0.) Similarly order the indices of the remaining vertices $b_j, 1 \le j \le |B|-1$, of B subordinate to their Λ-distance from a_0. After this reordering, place each a_i at the point $(i/|A|, -1, 0)$ and each b_j at the point $(j/|B|, 1, 0)$.

At each vertex of Λ add a vertical (i.e. z-parallel) arc of length 1. That is, attach to each $(i/|A|, -1, 0)$ the arc $\{(i/|A|, -1)\} \times [0, 1]$ and to each $(j/|B|, 1, 0)$ the arc $\{(j/|B|, 1)\} \times [0, 1]$. In order to somewhat simplify the description of Λ, the edges of Λ will originally be placed so that they are horizontal (i.e. parallel to the $x - y$ plane) with ends on these vertical arcs. Λ is then finally recovered from the simplified description by collapsing the vertical arcs $\{(i/|A|, -1)\} \times [0, 1]$ and $\{(j/|B|, 1)\} \times [0, 1]$ back down to A and B, respectively.

Let ℓ be the maximal Λ-distance of any vertex in Λ from a_0. We will place the edges of Λ in a sequence of ℓ stages; the edges placed at the k^{th} stage lie near the horizontal square $I \times [-1, 1] \times \{k/\ell\}$. Specifically, at the k^{th} stage select all edges of Λ which have the property that their most Λ-distant end is a Λ-distance k from a_0. (The other end of each selected edge must then be Λ-distance $k - 1$ from a_0.) If there are p such edges, select a sequence of p horizontal planes whose height (i.e. z-coordinate) is near k/ℓ and place each edge in a separate horizontal plane, as a linear edge connecting the appropriate a_i to the appropriate b_j, with parallel edges on adjacent horizontal planes. The linear embedding ensures that each edge is monotonic in the y-coordinate, a fact that is unchanged when the vertical arcs $\{(i/|A|, -1) \times [0, 1]$ and $(j/|B|, 1)\} \times [0, 1]$ are collapsed to A and B to create Λ. We have thereby described an embedding of Λ into the cube that clearly satisfies the first three requirements. See Figure 5.

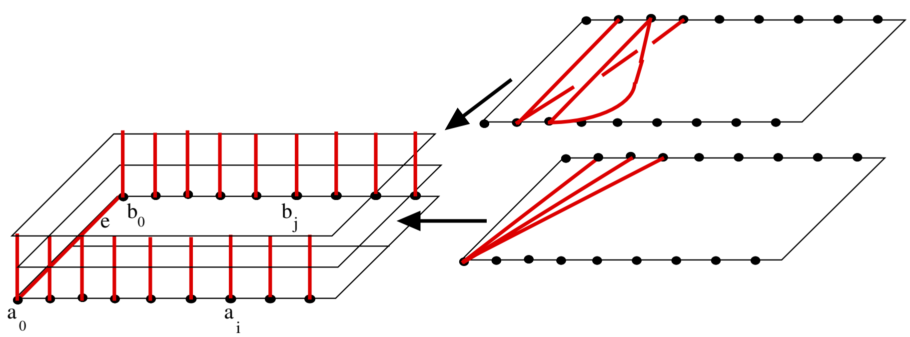

FIGURE 5. Putting Λ in the cube in layers

It remains to describe how the edges of Λ can be slid and isotoped, without moving the vertices A, B or the edge e, so that afterwards the resulting graph lies entirely in the face $I \times [-1, 1] \times \{0\}$. The description of this sliding mimics the k stages of the construction of Λ and we will describe them in the graph above as if we had not collapsed the vertical arcs, but also mostly focusing on the $x - y$ coordinates.

At the first stage of the construction above, exactly those edges with one end on a_0 are added, near the horizontal plane $z = 1/\ell$. By our choice of ordering of the b_j, the other ends of these edges lie exactly on the vertices $b_0, \ldots, b_q$, for some $q \geq 0$. (If any two of these edges are parallel, slide one over the other to form a tiny circle which we may henceforth ignore). Then, if $q > 0$ the rightmost edge, i.e. that connecting a_0 to $b_q = (q/|B|, 1)$, may be slid over the edge connecting a_0

to b_{q-1} until instead it is just the straight interval between b_{q-1} and b_q, i.e. the interval $[q-1,q] \times \{1\}$. Continue in this manner until all the edges but e have been slid to the line $y = 1$ to constitute the single interval $[0,q] \times \{1\}$, still in the plane $z = 1/\ell$. See Figure 6. Now slide all these edges up vertically to height just below $z = 2/\ell$ and begin the second stage.

