

HARMONIC MAPS $M^3 \rightarrow S^1$ AND 2-CYCLES, REALIZING THE THURSTON NORM

GABRIEL KATZ

ABSTRACT. Let M^3 be an oriented 3-manifold. We investigate when one of the fibers or a combination of fiber components, F_{best} , of a *harmonic* map $f : M^3 \rightarrow S^1$ with Morse-type singularities delivers the Thurston norm $\chi_-([F_{best}])$ of its homology class $[F_{best}] \in H_2(M^3; \mathbb{Z})$.

In particular, for a map f with connected fibers and any well-positioned oriented surface $\Sigma \subset M$ in the homology class of a fiber, we show that the Thurston number $\chi_-(\Sigma)$ satisfies an inequality

$$\chi_-(\Sigma) \geq \chi_-(F_{best}) - \rho^\circ(\Sigma, f) \cdot \text{Var}_{\chi_-}(f).$$

Here the variation $\text{Var}_{\chi_-}(f)$ is can be expressed in terms of the χ_- -invariants of the fiber components, and the twist $\rho^\circ(\Sigma, f)$ measures the complexity of the intersection of Σ with a particular set F_R of “bad” fiber components. This complexity is tightly linked with the optimal “ \tilde{f} -height” of Σ , being lifted to the f -induced cyclic cover $\tilde{M}^3 \rightarrow M^3$.

Based on these invariants, for any Morse map f , we introduce the notion of its *twist* $\rho_{\chi_-}(f)$. We prove that, for a harmonic f , $\chi_-([F_{best}]) = \chi_-(F_{best})$ if and only if $\rho_{\chi_-}(f) = 0$.

1. INTRODUCTION

Let M be a compact, oriented 3-manifold, possibly with a boundary. With any homology class $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$, one can associate a number of interesting invariants. The first one, $g([\Sigma])$, is the minimum genus of an embedded (immersed) oriented surface $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ realizing the homology class.

Let Σ° stand for the union of all components of Σ , excluding spheres and disks. Put

$$\chi_-(\Sigma) = |\chi(\Sigma^\circ)| \quad \text{and} \quad \chi_+(\Sigma) = |\chi(\Sigma \setminus \Sigma^\circ)|,$$

where $\chi(\sim)$ is the Euler number.

The Thurston norm $\|[\Sigma]\|_T$ is defined to be the minimum of $\{\chi_-(\Sigma)\}$ over all embedded surfaces Σ representing the homology class $[\Sigma]$.

In general, the correspondence $\{[\Sigma] \Rightarrow \chi_-([\Sigma])\}$ gives rise to a *semi-norm* on the vector space

$$H_2(M, \partial M; \mathbb{R}) \approx H^1(M; \mathbb{R}).$$

In many cases, $\|\sim\|_T$ is actually a norm, with the unit ball in the shape of a convex polyhedron.

Received by the editors May 15, 2002 and, in revised form, October 10, 2003.

2000 *Mathematics Subject Classification*. Primary 57M15, 57R45.

Let \mathcal{F} be a codimension one, oriented foliation with no Reeb components in M or along its boundary. Such foliations are characterized by a global property: every leaf of \mathcal{F} is hit by a loop *transversal* to the foliation, and a similar transversal loop condition is satisfied by $\mathcal{F}|_{\partial M}$. In an appropriate metric, the leaves of \mathcal{F} are *minimal hypersurfaces* [Su]. Foliation with this property are called *taut*.

In [T] Thurston showed that any compact leaf of a taut foliation \mathcal{F} attains the *minimal* value of $\chi_-(\sim)$ in its homology class. On the other hand, Gabai proved that, if a surface $\Sigma \subset M$ minimizes the χ_- -value in its non-trivial homology class and has no toral components, then it is a compact leaf of a smooth, taut foliation [G]. Thus, surfaces which realize the norm $\|\sim\|_T$, are compact leaves of taut foliations.

In general, taut foliations are hard to construct. In contrast, closed, or even harmonic differential forms are easy to produce. If an oriented foliation \mathcal{F} is generated by the kernels of a closed, *non-singular* 1-form ω , the foliation is automatically taut. Then, all the leaves of \mathcal{F} are non-compact, or alternatively, they all are compact. In the second case, M fibers over a circle, and \mathcal{F} is comprised of the fibers of the corresponding map $f_\omega : M \rightarrow S^1$. In fact, ω and f_ω are *harmonic* in an appropriate metric. In this setting, the harmonicity of a 1-form is equivalent to the tautness of the associated foliation. As a result, compact leaves of a foliation generated by a *harmonic non-singular* form realize the Thurston norm of their homology class.

Closed 1-forms with singularities produce *singular* foliations, which exhibit rich and drastically different behavior from the classical non-singular species [FKL]. In the paper, we will be concerned with the foliations generated by *harmonic* 1-forms with *Morse-type singularities* and *rational* periods. Although their topology is very different from the non-singular foliations, they still possess the transversal loop property [C], and their leaves are *near-minimal* [K] (the harmonically-generated singular foliations are “near-taut”).

There is a homological version of harmonicity (described in Theorem 4.7) which plays an important role in our arguments. Maps with one connected fiber are intrinsically harmonic.

Computing the Thurston norm $\|[\Sigma]\|_T$ in terms of the topology of M alone is very difficult. The idea is to employ an appropriate map $f : M \rightarrow S^1$ to get a handle on the problem. It is natural to start with maps whose fibers realize $[\Sigma]$. The geometry of f allows us to determine the χ_- -invariant of each fiber component. At least, among all the combinations of fiber components, we can pick a representative of $[\Sigma]$ with the minimal value of $\chi_-(\sim)$. This gives rise to an “ f -vertical” seminorm $\|\sim\|_{H^f}$ on the subspace $H_2^f \subset H_2(M, \partial M; \mathbb{R})$ spanned by the fundamental classes of various fiber components. We call a combination of fiber components which delivers $\|[\Sigma]\|_{H^f}$ the *best*, and denote it by F_{best} .

Our main goal is to understand the relation between the “incomputable” norm $\|[\Sigma]\|_T$ and the “computable” $\|[\Sigma]\|_{H^f}$. For instance, how can we tell when a map f has the property

$$\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}?$$

A somewhat different question can be investigated: “When is Σ delivering $\|[\Sigma]\|_{H^f}$ realizable by a genuine f -fiber, and not by a union of fiber components?” Figure 1 (in §2) shows a map which is not intrinsically harmonic and for which F_{best} —a union of two spherical fiber components—is distinctly different from any fiber. Answering both questions will allow us to characterize maps for which a fiber delivers the Thurston norm.

This article is a by-product of my unsuccessful attempts to establish an analog of the Thurston theorem for harmonically-generated foliations *with singularities*. An important case of such foliations is provided by generic harmonic maps to a circle, that is, by generic rational harmonic 1-forms on M .

For some time, I believed that, for a harmonic map $f : M \rightarrow S^1$, $\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}$ — the best union of fiber components realizes the Thurston norm of its homology class. All my efforts to prove this very naive conjecture (by employing the theory of minimal surfaces) failed, until I found a simple counter-example (cf. Example 4.10 —the *Harmonic Twister*). Although too weak on its own, some form of harmonicity seems to be a valuable ingredient in any “best fiber component theorem” (cf. Corollaries 8.3, 8.8): we always assume that our maps f have no local extrema.

In fact, the reality is as far from what I conjectured as it could be: there are harmonic maps $f : M \rightarrow S^1$ with very few singularities and with $\chi_-(F_{best})$ arbitrarily distant from the Thurston norm $\chi_-([F_{best}])$ (cf. Example 4.10).

The phenomenon occurs because maps can have arbitrary big “twists” $\rho_{\chi_-}(f)$. Crudely, the twist invariant $\rho_{\chi_-}(f)$ measures the minimal complexity of the intersection patterns of surfaces $\Sigma \subset M$, delivering the Thurston norm, with a generic fiber component. When $[\Sigma]$ is in the homology class of a fiber, $\rho_{\chi_-}(f)$ can be estimated from above by the minimal \tilde{f} -height of such a Σ , being appropriately lifted to the f -induced cyclic covering $\tilde{M} \rightarrow M$. Here \tilde{f} stands for a function on \tilde{M} covering the map f .

In a sense, one can also think about $\rho_{\chi_-}(f)$ as the S^1 -controlled size of a homotopy, which takes a given map f to a map f_1 with one of its fibers delivering the Thurston norm of its cohomology class (cf. Corollary 6.15).

Similar invariants can be introduced for any probe surface $\Sigma \subset M$ in a vertical homology class $[\Sigma] \in H_2^f$ (cf. Section 6). They are based on the twists $\rho(\Sigma, F)$ which measure the complexity of the intersection pattern $\mathcal{C} := \Sigma \cap F$ inside a generic fiber component F . The number $\rho(\Sigma, F) + 1$ does not exceed the number of components into which \mathcal{C} divides the surface F . By ignoring the components of \mathcal{C} which bound a disk in F , a modification $\rho^\circ(\Sigma, F)$ of $\rho(\Sigma, F)$ is introduced.

In the Introduction, we use $\rho^\circ(\Sigma, f)$ to denote $\max_F \{\rho^\circ(\Sigma, F)\}$, where F runs over all possible f -fiber components.

When $[\Sigma]$ is chosen to be the homology class of a fiber F , the quantity $\rho(\Sigma, F)$ admits an interpretation as the *breadth* $b(F, \Sigma)$ of a lifting $\hat{F} \subset \tilde{M}$ relative to a special lifting $\hat{\Sigma} \subset \tilde{M}$ of Σ (cf. Definition 6.7). It has an upper bound $h(\Sigma, f)$ which is defined to be the integral part of the minimal \tilde{f} -height of $\hat{\Sigma}$ plus one.

Given a Morse map $f : M \rightarrow S^1$, we consider a finite distribution of the values $\{\chi_-(f^{-1}(\theta))\}_{\theta \in S^1}$ along the circle. A number $\chi_-(f^{-1}(\theta))$ can jump only when θ crosses an f -critical value. We define the χ_- -variation of the function $\{\theta \rightarrow \chi_-(f^{-1}(\theta))\}$ by the formula (3.1). For maps f with no local extrema, the variation counts, so called, *non-bubbling* f -critical points (cf. Definition 3.2).