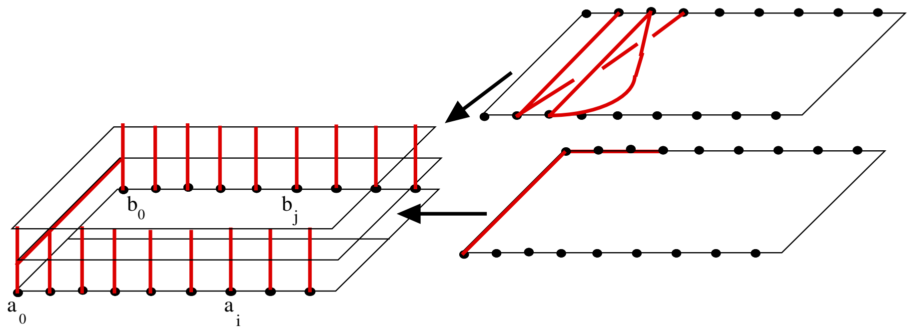

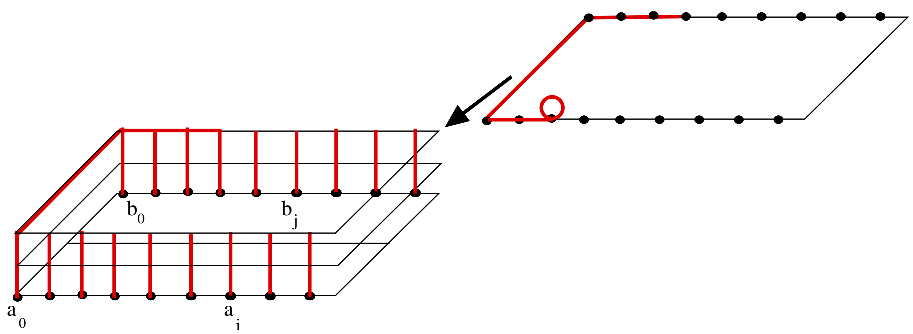

Figure 6. Sliding Λ in the cube by layer

Because of our ordering of the a_i, there is a $p \geq 1$ so that the vertices $a_1, \ldots, a_p$ constitute exactly the ends in A of edges included at the second stage. Moreover the other end of each such edge lies among the $b_0, \ldots, b_q$ which, after the slides we have done on the edges of the first stage, all lie on the L-shaped graph $e \cup ([0,q] \times \{1\})$. This L-shaped graph gives a way, much as above, of sliding the edges added at the second stage until they are either tiny circles (henceforth ignored) or constitute the straight line from a_0 to a_p, i.e. the line $[0,p] \times \{-1\} \times 2/\ell$. See Figure 6. Now slide this whole graph vertically up until it is near the plane $z = 3/\ell$ and continue the process. By the time we have reached the ℓ^{th} stage, the graph consists (now at height $z = 1$) of arcs in the lines $y = \pm 1$ that contain all the vertices, together with the original edge e between a_0 and b_0 (and some tiny circles), all of which then lie in the square $I \times [-1,1] \times \{1\}$. Now collapse the vertical direction, bringing the graph down to $I \times [-1,1] \times \{0\}$. This process (when reinterpreted as slides on the actual embedding of Λ, in which the vertical arcs do not appear) verifies the last numbered condition. $\square$

4. Braid equivalence and unknotted graphs

Suppose (M, h) is a planar presentation and $t \in R$ is a regular value of h. Cut M along P^t and reattach the two copies of P^t by an orientation preserving homeomorphism $P^t \to P^t$ that is the identity on the circles ∂P^t. The result is a possibly new manifold M' and a planar presentation $h' : M' \to R$. Note that $h'|\partial M' = h|\partial M$. The two planar presentations (M, h) and (M', h') are called *braid equivalent*. More generally, two planar presentations (M, h) and (M', h') are called braid equivalent if one is obtained from the other by a finite sequence of such operations, called braid moves.

Under such braid moves, many more 3-manifolds with planar presentation can be made unknotted graph complements. The following lemma illustrates why. The setting is this: Suppose N_A and N_B are each homeomorphic to either B^3 or $S^2 \times I$ and P_A (resp. P_B) is a sphere component of the boundary of N_A (resp N_B). Let N be obtained by identifying P_A with P_B (so in particular N is homeomorphic to S^3, B^3 or $S^2 \times I$). Suppose further that $\Lambda \subset N$ is a properly embedded graph that is in general position with respect to $P_A = P_B$; let $\Lambda_A = \Lambda \cap N_A$ and $\Lambda_B = \Lambda \cap N_B$.

Lemma 4.1. *If both Λ_A and Λ_B are unknotted graphs, then there is a homeomorphism $\phi : P_A \to P_B$ such that*

1. *ϕ coincides with the original identification of P_A with P_B near the points $\Lambda \cap P_A$ and*
2. *$\Lambda_A \cup \Lambda_B$ is an unknotted graph in $N_A \cup_\phi N_B$.*

Proof. The case in which both N_A and N_B are copies of $S^2 \times I$ is representative (and in fact the most difficult), and it will be convenient to take $N_A = S^2 \times [-2, 0]$ and $N_B = S^2 \times [0, 2]$.

Construct an abstract bipartite graph G with vertex sets A and B as follows: There is a vertex in A (resp. B) for every component of Λ_A (resp. Λ_B). There is an edge for every point c in $P_A \cap \Lambda = P_B \cap \Lambda$. Identify the ends of the edge corresponding to c to the points in A and B corresponding to the components in Λ_A and Λ_B on which c lies. Imbed G in the cube $I \times [-1, 1] \times I$ as described in Lemma 3.11 and embed the cube in $S^2 \times [-2, 2]$ with the $x - z$ square cross-section of the cube lying in the S^2 factor and the y-coordinate of the cube projecting to the interval factor in $S^2 \times [-1, 1] \subset S^2 \times [-2, 2]$.

For each vertex v in $\partial\Lambda \cap (S^2 \times \{-2\})$ add a monotone edge $e_v \subset S^2 \times [-2, -1]$ to G with one end of e_v on v and the other end on the vertex in A corresponding to the component of Λ_A on which v lies. Similarly, add a monotone edge in $S^2 \times [1, 2]$ for each vertex in $\partial\Lambda \cap (S^2 \times \{2\})$, with one edge on the vertex and the other on the appropriate vertex in B. Call the resulting graph G_+. See Figure 7.

The graph G_+, as embedded, has three important properties:

- It follows from the Remark following Lemma 3.11 that G_+ is an unknotted graph in N.
- It follows from Proposition 3.4 and Example 3.7 that (perhaps after wedging some tiny circles to G_+, and adding tiny bouquet-of-circle components to G_+) the graphs $G_+ \cap N_A$ and Λ_A are equivalent unknotted graphs, via an equivalence that is the identity near $\partial\Lambda_A$.
- Similarly (again perhaps after wedging on tiny circles and adding tiny bouquets of circles), the graphs $G_+ \cap N_B$ and Λ_B are equivalent unknotted graphs via an equivalence that is the identity near $\partial\Lambda_B$.

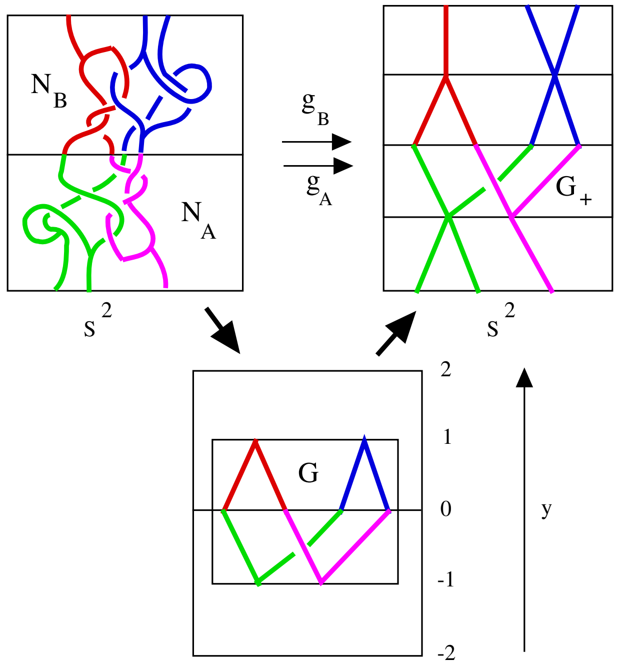

FIGURE 7. A braid move makes Λ unknotted

Let $g_A : P_A \to P_A$ and $g_B : P_B \to P_B$ be the homeomorphisms given by the latter two equivalences. Let $\phi = g_B g_A^{-1} : P_A \to P_B$. Then the construction $N_A \cup_\phi N_B$ changes Λ to a graph equivalent to G_+, which is unknotted. $\square$

This has as an immediate corollary, analogous to Example 3.10. Suppose (M, h) is a planar presentation of a manifold M and $M^{[a,c]}$ is a component of $h^{-1}([a, c])$. We again will use the following notation: for J a subinterval of $[a, c]$ let $M^J = M^{[a,c]} \cap h^{-1}(J)$ and for $t \in [a, c]$ let $Q^t = M^{[a,c]} \cap P^t$.

Corollary 4.2. *Suppose (M, h) is a planar presentation of a manifold M and $M^{[a,c]}$ is a component of $h^{-1}([a, c])$.*

Suppose that for some $b \in [a, c]$

1. *Q^b is connected,*
2. *the pair $(M^{[b,c]}, Q^b \cup Q^c)$ is an unknotted graph complement,*
3. *all saddles in $M^{[a,b]}$ are nested.*

Then $M^{[a,c]}$ is braid-equivalent to an unknotted graph complement.

Proof. The proof is by induction on the number of critical points of h on ∂M that occur in $M^{[a,b]}$. If there are none, then of course $M^{[a,c]} \cong M^{[b,c]}$ and there is nothing to prove. If the highest singularity in $M^{[a,b]}$ is a maximum or a minimum (necessarily an internal max or min since Q^b is connected and all saddles in $M^{[a,b]}$ are nested), then for t just below the corresponding critical value, $M^{[t,c]}$ is a standard graph complement, and we are done by induction. Similarly, if the highest critical value in $[a, b]$ is a (nested) upper saddle, then apply Example 3.10 to complete the inductive step.

The only remaining case is when the highest critical point is a lower saddle, i.e. it suffices to consider the case in which the only critical point in $M^{[a,b]}$ is a single nested lower saddle. But even in the more general case that all saddles in $M^{[a,b]}$ are nested lower saddles, the proof is an immediate consequence of Lemma 4.1 and Example 3.9 with the latter applied to $M^{[a,b]}$, which has no upper saddles. □

Next we hope to understand what happens to planar presentations of unknotted graph complements at unnested saddles. So let a be a critical value with corresponding critical point x_0, an unnested saddle. For small ϵ, let $M^{[a-\epsilon,a+\epsilon]}$ be the component of $h^{-1}([a-\epsilon, a+\epsilon])$ that contains x_0. Here we have taken ϵ so small that x_0 is the only critical point of $h|\partial M$ in $M^{[a-\epsilon,a+\epsilon]}$. Then for, say, a lower saddle, $P^{a+\epsilon}$ intersects $M^{[a-\epsilon,a+\epsilon]}$ in two connected planar surfaces denoted P_1 and P_2 and $P^{a-\epsilon}$ intersects $M^{[a-\epsilon,a+\epsilon]}$ in a single connected planar surface P_3. The roles of $\pm\epsilon$ are reversed for an upper saddle. We will be interested only in the case in which each P_i separates M. The component of $M - P_i$ not containing the saddle point will be denoted M_i. See Figure 8.