Of course, for any fibration f , $\text{var}_{\chi_-}(f) = 0$. However, if a harmonic f is not a fibration, then $\text{var}_{\chi_-}(f) = 0$ implies that the Thurston *semi*-norm is not a norm: some non-trivial classes in $H_2(M; \partial M; \mathbb{Z})$ are represented by 2-spheres or 2-disks.

We promote here a slogan:

“Maps f with the 0-variation are like fibrations over the circle”.

Finally, a ghost of our original *Best Fiber Conjecture* “ $\chi_-([F_{best}]) = \chi_-(F_{best})$ ” pays a visit:

Theorem 1.1. *Consider a Morse map $f : M \rightarrow S^1$ with all its fibers being connected.¹ Then:*

- *The χ_- -invariant of the best fiber F_{best} attains the minimal value among all the surfaces homologous to a fiber if and only if the twist $\rho_{\chi_-}(f) = 0$.*
- *In fact, $\text{var}_{\chi_-}(f) = 0$ implies $\rho_{\chi_-}(f) = 0$.*

Moreover, the same conclusions are valid for connected sums of maps with connected fibers.

This theorem is a very special case of our main results—Theorems 8.2, Corollary 8.3, Theorem 8.6, Corollary 8.7 and Corollary 8.14. To avoid technicalities, let us state these propositions for another special, but important class of maps to a circle—for the *self-indexing* harmonic maps (in fact, one can deform any map $f : M \rightarrow S^1$ into an intrinsically harmonic self-indexing map).

Given such a map f and any “probe” surface $\Sigma \subset M$ homologous to a fiber and well-positioned (cf. Definition 7.1) with respect to the “worst” fiber F_R , we have

$$(1.1) \quad \chi_-(\Sigma) \geq \chi_-(F_{best}) - \text{var}_{\chi_-}(f) \cdot \rho^\circ(\Sigma, f),$$

where $\text{var}_{\chi_-}(f)$ is the number of non-bubbling f -critical points.

Evidently, unless the defect $\text{var}_{\chi_-}(f) \cdot \rho^\circ(\Sigma, f)$ is smaller than $\chi_-(F_{best})$, this inequality is not very informative.

For harmonic self-indexing maps with a non-zero variation, (1.1) can also be viewed as giving some grip on the invariants $\rho_{\chi_-}(f)$ and $h_{\chi_-}(f)$ (cf. Section 6):

$$(1.2) \quad h_{\chi_-}(f) \geq \rho_{\chi_-}(f) \geq \frac{\chi_-(F_{best}) - \chi_-([F_{best}])}{\text{var}_{\chi_-}(f)}.$$

In other words, if $\chi_-(F_{best}) \gg \chi_-([F_{best}])$ and the number of f -critical points is small, then the χ_- -minimizing, well-positioned surface Σ must be very tall. That is, f “wraps” Σ many times around the circle. Also, such a Σ must have complex intersection patterns with a fiber — an intersection which is comprised of at least as many curves as the RHS of (1.2) requires.

Now we describe the organization of the paper. It is comprised of nine sections (including the Introduction) followed by a notation list.

In *Section 2* we consider *intrinsically harmonic* 1-forms and maps into a circle. One can associate a finite graph Γ_ω with any closed 1-form ω . The intrinsically harmonic forms and maps give rise to very special graphs. As a result, it is possible to faithfully express intrinsic harmonicity in pure combinatorial, graph-theoretical terms.

In *Section 3* we use graphs Γ_f as book-keeping devices to record the distribution of χ_- -invariants of the f -fibers and their connected components. The section deals with the effects of deforming a given map f on the graph Γ_f and these invariants.

In *Section 4* we develop further graph-theoretical manifestations of harmonicity (cf. Theorem 4.7). Examples of harmonic maps f , with only two singularities and with the norm $\|[F_{best}]\|_{H^f}$ being *arbitrarily distant* from the Thurston norm $\|[F_{best}]\|_T$, conclude Section 4. In these examples, which we call *harmonic twistors*,

¹Such maps f are intrinsically harmonic.

surfaces Σ , which deliver the Thurston norm have arbitrary big twists and heights (cf. Figure 15).

In *Section 5* we study a very special case of *self-indexing* maps $f : M \rightarrow S^1$ and surfaces Σ , which have the simplest intersection pattern with the “worst” fiber F_{worst} . For them, we establish the most desirable result: $\chi_-(\Sigma) \geq \chi_-(F_{\text{best}})$, $g(\Sigma) \geq g(F_{\text{best}})$ (cf. Theorem 5.2). Section 5 indicates the main ideas of our approach in a form which is divorced from the combinatorial complexities of the general case (presented in Section 8). After 2-surgery, the resolved surface is pushed into a neighborhood of F_{best} —a union of “good” fiber components—where it can be effectively compared with F_{best} .

Section 6 is the most tedious of them all. Here we develop the main technical tools: the notions of twist, breadth and height invariants of special surfaces $\Sigma \subset M$ in relation to a given map $f : M \rightarrow S^1$. Ultimately, the singularities of f are responsible for the non-triviality of these invariants.

In *Section 7* we aim to separate a generic embedded surface Σ , representing a given vertical homology class, from the union F_R of the “worst” fiber components. In a sense, such a separation will permit us to reduce the case of general maps f to the case treated in Section 5. It is achieved by *resolving* the intersections of a probe surface Σ with the fiber components from F_R (cf. Figure 11).

Section 8 contains the proofs of our main results—Theorems 8.2, 8.6, 8.13 and Corollaries 8.3, 8.7, 8.8, 8.14. Here we combine the strategy from Section 5 with the combinatorial tools developed in Sections 4 and 6 to derive generalizations of the inequalities (1.1) and (1.2). After proving a variety of “best fiber component theorems” generalizing Theorem 1.1, we proceed to apply these results to the χ_- -characteristic of special links (cf. Corollary 8.15).

Finally, *Section 9* deals with the way a surface Σ , homologous to a best combination F_{best} of fiber components *in the complement* to the f -singularities, can be tangent to the f -fibers. In a sense, we connect, via the twist invariants, the Morse theory of f with the induced Morse theory of $f|_{\Sigma}$ on a probe surface $\Sigma \subset M$.


To help our reader to cope with the expanding variety of notations, we conclude with a notation list.

2. INTRINSICALLY HARMONIC 1-FORMS AND THEIR GRAPHS

Next, we proceed with a description of a few facts, constructions and notations related to an intrinsic characterization of harmonic (rational) 1-forms. Actually, these facts are not specific to dimension three.

Let $\Sigma \subset M$ be an oriented surface, and let $[\omega] \in H^1(M; \mathbb{Z})$ be the class Poincaré-dual to $[\Sigma]$. It can be realized by a closed rational 1-form ω on M , or, equivalently, by a map $f : M \rightarrow S^1$ with Σ as one of its regular fibers. The two realizations are linked by the formula $f^*(d\theta) = \omega$, where $d\theta$ is the canonical 1-form on the oriented circle. For a given class $[\omega]$, one can choose its representatives f and ω with Morse-type singularities. Furthermore, if $[\omega] \neq 0$, through a deformation of f , the singularities of indices 0 and 3—the local minima and maxima of f —can be eliminated.

By considering harmonic maps $f : M \rightarrow S^1$ or, what is the same, harmonic 1-forms $\omega = f^*(d\theta)$, we exclude the singularities of indices 0 and 3—harmonic functions have no local maxima and minima. One might wonder if there are restrictions on the distribution of critical points of indices 1 and 2 imposed by the

FIGURE 1. Graph Γ_ω and the fiber components.

harmonicity and not prescribed by the topology. Fortunately, this question has a comprehensive answer.

In [FKL], for any closed 1-form ω on M , we have introduced a finite graph Γ_ω . The manifold M is canonically mapped onto the graph Γ_ω by a map p_ω . When the singular foliation \mathcal{F}_ω , determined by ω , has only compact leaves, the points of Γ_ω are just their *connected components*. In particular, each connected component of every fiber of a Morse mapping $f : M \rightarrow S^1$ corresponds to a single point in the graph $\Gamma_f = \Gamma_\omega$. Vertices of Γ_f correspond to critical points of f .

An example of Γ_f is given in Figure 1, which also depicts generic fiber components suspended over each edge of Γ_f .

The map f factors through a canonical projection $p_f : M \rightarrow \Gamma_f$, and thus generates an equally canonical map $\pi_f : \Gamma_f \rightarrow S^1$. In Figure 1, π_f is induced by the radial projection.

For Morse maps f with distinct critical values in S^1 and no local extrema, some vertices of Γ_f mark critical points x_* , such that crossing the critical value $\theta_* = f(x_*)$ causes the fiber to change the number of its connected components. Such singularities correspond to the *trivalent* vertices in Γ_ω . The rest of the singularities correspond to *bivalent* vertices of the graph.

Definition 2.1. A 1-form ω or a related map $f : M \rightarrow S^1$ is called *intrinsically harmonic*, if it is harmonic with respect to *some* metric on M .

In [C], Calabi established the following global criterion:

Theorem 2.2. A closed 1-form ω on a closed manifold M is *intrinsically harmonic* if and only if, through any point in M different from the singularities of ω , there is a loop γ along which ω is strictly positive, that is, $\omega(\dot{\gamma}) > 0$.

In [FKL], we observed that a similar property can be formulated for the graph Γ_ω . In fact, for closed manifolds, the intrinsic harmonicity of ω becomes equivalent to the following *positive loop property* of Γ_ω : through any point of Γ_ω , one can draw an oriented loop, comprised of ω -oriented edges. We call such Γ_ω 's *Calabi graphs*.

Thus, the notion of Calabi graphs provides us with a completely combinatorial description of intrinsic harmonicity.

In particular, if for some arc in S^1 , its pre-image in Γ_f is a *single* edge, then the graph automatically satisfies the positive loop property, and the map f to S^1 is intrinsically harmonic.

The graph in Figure 1 violates the positive loop property, and the corresponding map f is not intrinsically harmonic.

In [FKL], we proved Theorem 2.3 below for two extreme model cases: the case when all the leaves of the foliation \mathcal{F}_ω are compact and the case, when none of the leaves is compact. In [Ho], Honda established the general case.

Theorem 2.3. *Let ω be a closed 1-form on a closed n -manifold M with Morse-type singularities. Assume that ω has no critical points of indices 0 and n . Then one can deform ω (through the space of closed 1-forms) to an intrinsically harmonic form $\tilde{\omega}$ which has the same collection of singularities.* \square

3. THE GENERA OF FIBERS AND THE COMBINATORICS OF HANDLE MOVES

We examine some graph-theoretical descriptions of Morse maps $f : M \rightarrow S^1$ in connection with genera and χ_- -invariants of their fiber components. We analyze how f -deformations affect these combinatorial descriptions.