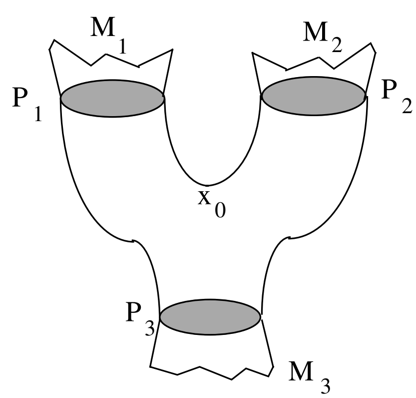

FIGURE 8.

The following two lemmas can be informally described as showing that when two of the M_i are known to be unknotted graph complements, so is the result of their adjunction at the saddle point. The first (Lemma 4.3) does the easy case: when the M_i that are known to be unknotted graph complements both lie above (or both lie below) the saddle point. The second (Lemma 4.4) covers the more complicated case when one of the two M_i known to be an unknotted graph complement lies above the saddle and the other one lies below. In order to best connect to the framework of our earlier notation and discussion, it will be advantageous to argue the former case at a *lower* saddle and the second case at an *upper* saddle, but that strategy is just a convenience. The distinction between the lemmas is not whether we are at an upper or a lower saddle, but rather whether we are given information about components M_i lying on the same side (i.e. both above or both below) of the saddle or on opposite sides of the saddle.

Lemma 4.3. *If (M_1, P_1) and (M_2, P_2) are unknotted graph complements in B^3, then so is $(M - interior(M_3), P_3)$, the component of the complement of P_3 that contains both M_1 and M_2.*

Proof. As remarked above, we will construct the argument for the case of a lower saddle, but the proof extends immediately to an upper saddle by just flipping everything over.

A useful model of an unknotted graph in B^3 is this: In a cube $I \times I \times I$, let Λ be a subgraph of the square $I \times I \times \{1/2\}$ with a single boundary vertex on the top $I \times \{1\} \times \{1/2\}$ and the rest on the bottom $I \times \{0\} \times \{1/2\}$. (Here projection to the y-coordinate models height h). The single vertex on the top reflects the description in Lemma 3.2 of how a planar presentation with a single external maximum gives rise to an unknotted graph complement in a ball (the center of the ball corresponds to the highest vertex). Furthermore, if we think of the vertex at the top of the box as stretching over the top and all the sides of the box, the planar part of the complement of Λ is precisely the bottom of the box, namely $(I \times \{0\} \times I) - \eta(\partial\Lambda)$. See Figure 9.

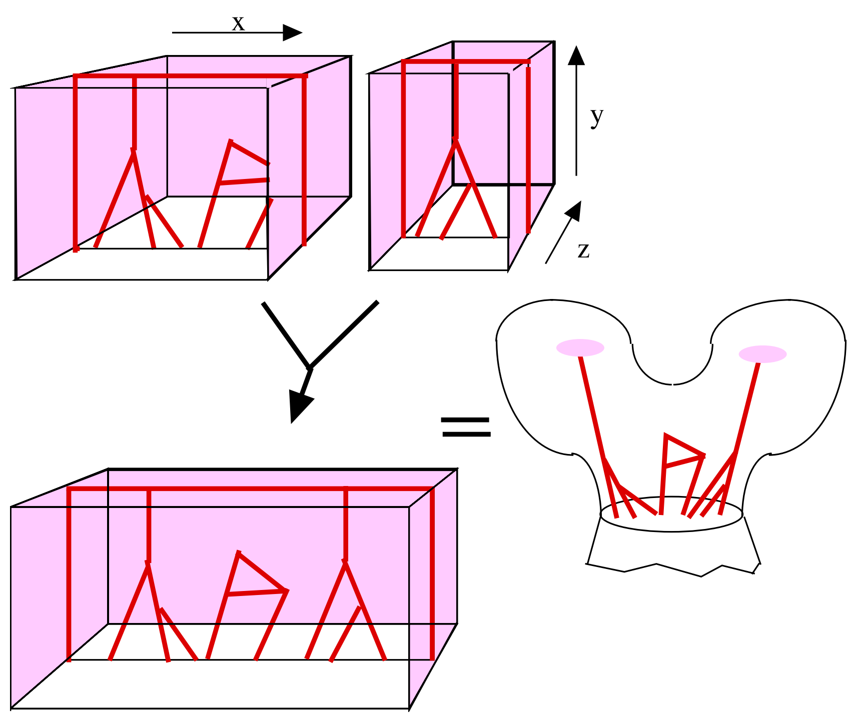

FIGURE 9. Modelling unknotted graph complements occurring above a lower saddle

The effect of passing through an unnested lower saddle is to take two such boxes (each containing one P_i on its bottom) and glue the side $\{1\} \times I \times I$ of one to the side $\{0\} \times I \times I$ of the other, obtaining a graph complement with planar part the boundary sum of the original two planar parts. The result is again a cube with the same sort of graph deleted, with the sole difference that now there are two boundary vertices of the graph on the top of the box. But since the top of the box is entirely disjoint from the planar part of the graph complement, up to graph complement equivalence, nothing is changed by sliding one top boundary vertex to the other along the top arc $I \times \{(1, 1/2)\}$, and then sliding an end of one edge down the end of the other, after which there is again a single boundary vertex on the top. In particular, the result is again an unknotted graph complement in the cube. □

A much harder situation to analyze is this:

Lemma 4.4. *If (M_3, P_3) and (M_2, P_2) are unknotted graph complements in B^3, then the pair $(M - interior(M_1), P_1)$ is braid equivalent to an unknotted graph complement in B^3.*

Proof. As remarked above, we will construct the argument for the case of an upper saddle, but the proof extends immediately to a lower saddle by just flipping everything over.