The genus $g(\Sigma)$ of a surface Σ is defined to be half of the rank of the homology group $H_1(\Sigma; \mathbb{Z})$.² If Σ consists of several components, its genus is the sum of the components' genera. The Euler characteristic $\chi(\sim)$ of a surface is the sum of the Euler characteristics of its components. Finally, the Thurston $\chi_-(\sim)$ -characteristic of a surface is the *absolute value* of the Euler number of the union of all its components, *excluding* the 2-spheres and, in the case of surfaces with boundary, also excluding the 2-disks.

For a closed oriented surface Σ we have $\chi_-(\Sigma) = 2|\nu(\Sigma) - g(\Sigma)|$, where $\nu(\Sigma)$ stands for the number of non-spherical connected components in Σ .

Let γ be a simple loop on an oriented surface Σ . Performing a 2-surgery on Σ along γ has the following effect on the three invariants.

2-Surgery List A:

- When the loop γ is null-homotopic, then the surgery will have the following effect:
 - (1) the genus will remain the same,
 - (2) the Euler characteristic will increase by 2, and
 - (3) χ_- will remain the same.
- When the loop γ separates Σ and is not null-homotopic, then the surgery will have the following effect:
 - (1) the genus will remain the same,
 - (2) the Euler characteristic will increase by 2, and
 - (3) χ_- will decrease by 2.
- When the loop γ does not separate Σ , then the surgery will change:
 - (1) the genus by subtracting 1,
 - (2) the Euler characteristic by adding 2, and
 - (3) χ_- will remain the same if the surface is a torus, and will decrease by 2 otherwise.

²For closed surfaces it is an integer; for surfaces with boundary, it might be a half integer.

Let γ be a simple arc on an oriented surface Σ , connecting two points on its boundary $\partial\Sigma$. Performing a relative 2-surgery on Σ along γ produces changes described in

2-Surgery List B:

- When the arc γ is null-homotopic modulo $\partial\Sigma$, then the surgery will have the following effect:
 - (1) the genus will remain the same,
 - (2) the Euler characteristic will increase by 1, and
 - (3) χ_- will remain the same.
- When the arc γ separates Σ and is not null-homotopic modulo $\partial\Sigma$, then
 - (1) the genus will remain the same,
 - (2) the Euler characteristic will increase by 1, and
 - (3) χ_- will decrease by 1.
- When the arc γ does not separate Σ , then
 - (1) the genus will drop by 1,
 - (2) the Euler characteristic will increase by 1, and
 - (3) χ_- will remain the same, if the surface is an annulus, and will decrease by 1 otherwise.

These observations can be summarized in

Lemma 3.1. *Under 2-surgery, the genus of a surface and its $\chi_-(\sim)$ -characteristic are non-increasing quantities. The Euler characteristic is strictly increasing. \square*

In the vicinity of each vertex and over a small arc of the circle centered on the critical value, Γ_f looks as depicted in the four diagrams of Figure 2. The labels 1 and 2 indicate the Morse index of the critical point. Each vertex is a *bivalent* or a *trivalent* one. The corresponding critical points also are called bivalent or trivalent. The diagrams do not include the case of critical points of indices 0 and 3. They can be depicted by an oriented edge, emanating from or terminating at a vertex of multiplicity 1.




FIGURE 2. Bivalent and trivalent singularities and the 1-chain $\tau_g(f)$.

Using f , one can produce a 1-chain $\tau_g(f)$ on the graph Γ_f : just assign to each edge the genus of the generic fiber component over it. Examining Surgery Lists A and B above, we see that the chain $\tau_g(f)$ is a *relative cycle* modulo the bivalent vertices. When $f : \partial M \rightarrow S^1$ is a fibration, crossing a bivalent vertex of index 1 from left to right results in an increase of the $\tau_g(f)$ -value by 1, and crossing a bivalent vertex of index 2 from left to right results in a decrease of the $\tau_g(f)$ -value by 1.

For a given smooth map $f : M \rightarrow S^1$ with a finite number of distinct critical values $\{\theta_i^*\}_{0 \leq i \leq k}$, we define the variation $var_{\chi_-}(f)$ as a *half* of the cyclic sum

$$(3.1) \quad |\chi_-(f^{-1}(\theta_0)) - \chi_-(f^{-1}(\theta_k))| + \sum_{i=0}^{k-1} |\chi_-(f^{-1}(\theta_{i+1})) - \chi_-(f^{-1}(\theta_i))|,$$

where a typical point θ_i belongs to the open arc $(\theta_i^*, \theta_{i+1}^*)$. We introduce the oscillation $osc_{\chi_-}(f)$:

$$(3.2) \quad osc_{\chi_-}(f) := \max_{\theta} \{\chi_-(f^{-1}(\theta))\} - \min_{\theta} \{\chi_-(f^{-1}(\theta))\}.$$

By definition, $var_{\chi_-}(f) \geq osc_{\chi_-}(f)$.

Note that $\{var_{\chi_-}(f) = 0\}$ and $\{osc_{\chi_-}(f) = 0\}$ are equivalent conditions imposed on f . Evidently, for a fibration f , $var_{\chi_-}(f) = osc_{\chi_-}(f) = 0$.

We say that a Morse map $f : M \rightarrow S^1$ is *self-indexing*, if there is a point $\theta_b \in S^1$ so that, moving from θ_b along the oriented circle, the critical values of critical points of lower indices precede the ones of higher indices. Most of the time, we assume that the f -critical values are all distinct.

Given a self-indexing map with no critical points of indices 0 and 3, one can find two points $\theta_b, \theta_w \in S^1$ such that the oriented arc (θ_b, θ_w) contains all critical values of index 1, and the complementary arc (θ_w, θ_b) contains all critical values of index 2. Let $F_{best} = f^{-1}(\theta_b)$ and $F_{worst} = f^{-1}(\theta_w)$.

For a self-indexing map $f : M^3 \rightarrow S^1$ with no local maxima and minima, the variation $var_{\chi_-}(f)$ equals $osc_{\chi_-}(f) := \chi_-(F_{worst}) - \chi_-(F_{best})$: the invariant χ_- is non-increasing under 2-surgery.

Next, we examine the effect of deforming Morse maps $f : M \rightarrow S^1$ on the invariants $\chi_-(F_{worst}), \chi_-(F_{best})$ and $var_{\chi_-}(f)$. We use graphs Γ_f , equipped with a canonical map π_f to a circle, as book-keeping devices.

We notice that the boundary $\partial\tau_g(f)$ of the 1-chain $\tau_g(f)$ is a 0-chain on Γ_f supported on the bivalent vertices. Its l_1 -norm is the variation $var_g(f)$.

In a similar way, one can introduce a 1-chain $\tau_{\chi_-}(f)$ on Γ_f by assigning to each edge the χ_- -characteristic of the corresponding fiber component. Figure 3 deals with the case when M is closed, or when $\partial M \rightarrow S^1$ is a fibration. It is divided into generic and special patterns. Special patterns arise when at least one of the fiber components is a sphere.

Definition 3.2. A critical point is called *bubbling*, if there is at least one spherical or disk fiber component in its vicinity.

Note that only at the bubbling vertices does the chain $\tau_{\chi_-}(f)$ satisfy the cycle condition. The boundary $\partial\tau_{\chi_-}(f)$ of the 1-chain $\tau_{\chi_-}(f)$ is a 0-chain on Γ_f supported on the non-bubbling vertices. Its l_1 -norm is the variation $var_{\chi_-}(f)$.

Deforming f causes Γ_f to go through a number of transformations that can be decomposed in a few basic moves. Before and after deformations, all the π_f -images of the vertices in Γ_f are assumed to be distinct in S^1 .

The five diagrams in Figure 4 depict the effect on the chain $\gamma_g(f)$ of deforming a Morse function, so that the critical value of index 1 is placed *below* the critical level of index 2 (equivalently, of sliding a 1-handle “below” a 2-handle). The diagrams are produced by combining the patterns from Figure 2 in pairs.

We notice that such an operation is always possible [M]. Unfortunately, it could only *increase* the l_1 -norm of the chains $\tau_g(f)$ and $\partial\tau_g(f)$, or the value $g(F_{best})$.




FIGURE 3. Bivalent and trivalent singularities and the 1-chain $\tau_{\chi_-}(f)$. The dotted lines mark the spherical fiber components.

From this perspective, the inverse operations (acting from the right to the left configurations) are desirable, but not always geometrically realizable! On the other hand, any two critical points of the same index can be re-ordered.




FIGURE 4. Generic patterns of moving a critical point of index 2 above a critical point of index 1 and the effect of these moves on the genera of fiber components.

Consider a portion W of M represented by the diagrams in Figure 4 and view W as a cobordism between two surfaces Σ_0 and Σ_1 represented by the left and right ends of the diagrams. Examining the five moves, we notice that, in the configurations B through E, the two critical points cannot cancel each other *locally*, that is, by a deformation which is constant on $\Sigma_0 \cup \Sigma_1$ —the trivial cobordism that would result is inconsistent with: 1) the prescribed connectivity of $\Sigma_0 \cup \Sigma_1$ (cases B and C) or with 2) the connectivity of W (case D), or with the non-triviality of

$H_1(W, \Sigma_0)$ (cases D, E). As the diagrams testify, the local cancellation of a 1-handle and a 2-handle is only possible among the vertices in diagram A.

Note that the contribution to the variation $var_g(\sim)$ in four diagrams A–D remains invariant under the moves! In fact, $var_g(\sim)$ can be changed only through the cancellation of singularities in pairs, or through the E-moves. In the first case it drops by 1, in the second it rises by 1.




FIGURE 5. Generic patterns of moving a critical point of index 2 above a critical point of index 1 and the effect of these moves on the χ_- -characteristics of the fiber components.

Figure 5 shows the effect of the same deformations on the chain $\tau_{\chi_-}(f)$. It depicts cases where the spherical and disk components are not involved. The diagrams are produced by combining in pairs the first two patterns in Figure 3.

The 14 bubbling (that is, special spherical) patterns are the result of combining in pairs the patterns in Figure 3. We leave their depiction to the reader.