The initial difficulty is to determine a good model for what we are trying to show, analogous to the model in Lemma 4.3. Let Λ_A, Λ_B be unknotted graphs in the 3-ball whose complements give M_2 and M_3 respectively. Inspired by the model above (with the y-coordinate again modelling the height function h, but this time for an upper saddle) choose two cubes C_2, C_3 in R^3, as follows (see Figure 10):

(1) $C_2 = [0, 1] \times [-2, 0] \times [-1, 1]$,
(2) $C_3 = [0, 2] \times [0, 2] \times [-1, 1]$.

Let $C_\cup = C_2 \cup C_3$, which is itself homeomorphic to a 3-ball.

Construct an abstract bipartite graph G with vertex sets A and B as follows: There is a vertex in A (resp. B) for every component of Λ_A (resp. Λ_B). There is an edge for every component c of ∂P_2. Identify the ends of such an edge to the points in A and B that represent the components of Λ_A and Λ_B on which c lies. Imbed G in the cube $I \times [-1, 1] \times I \subset C_\cup$ as described in Lemma 3.11 with the special edge e chosen to be that which corresponds to the boundary component of P_2 which is incident to the saddle singularity. (Notice that e lies on the face $\{0\} \times I \times I$ of $\partial C_\cup$.)

The vertices of B are strung out along the interval $I \times \{1\}$ in the $x - y$ plane, and all of them but a vertex of e lie in the interor of $C_\cup$. Add edges to G that connect these vertices of B linearly to the corresponding vertices in the interval $\{1\} \times [1, 2]$ in the $x - y$ plane. Explicitly, add an edge $e_j, j = 1, \ldots, |B| - 1$, that connects the point $b_j = (j/|B|, 1, 0) \in B$ to the point $(1, 2 - j/|B|, 0)$. Next add edges that connect these points linearly to a collection B' of points in the line $[1, 2] \times \{(0, 0)\} \subset \partial C_3$. This collection B' is chosen so that each point in B' corresponds to a boundary component of P_1, other than the one containing the saddle singularity. Equivalently, each point $b' \in B'$ corresponds to a vertex in $\partial\Lambda_B$ that does not also naturally correspond to a vertex in $\partial\Lambda_A$. Such a boundary vertex lies on a component of Λ_B to which a vertex b_j has been assigned; append a linear edge in $[1, 2] \times [0, 2] \times [-1, 1]$from the *other* end of e_j to b'. (We pick the ordering of B' in the interval $[1, 2] \times \{(0, 0)\}$ so that these edges do not intersect.) Finally, append an appropriate number of tiny circles to $G \cap C_2$ and $G \cap C_3$ so that each component has the same Euler number as the corresponding component of Λ_A and Λ_B. Let G_+ be the graph in $C_\cup$ given by this construction. Note that it is a proper graph in $C_\cup$ whose planar part $P_\cup$ we take to be $([1, 2] \times \{0\} \times [-1, 1]) - \eta(G_+)$, i.e. the complement of $G_+ \cup C_2$ in the bottom face of C_3. See Figure 10.

Let $P_{x-z} \subset R^3$ denote the plane $y = 0$. The graph G_+ has been constructed to have these properties:

(1) For $i = 2, 3$, the graph $\Gamma_i = G_+ \cap C_i$ is unknotted, with planar part $(P_{x-z} \cap C_i) - \eta(G_+)$, by Example 3.6.

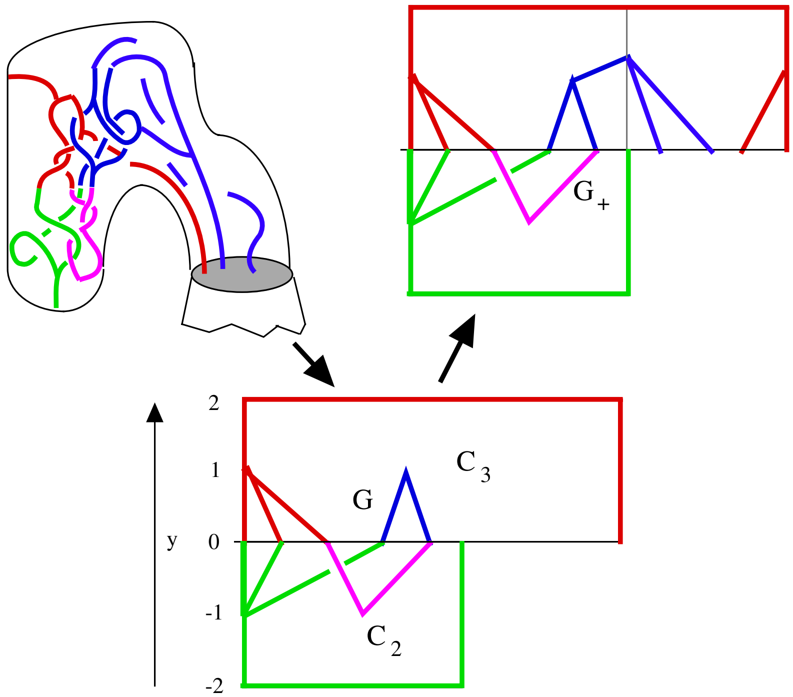

FIGURE 10. Modelling unknotted graph complements above and below an upper saddle

(2) Each component of Γ_2 (resp. Γ_3) is homeomorphic to a corresponding component of Λ_A (resp. Λ_B) so that the homeomorphisms agree, where they are simultaneously defined, namely on $\partial\Gamma_2 \subset \partial\Gamma_3$.
(3) The graph $G_+ \subset C_\cup$ is unknotted by Lemma 3.11.