Note, that the variation $var_{\chi_-}(f) = \frac{1}{2} \|\partial \tau_{\chi_-}(f)\|_{l_1}$ is preserved under all the moves in Figure 5. It can increase by 2 only under special moves which involve bubbling singularities. The variation $var_{\chi_-}(f)$ is non-decreasing under all the moves.

The diagrams in Figures 2–5 reflect our fundamental assumption that $f : \partial M \rightarrow S^1$ is a fibration.

These observations can be summarized in Lemmas 3.3–3.5 below. We assume that all the deformations described in the lemmas take place in the space $\mathcal{M}(M, S^1)$ of smooth maps $f : M \rightarrow S^1$ with *no local maxima and minima* and with generalized Morse singularities (these are the usual quadratic Morse singularities and the generic cubic singularities resulting from the merge of two Morse singularities). Let $\mathcal{M}^\circ(M, S^1) \subset \mathcal{M}(M, S^1)$ be the subspace of maps with no spherical or disk components in their fibers. Considering maps from $\mathcal{M}^\circ(M, S^1)$ eliminates the special bubbling patterns. For instance, if no non-trivial element from $H_2(M, \partial M; \mathbb{Z})$ admits a representation by a sphere or a disk, then any harmonic map $f : M \rightarrow S^1$ belongs to $\mathcal{M}^\circ(M, S^1)$.

Lemma 3.3. *The number of bivalent, index-one vertices of the graph Γ_f is equal to the variation $var_g(f)$ of the fiber genus, which, in turn, coincides with the half of the l_1 -norm $\|\partial \tau_g(f)\|_{l_1}$ of the 0-cycle $\partial \tau_g(f)$.*

The numbers of trivalent and bivalent vertices in Γ_f are deformation invariants of the Morse maps which avoid the cancellation of singularities and the E-moves. Such deformations can be decomposed into a sequence of the four basic moves A–D, depicted in Figure 4, and their inverses (when those are realizable), plus the moves which reorder the adjacent vertices of the same Morse index.³

Under an E-move, the norms $\|\partial\tau_g(f)\|_{l_1}$ and $\|\tau_g(f)\|_{l_1}$ both increase by 2. \square

Lemma 3.4. Under the cancellation of singularities of indices 1 and 2, the norms $\|\tau_g(f)\|_{l_1}$, $\|\tau_{\chi_-}(f)\|_{l_1}$ and the variations $\|\partial\tau_g(f)\|_{l_1}$, $\|\partial\tau_{\chi_-}(f)\|_{l_1}$ are decreasing. \square

Lemma 3.5. The number of non-bubbling vertices of the graph Γ_f is equal to the χ_- -variation $\text{var}_{\chi_-}(f) = \|\partial\tau_{\chi_-}(f)\|_{l_1}$.

The variation is preserved under the generic moves,⁴ as shown in Figure 5, together with their inverses. The special moves (involving bubbling singularities) can increase the variation by 2.

In particular, $\text{var}_{\chi_-}(f)$ is invariant under the deformations within the space of Morse maps with a fixed list of index 1 and 2 singularities and with no spherical fiber components.

Under the generic moves, $\|\tau_{\chi_-}(f)\|_{l_1}$ increases at least by 4. Under the special moves, it increases by 2, or is preserved. \square

4. GRAPH-THEORETICAL MANIFESTATIONS OF HARMONICITY

In this section we further examine harmonicity in terms of graph theory.

Let $\pi_f : \Gamma_f \rightarrow S^1$ be a map of graphs corresponding to a given Morse map $f : M \rightarrow S^1$. For a map f with no local maxima and minima, the vertices of Γ_f all are bivalent or trivalent. In addition, they come in two flavors: indexed by 1 or 2 depending on the Morse index of the corresponding critical point.

There are a few restrictions on the distribution of vertices of indices 1 and 2 in Γ_f . They are shown in Figure 6.

Lemma 4.1. Assume that $f : M \rightarrow S^1$ has no local extrema. If a simple positive loop $\gamma \subset \Gamma_f$ contains vertices only of a particular index, then all the f -critical points corresponding to these vertices are bubbling.

Any loop $\gamma \subset \Gamma_f$ which is not mapped by π_f into S^1 in a monotone fashion contains at least one vertex of index 1 and at least one vertex of index 2.




FIGURE 6.

³They do not affect the graph Γ_f and the chain $\tau_g(f)$.

⁴Those avoid bubbling singularities.

Proof. Follow the change in the χ_- -values or genera of fiber components along γ and use the principle “what goes up has to come down”. Figures 2, 3 and 6 will facilitate the argument. Because all vertices of Γ_f are bivalent or trivalent, each local maximum of $\pi_f : \gamma \rightarrow S^1$ corresponds to a trivalent vertex of index 1 and each local minimum corresponds to a trivalent vertex of index 2 (cf. Figure 8). These local maxima and minima alternate along γ and their cardinalities are equal. \square

We introduce two finite subsets A and R of Γ_f which will play an important role in the paper. The elements of A will be called *attractors*, and the elements of R will be called *repellers*.⁵ Each π_f -oriented edge of Γ_f with its left vertex being of index 1 and its right vertex of index 2 acquires exactly one repeller; each oriented edge with its left vertex being of index 2 and its right vertex of index 1 acquires exactly one attractor (cf. Figure 7).

We observe that changing f to $-f \bmod 2\pi$ switches the orientations of the edges in Γ_f and turns points of index 1 into points of index 2. Therefore, Γ_f and Γ_{-f} share the same sets of attractors and repellers.

It is worth noticing that the elementary moves from Figure 4 all increase the number of repellers by 1. So, it is easy to increase the size of R . To decrease it is a very different story. For instance, one might wonder: *what is the minimal number of repellers in a given 2-homology class?*




FIGURE 7.

Lemma 4.2. *Let $\pi_f : \Gamma_f \rightarrow S^1$ be such that no positive loop in Γ_f contains vertices only of a particular index and no vertex of Γ_f is univalent.⁶ Then each $r \in R$ gives rise to a pair of subtrees $T_r^+, T_r^- \subset \Gamma_f$ with the common root at r and their leaves belonging to A . The branches of T_r^+ (T_r^-) are formed by paths emanating from r in the positive (negative) direction and terminating at the first point from A they encounter. The trees $\{T_r^+, T_r^-\}_{r \in R}$ form a cover of Γ_f .*

In particular, these conclusions hold for any Γ_f produced by a harmonic map f of a manifold M with the property: no non-trivial class in $H_2(M, \partial M; \mathbb{Z})$ admits a spherical or disk representative.

Proof. Assume that the subgraph T_r^+ contains a loop. This could happen in a number of ways.

1) There exists a π_f -positive path ξ in Γ_f which leaves r and closes on itself at a vertex x without encountering an attractor. This generates a positive loop $\tau \subset \xi$ which contains x . If τ does not contain r , then x must be of index 1. By

⁵The names are inspired by the roles these combinatorial devices will play in our method of tackling the best fiber component problem.

⁶This is the case when $f : M \rightarrow S^1$ does not have extrema and bubbling singularities.

the lemma's hypotheses, τ must contain at least one vertex of index 2. Hence, an attractor must exist on τ . This contradicts assumption 1). If $r \in \tau$, ξ is a positive closed path. It contains an oriented edge $[y, z]$, where the vertex y is of index 2, the vertex z of index 1 and $r \in [y, z]$. Therefore, the loop ξ must also contain an oriented edge $[y', z']$, where the vertex y' is of index 1 and the vertex z' is of index 2. So, ξ must contain an attractor, which contradicts assumption 1).

2) The second option for T_r^+ to contain a loop arises when there are two distinct positive paths emanating from r and terminating at the same attractor a . We can assume that, for both paths, a is the first attractor after r . Since the two paths must first separate at some point x of index 2 (which succeeds r) and then join at another point y of index 1 (which precedes a), each of the paths must contain at least one attractor distinct from a and which precedes it. Thus, the two paths must terminate at these two distinct attractors *before* they reach a . This contradiction rules out loops of the second type.

Finally, we need to show that any point $x \in \Gamma_f$ belongs to some tree T_r^+ or T_r^- . Consider a positive path ξ through x which does not admit any extension. Such a path must be closed or must close on itself in both positive and negative directions. If ξ is closed, by the lemma's hypothesis, it must contain at least one vertex of index 1 and at least one vertex of index 2, unless f is a fibration. Therefore, ξ will contain at least one attractor and one repeller. Thus, moving from x along ξ in the positive or negative direction we must encounter a repeller r . Evidently, $x \in T_r^\pm$ for the first repeller.

The case when ξ through x closes on itself in both directions already has been analyzed in 1). Again, the two loops at the “ends” of ξ each must contain an attractor-repeller pair. In the worst case, at least there we will find the right repeller—a repeller r whose trees T_r^\pm contains x .

It remains to notice that harmonic maps f do not have local extrema, thus excluding univalent vertices in Γ_f . Also, no bubbling singularities can occur, because the relevant spherical fiber components of harmonic maps must generate non-trivial elements in the 2-homology of M (contrary to the hypotheses about M). \square

Lemma 4.3. *For any loop $\gamma \subset \Gamma_f$, $\int_\gamma A = \int_\gamma R$, where $\int_\gamma A$ and $\int_\gamma R$ stand for the sums of (\pm) -weighted points from A and R along the loop γ . The sign of a point $x \in \gamma$ is produced by comparing the orientation of γ with the π_f -induced orientation of the edge containing x (cf. Figure 8).*




FIGURE 8. $\int_\gamma [A] = \int_\gamma [R]$.

Proof. The numbers of local maxima and minima of the function $\pi_f|_\gamma$ along any loop γ are equal. Maxima occur at vertices of index 2 and minima occur at vertices

of index 1. The signs attached to repellers and attractors along each arc of γ between two consequent extrema are the same and alternate as one crosses from an arc to an arc. Also, since along each arc attractors and repellers alternate and since the cardinality of attractors exceeds the cardinality of repellers by 1, the attractors contribute to the integral 1 more (less) than the repellers along any ascending (descending) arc. Because the number of ascending and descending arcs are equal, the total contributions of R and A to the integral are equal as well. \square

Definition 4.4. A homology class in $H_2(M, \partial M; \mathbb{Z})$ is called *f-vertical* if it can be represented as a \mathbb{Z} -linear combination of the fundamental classes of the fiber components (which do not necessarily belong to the same f -fiber).