The first two properties guarantee (via Proposition 3.4) that there is a homeomorphism of pairs $(M_i, P_i) \cong (C_i - \eta(G_+), (P_{x-z} \cap C_i) - \eta(G_+))$. In particular, much as in Lemma 4.1, M_3 can be cut off from M and reattached so that the pair $(M - interior(M_1), P_1)$ becomes pairwise homeomorphic to $(C_\cup - \eta(G_+), P_\cup)$. But since G_+ is unknotted, the latter is an unknotted graph complement. Hence $(M - interior(M_1), P_1)$ is braid equivalent to a standard graph complement. □

5. HEEGAARD REIMBEDDING

Theorem 5.1. *Suppose (M, h) is a planar presentation of a 3-manifold with connectivity graph a tree. Then (M, h) is braid-equivalent to an unknotted graph complement.*

Proof. If the connectivity graph Γ is a vertex (i.e. all saddles are nested) the result follows easily from Corollary 4.2. So we will assume that Γ has at least one edge. In that case, Lemma 4.1 demonstrates that the proof of the theorem will follow from the proof of the following relative version. □

Proposition 5.2. *Suppose (M, h) is a planar presentation of a 3-manifold and Γ is its connectivity graph. Suppose $\gamma \subset \Gamma$ is an edge such that a component $\Gamma_0 \subset \Gamma$ of the complement of γ is a tree. Let $P_\gamma \subset M$ be the planar surface corresponding to γ and let $M_0 \subset M$ be the component of $M - P_\gamma$ that corresponds to Γ_0. Then (M_0, P_γ) is braid equivalent to an unknotted graph complement.*

Proof. The proof will be by induction on the number of edges in Γ_0. Let v be the vertex of Γ_0 that is incident to γ and, in the terminology of Lemma 2.1, let M_v be the component of $M - \bigcup_{i=1}^{n} P^{s_i}$ corresponding to v, with $h(M_v) = [s_i, s_{i+1}]$. We will assume that the saddles at heights s_i and s_{i+1} both involve the particular component M_v, since the argument is easier if either or both do not.

Without loss of generality we will assume that the planar surface corresponding to the edge γ is at the bottom of M_v, i.e. near height s_i. Consider first the saddle x_+ at height s_{i+1}. Let Q be the connected planar surface $M_v \cap P^{(s_{i+1}-\epsilon)}$ and let M_Q be the component of $M - Q$ that contains x_+. If x_+ is an upper (unnested) saddle, then Q corresponds to an edge in Γ_0 and so by inductive assumption the pair (M_Q, Q) is braid equivalent to an unknotted graph complement. See Figure 11a. If x_+ is a lower saddle, then the two contiguous components of $h^{-1}(s_{i+1} + \epsilon)$ each represent edges in Γ_0 and (M_Q, Q) is again an unknotted graph complement by inductive assumption combined with Lemma 4.3. See Figure 11b. So in any case, (M_Q, Q) is braid equivalent to an unknotted graph complement.

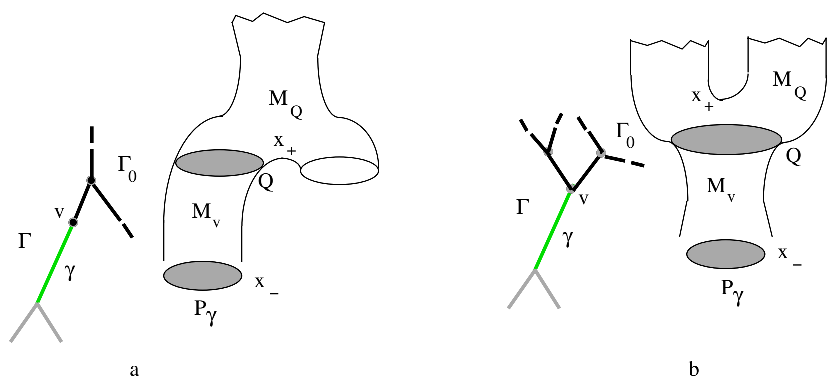

FIGURE 11.

Now consider the saddle x_- at height s_i. See Figure 12. If it is a lower saddle, then the planar surface $P^{s_i+\epsilon} \cap M_v$ is P_γ, the planar surface corresponding to the edge γ and the proposition follows from Corollary 4.2. If the saddle x_- is an upper saddle, then P_γ is one of the two connected planar surfaces in $P^{s_i-\epsilon}$ contiguous to the saddle. Let $P_{\gamma'}$ be the other one, with corresponding edge $\gamma' \subset \Gamma_0$, and let P be the connected planar surface $M_v \cap P^{s_i+\epsilon}$. Now by inductive assumption, the component of $M_0 - P_{\gamma'}$ not containing x_- is an unknotted graph complement and by Corollary 4.2 so is the component of $M_0 - P$ not containing x_-. Then the proposition follows from Lemma 4.4. $\square$

Corollary 5.3. *Suppose $N \subset S^3$, $p : S^3 \to R$ is the standard height function, N contains both poles, and the connectivity graph of $S^3 - N$ is a tree (so in particular*

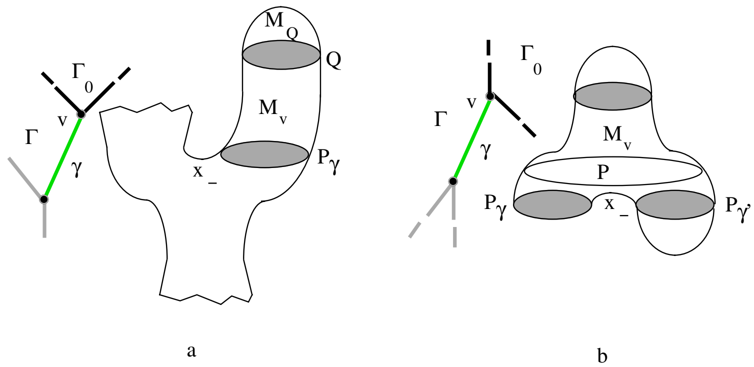

FIGURE 12.

$S^3 - N$ is connected). Then there is an embedding $f : N \to S^3$ so that

(1) *$p = pf$ on N, i.e. f preserves height and*
(2) *$S^3 - f(N)$ is a connected sum of handlebodies.*

Proof. It follows from Theorem 5.1 that $M = S^3 - N$ is braid-equivalent to a connected sum of handlebodies. We will show that a braid move on M defines a reimbedding of N.

Let S^t be the 2-sphere $p^{-1}(t)$ and let $P^t = S^t - N = S^t \cap M$. Then a braid move of M at a generic level t is given by cutting M open along P^t and then reattaching P^t to itself by a homeomorphism $\phi : P^t \to P^t$ that is the identity on ∂P^t. In particular, the homeomorphism ϕ extends via the identity on $S^2 - P^t$ to a self-homeomorphism of S^2. But any (orientation preserving) self-homeomorphism of the sphere is isotopic to the identity, so in fact there is a level-preserving self-homeomorphism $S^2 \times [t - \epsilon, t + \epsilon]$ that is the identity on one end and the extended ϕ on the other. Use this self-homeomorphism to redefine the embedding of N in the region $h^{-1}[t - \epsilon, t + \epsilon]$. The effect on the complement M is to do the original braid move. □

Corollary 5.4. *Suppose $p : S^3 \to R$ is the standard height function and $H \subset S^3$ is a handlebody for which horizontal circles constitute a complete collection of meridian disk boundaries. Then there is a reimbedding $f : H \to S^3$ so that*

(1) *$p = pf$ on N, i.e. f preserves height and*
(2) *$H \cup (S^3 - H)$ is a Heegaard splitting of S^3.*

Proof. As noted before Proposition 2.3, we may as well assume that H contains both poles. The condition on horizontal disks guarantees, via Proposition 2.3, that the connectivity graph of $S^3 - H$ is a tree. Then Corollary 5.3 says there is a height-preserving reimbedding of H so that $S^3 - H$ is a connected sum of handlebodies. But since ∂H is connected, $S^3 - H$ is in fact simply a handlebody. □

6. KNOT WIDTH

For standard definitions about knots in S^3, see [BZ], [L] or [R].

Definition 6.1. As above, let $p : S^3 \to R$ be the standard height function and let S^t denote $p^{-1}(t)$, a sphere if $|t| < 1$. Let $K \subset S^3$ be a knot in general position with respect to p and let $c_1, \ldots, c_n$ be the critical values of $h = p|K$ listed in increasing order; i.e., so that $c_1 < \cdots < c_n$. Choose $r_1, \ldots, r_{n-1}$ so that $c_i < r_i < c_{i+1}, i = 1, \ldots, n-1$. The *width of* K *with respect to* h, denoted by $w(K, h)$, is $\sum_i |K \cap S^{r_i}|$. The *width of* K, denoted by $w(K)$, is the minimum of $w(K', h)$ over all knots K' isotopic to K. We say that K is in *thin position* if $w(K, h) = w(K)$.

We note as an aside that there is an alternative way to calculate width, inspired by a comment of Clint McCrory. For the levels r_i described above, call r_i a *thin level* of K with respect to h if c_i is a maximum value for h and c_{i+1} is a minimum value for h. Dually r_i is a *thick level* of K with respect to h if c_i is a minimum value for h and c_{i+1} is a maximum value for h. Since the lowest critical point of h is a minimum and the highest is a maximum, there is one more thick level than thin level.

Lemma 6.2. *Let $r_{i_1}, \ldots, r_{i_k}$ be the thick levels of K and let $r_{j_1}, \ldots, r_{j_{k-1}}$ be the thin levels. Set $a_{i_l} = |\ K\ \cap S^{r_{i_l}}|$ and $b_{j_l} = |\ K\ \cap S^{r_{j_l}}|$. Then*

$$2w(K) = \sum_{l=1}^{k} a_{i_l}^2 - \sum_{l=1}^{k-1} b_{j_l}^2.$$

Proof. This can be proven by a direct computation and repeated use of the Gauss Summation Formula. It is illustrated in Figure 13. Each dot represents two points of intersection with a regular level surface between two critical level surfaces. For instance, the dots in Figure 13 represent the case in which the critical values, listed from the highest to the lowest, are a maximum, maximum, maximum, maximum, minimum, minimum, maximum, maximum, minimum, minimum, maximum, maximum, maximum, minimum, minimum, minimum, minimum, minimum. □

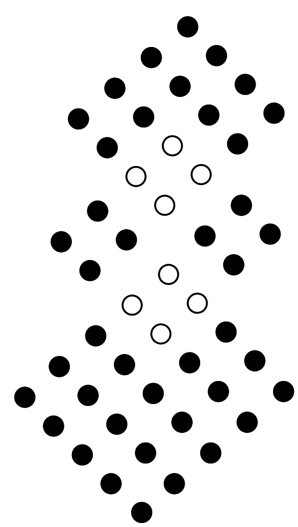

FIGURE 13. Dark dots indicate squares that are added; white dots indicate overlap squares that are subtracted