We denote the subgroup of f -vertical classes by H_2^f . It contains a positive cone H_2^{f+} generated by non-negative linear combinations of the fundamental classes of the fiber components.⁷

Lemma 4.5. *The elements of H_2^f are detected by their intersection numbers with loops $\{C_k \subset M\}_k$ whose images in Γ_f form a basis of $H_1(\Gamma_f; \mathbb{Z})$. In fact, H_2^f is Poincaré dual to the image of $H^1(\Gamma_f; \mathbb{Z})$ in $H^1(M; \mathbb{Z})$ induced by the canonic projection $p_f : M \rightarrow \Gamma_f$.*

Proof. By the Poincaré duality, an element of $H_2(M, \partial M; \mathbb{Z})$ is determined by its intersection numbers with loops in M . The intersection number of a loop $C \subset M$ with a vertical 2-cycle $[\Sigma]$ is equal to the intersection number of its weighted finite support σ in Γ_f with the image C' of C under the map $p_f : M \rightarrow \Gamma_f$. Therefore, $\Sigma \circ C = \sigma \circ C' = \sigma^*(C')$, where σ^* stands for the 1-cocycle in Γ_f dual to the 0-chain σ . Thus, $\Sigma \circ C$ reduces to the natural non-degenerated pairing between $H_1(\Gamma_f; \mathbb{Z})$ and $H^1(\Gamma_f; \mathbb{Z})$. Therefore, $H_2^f \approx p_f^*(H^1(\Gamma_f; \mathbb{Z}))$. \square

Let $\mathbb{Z}[A]$ be the free \mathbb{Z} -module generated by the attracting set $A = \{a\}$. Elements of $\mathbb{Z}[A]$ can be viewed as functions $\kappa : A \rightarrow \mathbb{Z}$ or, equivalently, as formal sums $\sum_{a \in A} \kappa_a \cdot a$ with integral coefficients $\{\kappa_a\}$. Combinations with non-negative coefficients generate a positive cone $\mathbb{Z}_+[A] \subset \mathbb{Z}[A]$.

Under the hypotheses of Lemma 4.2, each point in Γ_f which is not a vertex serves as a root of a subtree in Γ_f with its leaves in A . As a result, any fiber component is cobordant to a union of fiber components indexed by elements of A . Therefore,

Lemma 4.6. *Let $\pi_f : \Gamma_f \rightarrow S^1$ be such that no positive loop in Γ_f contains vertices only of a particular index. Then the obvious homomorphisms $\mathbb{Z}[A] \rightarrow H_2^f$ and $\mathbb{Z}_+[A] \rightarrow H_2^{f+}$ are onto.* \square

Theorem 4.7. *A map $f : M \rightarrow S^1$ with no local extrema and no bubbling singularities is intrinsically harmonic if and only if $\text{Ker}\{\mathbb{Z}_+[A] \rightarrow H_2(M, \partial M; \mathbb{Z})\} = 0$. That is, no positive combination of the f -oriented fiber components from A is homologous to zero in M modulo ∂M .*

In particular, if no non-trivial class of $H_2(M, \partial M; \mathbb{Z})$ has a spherical or disk representative⁸ and f is harmonic, then $\text{Ker}\{\mathbb{Z}_+[A] \rightarrow H_2(M, \partial M; \mathbb{Z})\} = 0$.

⁷The orientations of the fiber components are determined via f by the preferred orientations of M and S^1 .

⁸Equivalently, when the Thurston seminorm is a norm.

Proof. Calabi's positive loop property is equivalent to the intrinsic harmonicity of f ([C]). It implies that, for any f -oriented fiber component F , there is a positive loop C so that $F \circ C > 0$. Furthermore, for every other fiber component F' , $F' \circ C \geq 0$. Therefore, by Lemma 4.5, $\text{Ker}\{\mathbb{Z}_+[A] \rightarrow H_2(M, \partial M; \mathbb{Z})\} = 0$.

On the other hand, if the Calabi positive loop property fails for a point $x \in M$, then the upper world U_x of x — the set of points in M which can be reached from x following an f -positive path — is bounded by several fiber components (one of which contains x and the rest contain some singularities of f) [C]. Along these components the gradient of f is directed inwards to U_x . Hence, the union of these f -oriented components produces a trivial element in $H_2(M, \partial M; \mathbb{Z})$. Since, by Lemma 4.6, any f -oriented component is cobordant to a union of a few components indexed by elements of A , we have produced a nontrivial element in the kernel $\text{Ker}\{\mathbb{Z}_+[A] \rightarrow H_2(M, \partial M; \mathbb{Z})\}$.

To validate the last statement of the proposition, we notice that for harmonic maps spherical fiber components must be homologically non-trivial, which contradicts the postulated nature of M . \square

Lemma 4.8. *For an intrinsically harmonic f , $H_1(\Gamma_f; \mathbb{Z})$ admits a basis represented by π_f -positive loops.*

If no non-trivial class of $H_2(M, \partial M; \mathbb{Z})$ has a spherical or disk representative, then the Calabi graph Γ_f has no positive loops with vertices only of a particular index.

Proof. Since through any point $x \in \Gamma_f$ there exists a positive loop, the statement follows by induction on the number of edges in the complement to a maximal tree in Γ_f . The second statement follows from Lemma 4.1. \square

Now we introduce two quantities which (in Section 8) will play an important role in our arguments. The first is the difference $\chi_-(F_R) - \chi_-(F_A) = \|F_R\| - \|F_A\|$.

Lemma 4.9. *For any Γ_f as in Lemma 4.2,*

$$(4.1) \quad \text{var}_{\chi_-}(f) = \chi_-(F_R) - \chi_-(F_A).$$

Proof. We notice that $\chi_-(F_R) - \chi_-(F_A) = \sum_{r \in R} [\chi_-(F_r) - \sum_{a \in T_r^+} \chi_-(F_a)]$. Note that $[\chi_-(F_r) - \sum_{a \in T_r^+} \chi_-(F_a)]$ is twice the number of non-bubbling vertices of index 2 on the tree T_r^+ . Hence, $\chi_-(F_R) - \chi_-(F_A)$ equals the total number of non-bubbling f -singularities of index 2, while $g(F_R) - g(F_A)$ equals the total number of “bivalent” (i.e. locally non-separating) f -singularities of index 2 (cf. Lemmas 3.3–3.5).

Similar counting which employs the negative trees $\{T_r^-\}$ will reveal $\chi_-(F_R) - \chi_-(F_A)$ as the total number of non-bubbling f -singularities of index 1, and $g(F_R) - g(F_A)$ as the total number of “bivalent” f -singularities of index 1. Again using Lemmas 3.4 and 3.5, we see that formula-definition (3.1) gives still another count of non-bubbling singularities. \square

Denote by $F_{\text{best}}^{[F_A]}$ a surface $\coprod_{a \in A} \kappa_a \cdot F_a$ — an oriented union of fiber components — which represents the homology class $[F_A] := \sum_{a \in A} [F_a]$ and delivers the minimal value of $\sum_{a \in A} |\kappa_a| \cdot \chi_-(F_a)$ among such representatives.

We introduce $\text{Var}_{\chi_-}(f)$, a modification of $\text{var}_{\chi_-}(f)$, via the formula

$$(4.2) \quad \text{Var}_{\chi_-}(f) = \chi_-(F_R) - \chi_-(F_{\text{best}}^{[F_A]}) = \|F_R\| - \|[F_A]\|_{H^f}.$$

Evidently, $Var_{\chi_-}(f) \geq var_{\chi_-}(f)$. Although $var_{\chi_-}(f)$ is a more pleasing invariant (it connects in a more direct way with the f -singularities), actually, it is $Var_{\chi_-}(f)$ which will participate more often in our estimates.

Example 4.10. The harmonic twister. Let us consider the case when the graph Γ_f is an oriented loop with two vertices a and b of indices 1 and 2 respectively. The map $\pi_f : \Gamma_f \rightarrow S^1$ is the obvious one. The 1-chain $\tau_g(f)$ on Γ_f is defined to take the value n on the oriented arc (b, a) and the value $n + 1$ on the oriented arc (a, b) . It is easy to realize these data by a map $f : M^3 \rightarrow S^1$, M^3 being a closed manifold. By applying a move A from Figure 4 we send vertex b on a “round trip” and homotope f to a new map f_1 . The new map will generate a new chain $\tau_g(f_1)$, taking the value $n + 1$ on (b, a) and the value $n + 2$ on (a, b) . This deformation can be repeated again and again to produce maps with arbitrarily large genera $n + k$ of the best fiber. Of course, the original best fiber of genus n is still residing in M^3 ; however, it is *invisible* on the level of the new graphs Γ_{f_k} (in this case, identical with the original one) and new chains $\tau_g(f_k)$.

A similar argument applies to the χ_- -invariants and the chain $\tau_{\chi_-}(f)$.

As Corollary 8.9 and Example 8.12 imply, the original f -fiber F_0 will intersect the new f_k -fiber along a complex pattern of loops, none of which bounds a disk in the new fiber (cf. Figure 15). Furthermore, these intersections cannot be removed even by an isotopy of F_0 .

Since all the graphs Γ_{f_k} satisfy the Calabi positive loop property, it follows that all the maps f_k are intrinsically harmonic [FKL], and all the f_k -fibers are near-minimal surfaces [K]. This means that, for any choice of two disjoint 3-disks D_1 and D_2 centered on the two singularities, and any $\epsilon > 0$, there exists a riemannian metric on M^3 with the following properties: 1) the map f_k is harmonic; 2) the f_k -fibers are minimal surfaces outside of the disks; and 3) the area of the portion of each fiber inside the disks is smaller than ϵ (in other words, by the choice of metric, the deviation of fibers from minimality can be localized around the singularities and made numerically insignificant).

In contrast with the fibrations over a circle, as the twister example demonstrates, this “near-tautness” of the singular foliation \mathcal{F}_{f_k} does not imply the minimality of the χ_- -characteristic of the best fiber in its homology class. \square

A challenging problem is how to “untwist” a given map $f : M \rightarrow S^1$ and to lower the l_1 -norms of the characteristic chains $\tau_{\chi_-}(f)$, $\partial\tau_{\chi_-}(f)$, or the value $\chi_-(F_{best})$. The twist invariant $\rho_{\chi_-}([\Sigma], F_R)$ from Section 6 does measure the “twist” of f (relative to $[\Sigma]$). Regrettably, I do not know how to produce maps with the zero twist out of a variational principle.

Deforming a generic map $f : M \rightarrow S^1$ to a harmonic map, while preserving the list of its singularities, as it is done in the proof of Theorem 1, pages 474-475, in [FKL], requires elementary moves D and E from Figure 5. Unfortunately, they have the potential to increase $\chi_-(F_{best})$ and the variation. At the same time, some form of harmonicity seems to be an essential ingredient in our arguments, especially if one expects a *fiber* to deliver the Thurston norm (cf. Figure 1 depicting a non-harmonic map whose fibers fail to deliver the Thurston norm). This tension between our desire to lower the value $\chi_-(F_{best})$ and the need to harmonize the map f calls for an investigation beyond the scope of this paper.

5. SELF-INDEXING MAPS TO S^1 AND THE χ_- -MINIMIZING 2-CYCLES

To avoid combinatorial complications and to make the future arguments more transparent, first we treat the case of *self-indexing* maps $f : M \rightarrow S^1$ with no local extrema. Automatically, such maps are intrinsically harmonic. For a self-indexing f , unless it is a fibration, Γ_f is a union of two trees which share the same root and the same set of leaves. More general maps are considered in Section 8.

Let F and Σ be oriented embedded surfaces which intersect transversally in M . When $\partial M = \emptyset$, we assume the surfaces are closed; otherwise, their boundaries are contained in ∂M . The intersection $\mathcal{C} = F \cap \Sigma$ consists of a number of oriented loops and arcs. Their orientations are induced by the orientations of F , Σ and M . As we modify the intersection, we still call it \mathcal{C} .

Definition 5.1. The oriented intersection pattern $\mathcal{C} = F \cap \Sigma$ is *totally reducible*, if it is comprised of curves which bound disks in F or of arcs which bound relative disks in $(F, \partial F)$.

Theorem 5.2. Let $f : M \rightarrow S^1$ be a map from an oriented 3-manifold M to an oriented circle. Assume that:

- f has no critical points of indices 0 and 3;
- if $\partial M \neq \emptyset$, then $f : \partial M \rightarrow S^1$ is a fibration;
- along the circle, the critical values of critical points of index 1 belong to an oriented arc (θ_b, θ_w) , and the critical values of points of index 2 belong to a complementary oriented arc (θ_w, θ_b) .⁹

Let $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ be an embedded oriented surface, homologous rel. ∂M to a fiber. Assume that Σ has a totally reducible intersection \mathcal{C} with the fiber $F_{\text{worst}} = f^{-1}(\theta_w)$.

Then $\chi_-(\Sigma) \geq \chi_-(F_{\text{best}})$ and $g(\Sigma) \geq g(F_{\text{best}})$, where $F_{\text{best}} = f^{-1}(\theta_b)$.

Proof. Let $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ be an oriented regularly embedded surface, homologous to a fiber modulo ∂M , and such that $\mathcal{C} = \Sigma \cap F_{\text{worst}}$ is totally reducible. If a loop γ from \mathcal{C} bounds a disk $D \subset F_{\text{worst}}$ which is free of any other intersection curves, then we can perform a 2-surgery on Σ along γ with D as the core of a 2-handle. The resulting surface Σ' is homologous to Σ and has $\mathcal{C} \setminus \gamma$ as its intersection set with F_{worst} . According to Surgery List A, $g(\Sigma') \leq g(\Sigma)$ and $\chi_-(\Sigma') \leq \chi_-(\Sigma)$. A similar argument, based on Surgery List B, leads to a similar conclusion for a surface resulting from surgery along an arc γ bounding a disk modulo $\partial F_{\text{worst}}$. When γ bounds in F a disk which contains other curves from \mathcal{C} , we perform 2-surgery on Σ starting with the “most interior” disks and gradually “moving outwards”.

Because $\Sigma \cap F_{\text{worst}}$ is totally reducible, we can produce a new surface Σ' in the homology class of the original Σ having an *empty* intersection with F_{worst} . Moreover, $g(\Sigma') \leq g(\Sigma)$ and $\chi_-(\Sigma') \leq \chi_-(\Sigma)$. So, if we can prove the desired inequalities for Σ' , then they will be valid for the original Σ as well.

Now we revert to our generic notation Σ for this new surface Σ' . Starting at θ_b and moving along the oriented circle, first we cross the critical values of *all* index-one critical points. Similarly, starting at θ_w and moving along the oriented circle, we first meet the critical values of *all* index-two critical points.

We cut the manifold M open along $F_{\text{best}} = f^{-1}(\theta_b)$. The boundary of the resulting manifold \hat{M} consists of two copies of F_{best} —the surfaces F_{best}^0 and F_{best}^1 .

⁹In other words, f is a *self-indexing* map.

Denote by \tilde{M} the cyclic covering space of M induced by the map f . We can regard the \hat{M} as a fundamental domain of the cyclic action on \tilde{M} . The map f generates a Morse function $\hat{f} : \hat{M} \rightarrow [0, 1]$, so that $\hat{f}^{-1}(0) = F_{best}^0$ and $\hat{f}^{-1}(1) = F_{best}^1$. The original f induces a \mathbb{Z} -equivariant Morse function $\tilde{f} : \tilde{M} \rightarrow \mathbb{R}$ with no singularities along $\partial\tilde{M}$. Here $\partial\tilde{M}$ denotes the preimage of ∂M under the covering map $\tilde{M} \rightarrow M$. We can view \hat{f} as a restriction of \tilde{f} to the fundamental domain.

By the construction, all \hat{f} -critical points of index 1 lie below the surface $F_{worst} = f^{-1}(\theta_w)$, and all critical points of index 2 lie above it.

We consider an f -gradient-like vector field X on M and its lifting \tilde{X} on the covering space \tilde{M} . When necessary, these fields will be adjusted. Since $f : \partial M \rightarrow S^1$ is a fibration, we always can assume that X is *tangent* to the boundary ∂M and does not vanish there.

Let the index α enumerate the critical points $\{x_\alpha\}$ of index 1, and the index β the critical points $\{x_\beta\}$ of index 2 in M . Denote by $\{\hat{x}_\alpha\}$ and $\{\hat{x}_\beta\}$ the corresponding \hat{f} -critical points in $\hat{M} \subset \tilde{M}$.

We denote by \tilde{D}_α^1 the two descending trajectories of the field \tilde{X} which emanate from the singularity $\{\hat{x}_\alpha\}$, and by \tilde{D}_α^2 the union of all ascending trajectories. Similarly, we denote by \tilde{D}_β^1 the two ascending trajectories, and by \tilde{D}_β^2 the union of all descending trajectories which emanate from the singularity $\{\hat{x}_\beta\}$.

Let $\hat{D}_\alpha^1 = \tilde{D}_\alpha^1 \cap \hat{M}$, $\hat{D}_\alpha^2 = \tilde{D}_\alpha^2 \cap \hat{M}$, $\hat{D}_\beta^1 = \tilde{D}_\beta^1 \cap \hat{M}$, $\hat{D}_\beta^2 = \tilde{D}_\beta^2 \cap \hat{M}$.

Since \tilde{X} is tangent to the boundary $\partial\tilde{M}$ and does not vanish there, all the sets $\{\tilde{D}_\alpha^i\}$ and $\{\tilde{D}_\beta^i\}$ ($i = 1, 2$) have an *empty intersection* with $\partial\tilde{M}$.

Note that the portions of the sets \tilde{D}_β^2 lying *above* the surface F_{worst} , and of the sets \hat{D}_α^2 lying *below* the surface F_{worst} , are diffeomorphic to two-dimensional disks.

Denote by $\hat{\Sigma}$ the preimage of Σ under the natural map $p : \hat{M} \rightarrow M$. Clearly, $\hat{\Sigma}$ has an empty intersection with the surface $F_{worst} \subset \hat{M}$ and, therefore, $\hat{\Sigma}$ is divided into two disjoint pieces: $\hat{\Sigma}^0$ lying below F_{worst} and $\hat{\Sigma}^1$ lying above it. We aim to push $\hat{\Sigma}^0$ towards F_{best}^0 and *below* any critical value produced by the singularities of index 1. At the same time, we will try to push $\hat{\Sigma}^1$ towards F_{best}^1 and *above* any critical value, produced by the singularities of index 2. In general, both desired isotopies are obstructed by the critical points $\{\hat{x}_\alpha\}$ and $\{\hat{x}_\beta\}$; however, after 2-surgery on $\hat{\Sigma}$, as Figure 9 indicates, the two deformations will become possible. (Note that, Figure 9 depicts the simplest case, when the intersection of \hat{D}_β^2 with $\hat{\Sigma}$ consists of a single loop.)

Consider the intersection of each set \tilde{D}_β^2 with the surface $\hat{\Sigma}^1$. Above the level of F_{worst} , the set \tilde{D}_β^2 is a smooth disk. Since $\hat{\Sigma}^1$ lies above F_{worst} , by a choice of the gradient vector field X (equivalently, by a small isotopy of Σ) the intersection of $\hat{\Sigma}^1$ with the disk can be made transversal. It will consist of a number of closed simple curves $S_{\beta,k}^1$ shared by the disk and the surface. The curves are closed because the disks \tilde{D}_β^2 cannot reach the boundary $\partial\tilde{M}$ —the gradient field has been chosen to be tangent to the boundary and has no zeros there. We will use these loops to perform 2-surgery on $\hat{\Sigma}^1$ inside \hat{M} . We start with the most “inner” loop in the disk, say, with $S_{\beta,1}^1$. Inside of the set \tilde{D}_β^2 it bounds a disk $D_{\beta,1}^2$, which is embedded in the ambient \hat{M} . The normal frame to $S_{\beta,1}^1$ in $\hat{\Sigma}^1$ extends to a normal frame of $D_{\beta,1}^2$ in \hat{M} . We perform a surgery on $\hat{\Sigma}^1$ along $S_{\beta,1}^1$ by attaching a 2-handle




FIGURE 9. Moving Σ above the critical level of index 2 after a 2-surgery.

$D_{\beta,1}^2 \times [0, 1]$. The surgery eliminates the intersection $S_{\beta,1}^1$ and does not produce new intersections with the “disk” \tilde{D}_β^2 . This procedure is repeated until all the intersection loops $\{S_{\beta,k}^1\}_{\beta,k}$ are eliminated. The resulting surface $\hat{\Sigma}_\star^1$ has an empty intersection with all the sets $\{\tilde{D}_\beta^2\}$, and therefore can be isotoped along the gradient flow to a location *above* all critical points $\{\hat{x}_\beta\}$. In the process, the boundary $\partial\hat{\Sigma}_\star^1$ glides along the boundary $\partial\hat{M}$. The isotopy can be chosen to be fixed in a small neighborhood of F_{best}^1 .