Corollary 6.3. *Suppose K is a knot in an unknotted solid torus $W \subset S^3$. Suppose $f : W \to S^3$ is a knotted embedding and $K' = f(K)$. Then $w(K') \geq w(K)$.*

Proof. Let $p : S^3 \to R$ be the standard height function. Isotope K' so as to minimize its width with respect to this height function and let H denote the image of $f(W)$ after this isotopy. Each generic 2-sphere $S^t = p^{-1}(t)$ intersects ∂H in a collection of circles, each of them unknotted since they all lie in S^t. By standard Morse theory, there must be a generic value of t for which one of the circles $c \subset \partial H \cap S^t$ is essential in ∂H, and that circle cannot be a longitude, since H is a knotted torus. Hence c must be a meridian circle. It follows from Corollary 5.4 that there is a reimbedding g of H in S^3 that preserves height but after which H is unknotted. The reimbedding is defined via braid moves on $M = S^3 - H$; after perhaps adding a number of Dehn twists to one of the braid moves near a meridinal boundary component of $P^t = M \cap S^t$, we can take this reimbedding to preserve a longitude of H. So in particular, $g(K')$ is isotopic to K in S^3 and still has the width of K'. □

Corollary 6.3 can be applied to composite knots, via the following standard construction. Let $K = K_1 \# K_2$ be a composite knot with decomposing sphere S. Then $\partial(S^3 - \eta(K \cup S))$ has two components. Each of these components is a torus, called a *swallow-follow torus*. Each of these tori bounds a solid torus in S^3 that contains K; the torus T_1 whose core is parallel to K_1 is said to follow K_1 and swallow K_2. Similarly, the other torus T_2 follows K_2 and swallows K_1. The torus T_1 exhibits K as a satellite knot of K_1 with pattern K_2, and symmetrically for T_2. Therefore, when Corollary 6.3 is applied to each T_i in turn, we get

Corollary 6.4. *For any two knots K_1, K_2,*

$$w(K_1 \# K_2) \geq max\{w(K_1), w(K_2)\} \geq \frac{1}{2}(w(K_1) + w(K_2)).$$

Of course the construction can be iterated to give

Corollary 6.5. $w(K_1 \# \cdots \# K_n) \geq max\{w(K_1), \ldots, w(K_n)\} \geq \frac{1}{n}(w(K_1) + \cdots + w(K_n))$.

Proof. For each K_i there is a torus that swallows K_i and follows the connected sum of the remaining summands. □

It appears that the inequality $w(K_1 \# K_2) \geq max\{w(K_1), w(K_2)\}$ may in fact be the best possible. It is argued in [ST] that there are knots K so that for any two-bridge knot L, $w(K \# L) = w(K)$ and, less persuasively, that similar examples can be constructed for any given bridge number higher than two.

References

[BZ] G. Burde, H. Zieschang, Knots, de Gruyter Studies in Mathematics 5, Walter de Gruyter & G., Berlin, 1985, ISBN: 3-11-008675. MR0808776 (87b:57004)

[Fo] R. H. Fox, *On the imbedding of polyhedra in 3-space*, Ann. of Math. **49** (1948), 462–470. MR0026326 (10:138c)

[G] D. Gabai, *Foliations and the topology of 3-manifolds. III*, J. Differential Geom. **26** (1987), 3, 479–536. MR0910018 (89a:57014b)

[L] W.R.B.R. Lickorish, An introduction to knot theory, Graduate Texts in Mathematics, 175, Springer-Verlag, New York, 1997, ISBN: 0-387-98254-X. MR1472978 (98f:57015)

[Mo] K. Morimoto, *There are knots whose tunnel numbers go down under connected sum*, Proc. Am. Math. Soc, **123** (1995), no. 11, 3527–3532. MR1317043 (96a:57022)

[MS] K. Morimoto, J. Schultens, *Tunnel numbers of small knots do not go down under connected sum*, Proc. Am. Math. Soc, **128** (2000), no. 1, 269–278. MR1641065 (2000c:57014)

[RS] Y. Rieck, E. Sedgwick, *Thin position for a connected sum of small knots*, Algebraic and Geometric Topology **2** (2002), 297-309. MR1917054 (2003d:57021)

[R] D. Rolfsen, Knots and Links, Mathematics Lecture Series, No. 7, Publish or Perish, Inc., Berkeley, Calif., 1976. MR0515288 (58:24236)

[Sc1] M. Scharlemann, *Handlebody complements in the 3-sphere: a remark on a theorem of Fox*, Proc. Amer. Math. Soc. **115** (1992), 1115–1117. MR1116272 (92j:57004)

[ScSc] M. Scharlemann, J. Schultens, *Annuli in generalized Heegaard splittings and degeneration of tunnel number*, Math. Ann. **317** (2000), no. 4, 783–820. MR1777119 (2001j:57013)

[ST] M. Scharlemann, A. Thompson, *On the additivity of knot width*, ArXiv preprint math.GT/0403326.

[S] H. Schubert, *Über eine numerische Knoteninvariante*, Math. Z. **61** (1954), 245-288. MR0072483 (17:292a)

[Sch] J. Schultens, *Additivity of bridge numbers of knots*, Math. Proc. Cambridge Philos. Soc. **135** (2003), 539–544. MR2018265 (2004i:57010)

[Th] A. Thompson, personal communication.

Department of Mathematics, University of California, Santa Barbara, California 93106

E-mail address: mgscharl@math.ucsb.edu

Department of Mathematics, University of California, Davis, California 95616

E-mail address: jcs@math.ucdavis.edu