A similar treatment can be applied to the surface $\hat{\Sigma}^0$ and the portions of $\{\tilde{D}_\alpha^2\}$ lying below the surface F_{worst} . It will produce a new surface $\hat{\Sigma}_\star^0$ located below all the critical points $\{\hat{x}_\alpha\}$, in a neighborhood of F_{best}^0 . Because the surgery did not touch a portion of the original $\hat{\Sigma}^0$ in a neighborhood of F_{best}^0 and a portion of the original $\hat{\Sigma}^1$ in a neighborhood of F_{best}^1 , the new $\hat{\Sigma}_\star^0$ and $\hat{\Sigma}_\star^1$ still define a surface Σ_\star in M , located in a regular tubular neighborhood of the fiber F_{best} .

Since Σ and Σ_\star are linked by 2-surgery inside M followed by a regular isotopy, they define the same class in $H_2(M; \partial M; \mathbb{Z})$.

On the other hand, by Lemma 3.1, their genera and χ_- -characteristics satisfy the inequalities $g(\Sigma) \geq g(\Sigma_\star)$ and $\chi_-(\Sigma) \geq \chi_-(\Sigma_\star)$.

Applying Lemma 5.3 below, with $\Sigma = \Sigma_\star$, $F = F_{best}$ and $d = 1$, we complete the proof of Theorem 5.1. \square

Lemma 5.3. *Let U be a regular neighborhood of a connected oriented surface F , and V a regular neighborhood of its boundary $\partial F \subset \partial U$. Let $(\Sigma, \partial\Sigma) \subset (U, V)$ be an embedded oriented surface. Then $\chi_-(\Sigma) \geq |d| \cdot \chi_-(F)$ and $g(\Sigma) \geq |d| \cdot g(F)$, where d is a total degree of the map $\Sigma \rightarrow F$ induced by the retraction $p : (U, V) \rightarrow (F, \partial F)$.*

Remark. The assumption that Σ is an embedded surface is important: it is easy to construct examples of immersed surfaces which violate the conclusion of the lemma. Take, for instance, a double cover of a torus T by another torus T_1 , immersed in $T \times [0, 1]$ with the projection $T \times [0, 1] \rightarrow T$ inducing the covering degree-two map $T_1 \rightarrow T$. Here $g(T_1) = g(T)$ and not twice $g(T)$.

Proof. Let $\Sigma = \coprod_j \Sigma_j$, where each Σ_j is connected and let $p_j : \Sigma_j \rightarrow F$ be a map of degree d_j induced by the retraction $U \rightarrow F$. Recall that if $d_j \neq 0$, then

the homomorphism $(p_j)_* : H_*(\Sigma_j; \mathbb{R}) \rightarrow H_*(F_0; \mathbb{R})$ is an epimorphism [W]. In particular, $(p_j)_* : H_1(\Sigma_j; \mathbb{R}) \rightarrow H_1(F_0; \mathbb{R})$ is an epimorphism. Hence, for such a component, $g(\Sigma_j) \geq g(F)$ and $|\chi(\Sigma_j)| \geq |\chi(F)|$.

Put $d = \sum_j d_j$. We claim that Σ must divide $U \approx F \times [0, 1]$ into at least $|d| + 1$ connected regions $\{U_i\}$. If the complement to Σ in U consists of less than $|d| + 1$ components, then it is possible to construct a path $\gamma \subset U$, which connects a point $a \in F \times \{0\}$ with a point $b \in F \times \{1\}$ and which has less than $|d|$ intersections with Σ . Existence of such a γ contradicts the fact that $[\Sigma] = \sum_j [\Sigma_j]$ and $d[F]$ are homologous in (U, V) .

Take the region U_1 adjacent to $F \times \{0\}$, and let $V_1 := U_1 \cap V$. The projection $p : U_1 \rightarrow F \times \{0\}$ maps at least one component of the surface $\partial U_1 \setminus V_1$, distinct from $F \times \{0\}$, by a degree 1 map. Indeed, consider the intersections of a generic segment $I = x \times [0, 1]$, $x \in F$, with $\partial U_1 \setminus V_1$. Among them pick the highest intersection, say a . Let S_1 be the component of Σ which contains a . At a the path I leaves the domain U_1 “forever”. Take a point $b \in I$ just below a and connect b by a path $\gamma \subset U_1$ with the base of I . One can construct γ in such a way that its p -projection is a loop homologous to zero in F : just add an appropriate kick in F to any candidate for γ . Denote by J the portion of I above b . The new path K —the union of J with γ —intersects S_1 at a single point a and shares its beginning and end with I . Furthermore, the loop $K \cup I$ is null-homologous in $F \times [0, 1]$. Since the algebraic intersection number of K with S_1 is 1, the intersection number of I with S_1 must also be 1; that is, $\deg(p|_{S_1}) = 1$.

Next, consider a region U_2 adjacent to U_1 along S_1 . The same reasoning now applies to $U_1 \cup U_2$. We conclude that it must have a boundary component $S_2 \neq S_1$ which projects into $F \times \{0\}$ by the map p of degree 1. This inductive process provides us with at least $|d|$ components $\{\Sigma_j = S_j\}$ of Σ , each of which maps onto $F \times \{0\}$ by a degree 1 map. Thus, $g(\Sigma) \geq |d| \cdot g(F)$. By the same token, each of the $|d|$ components Σ_j has the property $|\chi(\Sigma_j)| \geq |\chi(F)|$. Furthermore, $\chi_-(\Sigma_j) \geq \chi_-(F)$ —a sphere cannot be mapped by a non-zero degree map onto an orientable surface different from a sphere. Hence, $\chi_-(\Sigma) \geq |d| \cdot \chi_-(F)$. \square

The following statement is very much in line with the harmonic twister example: it shows that harmonicity *alone* is too weak and too flexible to insure the “best fiber theorem”. For instance, the hypothesis in Theorem 5.2, requiring the probing surface Σ to have a totally reducible intersection with the worst fiber, is essential.

Lemma 5.4. *Let M be an oriented closed 3-manifold. Any connected, orientable, embedded and non-separating surface $\Sigma \subset M$ can be viewed as the best fiber of a self-indexing map $f : M \rightarrow S^1$ with all the fibers being connected. Such an f is intrinsically harmonic.*

Proof. Apply the Thom-Pontryagin construction to a regular neighborhood U of $\Sigma \subset M$. Perturb f away from U to convert it into a Morse map. Then cut M open along Σ to get a Morse function $\hat{f} : \hat{M} \rightarrow [0, 1]$ which maps one copy of Σ , $\Sigma_0 \subset \partial \hat{M}$, to 0, and the other copy, $\Sigma_1 \subset \partial \hat{M}$, to 1. Through a standard deformation of \hat{f} , fixed on $\partial \hat{M}$, one can eliminate all local maxima and minima. This leaves us only with critical points of indices 1 and 2. Then we can deform \hat{f} into a self-indexing Morse function \hat{f}' without changing it at $\partial \hat{M}$ ([M]). Since Σ is connected, the self-indexing \hat{f}' must have only connected fibers (cf. Figure 2), and therefore it is intrinsically harmonic. By Lemma 3.1, Σ is the best fiber of f' . \square

6. THE TWIST OF A MAP $f : M \rightarrow S^1$ AND THE \tilde{f} -BREADTH OF SURFACES

In this section we introduce a number of invariants characterizing the intersection complexity of two (hyper)surfaces $\Sigma, F \subset M$ which are specially positioned with respect to a given map $f : M \rightarrow S^1$. In Section 8 these invariants will contribute to our estimates of the Thurston norm. We also introduce the “twist” of f in terms of the intersection complexity of a surface Σ , which delivers the Thurston norm, with a generic fiber component. Ultimately, it is the presence of the f -singularities which is responsible for the non-triviality of these invariants.

Let $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ be an oriented surface representing an f -vertical class $[\Sigma]$ (cf. Definition 4.4). Let F be a finite union of fiber components. For such a pair (Σ, F) , we introduce a non-negative integer $\rho(\Sigma, F)$. It will measure the complexity of the transversal intersection $\mathcal{C} = \Sigma \cap F$ inside F . The fact that M is 3-dimensional is not important here: similar invariants make sense for any map $f : M \rightarrow S^1$ and any pair of vertical hypersurfaces in M .

The pattern $\mathcal{C} = \bigcup C_i \subset F$ is comprised of oriented simple curves (arcs and loops). Each curve C_i in F is equipped with a normal framing induced by the preferred normal framing of Σ .

It is crucial to notice that the algebraic intersection number of any loop $\gamma \subset F$ with \mathcal{C} is zero. Indeed, the algebraic intersection $\gamma \circ \mathcal{C} = \gamma \circ \Sigma$ of such a γ with any surface Σ representing a vertical class $[\Sigma] \in H_2^f$ is zero: just consider Σ' homologous to Σ and comprised of fiber components distinct from those of F to conclude that $\gamma \cap \Sigma' = \emptyset$.

We consider an oriented graph $K_{\mathcal{C}}$ whose vertices correspond to connected components of $F \setminus \mathcal{C}$ and whose edges correspond to the connected components of \mathcal{C} . The orientation of the edges is prescribed by the preferred normal frames to the intersection curves. Because any loop in F has a trivial algebraic intersection with \mathcal{C} , each loop in $K_{\mathcal{C}}$ will have an equal number of “clockwise” and “counter-clockwise” oriented edges.

Consider a 1-cochain c on $K_{\mathcal{C}}$ which takes value 1 at each oriented edge of the graph. Since, by the argument above, c takes the zero value on every loop in $K_{\mathcal{C}}$, it is a coboundary: $c = \delta u$ for some 0-cochain u . The potential u is a function on the vertices $C_0(K_{\mathcal{C}})$ of $K_{\mathcal{C}}$ which prescribes the flow c through the edges. For each connected component of F , $u : C_0(K_{\mathcal{C}}) \rightarrow \mathbb{Z}$ is well-defined up to a choice of a constant. We can synchronize these choices by equating all the maximum values of u on different connected components of $K_{\mathcal{C}}$.

Denote by $\rho(\Sigma, F)$, or by $\rho(\mathcal{C})$ for short, *one less* the number of *distinct values* taken by the synchronized function u . This integer $\rho(\mathcal{C})$ will be our measure of complexity of the intersection \mathcal{C} .

For an oriented graph with all vertices being sources and sinks, u takes only two values and $\rho(\mathcal{C}) = 1$. In general, $\rho(\mathcal{C})$ does not exceed the number of connected components in $F \setminus (F \cap \Sigma)$ minus one.

One can think of the potential u as an integral-valued function on $F \setminus \mathcal{C}$, constant on its components. In this interpretation, the curves from \mathcal{C} can be imagined as dams erected on F , and the u -values as the water levels for each of the fields from $F \setminus \mathcal{C}$. Depending on the orientations, crossing a dam results in a change of the water level by ± 1 . In this model, $\rho(\mathcal{C})$ is the integral variation of the water level across the irrigation system \mathcal{C} . In other words, for each connected component F_{α} of F , we consider an integral 2-chain E_{α} whose boundary is the 1-cycle $\Sigma \cap F_{\alpha}$, and

define $\rho(\mathcal{C})$ as $\max_{\alpha}\{osc(E_{\alpha})\}$, where $osc(E_{\alpha})$ denotes the oscillation in the values of the coefficients in the chain E_{α} .

Assume that two curves C_1 and C_2 from \mathcal{C} can be linked by an oriented arc γ which has a single transversal intersection x with C_1 and a single transversal intersection y with C_2 , the two intersections being of *opposite* signs. Also, assume that γ misses the rest of the curves from \mathcal{C} . Then the water level u_{\star} along γ before it hits C_1 and after it hits C_2 must be equal. So, we can connect the corresponding fields by a canal following γ and fill it with water up to the level u_{\star} . This irrigation construction will merge the two fields into a single one, and replaces C_1 and C_2 with their connected sum $C_1 \# C_2$, but it will not change the value of $\rho(\sim)$.

As we isotope the surface Σ in M , its transversal intersections \mathcal{C} with F are subjected to an isotopy in F and occasional surgery of the types $C_1 \sqcup C_2 \Rightarrow C_1 \# C_2$, $C_1 \# C_2 \Rightarrow C_1 \sqcup C_2$, or of the birth-annihilation types $C \Rightarrow \emptyset$, $\emptyset \Rightarrow C$. Here the loop C bounds a disk in F (or in $(F, \partial F)$) and in Σ . Also, a different type of surgery can occur: it corresponds to connecting two points x and y on the *same* curve C by an oriented arc. It has an effect of separating C into two components C_1 and C_2 . Under the transformations $C_1 \sqcup C_2 \Rightarrow C$, $C \Rightarrow C_1 \sqcup C_2$, the value of $\rho(\sim)$ is preserved. Only the birth-annihilation surgery can change it.

A modified definition of $\rho(\Sigma, F)$ will be useful. In the modification, from the very beginning, we *exclude* all loops from the intersection $\Sigma \cap F$ which *bound a disk* in F . We also exclude arcs of $\Sigma \cap F$ which bound a disk in $(F, \partial F)$. This gives us a simpler intersection pattern \mathcal{C}° . Then we employ \mathcal{C}° to define $\rho^{\circ}(\Sigma, F)$ as $\rho(\mathcal{C}^{\circ})$. This quantity can also change under the surgery of the type $C_1 \sqcup C_2 \Rightarrow C_1 \# C_2$.

The definition below introduces new invariants which depend only on the homology class $[\Sigma] \in H_2^f$, a value $\chi_{-}(\Sigma)$, and a surface $F \subset M$ which is a union of fiber components (alternatively, whose fundamental class $[F]$ is proportional to $[\Sigma]$).

Definition 6.1. Let $\rho_{\chi_{-}}(\Sigma, F)$ be the *minimum* of $\{\rho(\Sigma', F)\}_{\Sigma'}$, where $\Sigma' \subset M$ is *homologous* to Σ and $\chi_{-}(\Sigma') \leq \chi_{-}(\Sigma)$.

Lemma 6.2. $\rho_{\chi_{-}}(\Sigma, F)$ is the *minimum* of $\{\rho^{\circ}(\Sigma', F)\}_{\Sigma'}$, where $\Sigma' \subset M$ is homologous to Σ and $\chi_{-}(\Sigma') \leq \chi_{-}(\Sigma)$.

Proof. By performing 2-surgery on any given Σ along curves from $\Sigma \cap F$ which bound disks in F , we can replace Σ with a new surface Σ' such that $[\Sigma'] = [\Sigma]$, $\chi_{-}(\Sigma') \leq \chi_{-}(\Sigma)$, and $\rho^{\circ}(\Sigma', F) = \rho(\Sigma', F)$. \square

Definition 6.3. We fix a vertical homology class $[\Sigma] \in H_2^f$ and a repeller set $R \subset \Gamma_f$. Consider all oriented surfaces $\{\Sigma \subset M\}$ which deliver the minimal value of $\chi_{-}(\sim)$ in the homology class $[\Sigma]$. Among them pick Σ 's with the minimal value of the twist $\rho(\Sigma, F_R)$.¹⁰ We denote this optimal value by $\rho_{\chi_{-}}([\Sigma], F_R)$ and call it the *R-twist* of f relative to the class $[\Sigma]$.

Since any two sets of repelling components are isotopic, $\rho_{\chi_{-}}([\Sigma], F_R)$ does not depend on a particular choice of $R \subset \Gamma_f$.

For technical reasons, Definition 6.3 employs a special union F_R of fiber components. Replacing the F_R in Definition 6.3 with “any fiber component F ”, one can introduce a modified definition which makes sense for any Morse map f .

When $[\Sigma] = [F]$ is the homology class of a fiber, we also will use the abbreviation $\rho_{\chi_{-}}(f)$ for $\rho_{\chi_{-}}([F], F_R)$.

¹⁰Due to Lemma 6.2, this is equivalent to minimizing $\rho^{\circ}(\Sigma, F_R)$.

We shall see that the S^1 -controlled size of a homotopy which links a given map f to the one whose best fiber delivers the Thurston norm gives an upper bound for $\rho_{\chi_-}(f)$.

Example 6.4 (Jerome Levine). Let $\tilde{M} \rightarrow M$ be a cyclic cover induced by $f : M \rightarrow S^1$. It is not true that for every f -vertical homology class $[\Sigma]$, each surface Σ which realizes $[\Sigma]$ admits a lift to \tilde{M} . Take, for instance, a connected sum $M = (F_0 \times S^1) \# (F_1 \times S^1)$, where F_0 and F_1 are oriented surfaces of your choice. Let $f : M \rightarrow S^1$ be a connected sum of obvious projections. Consider $\Sigma = F_0 \# (\partial D^2 \times S^1)$, where $D^2 \subset F_1$ is a 2-disk. Note that $[\Sigma] = [F_0]$ (clearly, a vertical class), but Σ cannot be lifted to \tilde{M} : it contains a loop γ which is mapped by the degree 1 map $f|$ onto S^1 . We notice that $\gamma \circ \Sigma = 0$, which is in agreement with the lemma below. The immersed surface $\Sigma \cup F_1$, realizing the homology class $[F_0] + [F_1]$ (which satisfies the lemma's hypotheses) demonstrates that Lemma 6.5 cannot be generalized for immersed surfaces.

Lemma 6.5. *If $[\Sigma]$ is proportional to the homology class $[F]$ of an f -fiber, then any surface Σ realizing $[\Sigma]$ admits a lifting to the space \tilde{M} . When Σ is connected, the lift $\hat{\alpha} : \Sigma \subset \tilde{M}$ is unique up to the deck transformations in \tilde{M} .*

Proof. By Poincaré duality, any 2-homology class is characterized by its intersections with loops forming a basis in $H_1(M; \mathbb{Z})$. In particular, the homology class $[F]$ of a fiber is characterized by the property $[F] \circ [\gamma] = f_*([\gamma])$, where $f_*([\gamma])$ stands for the image of $[\gamma]$ — an integer — under $f_* : H_1(M; \mathbb{Z}) \rightarrow H_1(S^1; \mathbb{Z}) \approx \mathbb{Z}$. Thus, any class $k[F]$ is determined by the property $k[F] \circ [\gamma] = k \cdot f_*([\gamma])$. By linearity, this is equivalent to the proposition: $[\Sigma]$ is proportional to $[F]$ iff $[\Sigma] \circ [\gamma] = 0$ for any $[\gamma] \in \text{Ker}(f_*)$. Evidently, any class $[\Sigma]$ which satisfies this criterion has the property: $\{f_*([\gamma]) \neq 0\} \Rightarrow \{[\Sigma] \circ [\gamma] \neq 0\}$.

Now, in order to prove the lemma, it suffices to show that the image of any loop $\gamma \subset \Sigma$ under f is null-homotopic in S^1 . Since Σ is two-sided, $[\gamma] \circ [\Sigma] = 0$. On the other hand, if the loop $f([\gamma]) \neq 0$, then by the argument above, $[\gamma] \circ [\Sigma] \neq 0$. This contradiction completes the proof. \square

Put $\hat{\Sigma} := \hat{\alpha}(\Sigma)$. If Σ consists of many components $\{\Sigma_j\}$, each of them admits its own lift $\hat{\alpha}_j : \Sigma_j \subset \tilde{M}$. However, not every combination $\hat{\Sigma}$ of $\{\hat{\Sigma}_j \subset \tilde{M}\}_j$ will serve our goals. We will be especially interested in liftings $\hat{\alpha}$ such that the (relative) 2-cycle $\hat{\Sigma}$ is the *boundary* of an integral 3-chain C in \tilde{M} modulo $\partial\tilde{M} \cup \tilde{M}^{+\infty} \cup \tilde{M}^{-\infty}$. Here $\tilde{M}^{+\infty}$ ($\tilde{M}^{-\infty}$) stands for the positive (negative) ends of \tilde{M} .¹¹ We denote the set of such special liftings by $\mathcal{B}(\Sigma)$.

The surface $\hat{\Sigma}$ divides \tilde{M} into a finite number of connected domains $\{U_l\}_l$. When $\hat{\Sigma} \in \mathcal{B}(\Sigma)$, the (infinite) 3-chain C can be chosen so that it attaches the same integral multiplicity u_l to every 3-simplex from the domain U_l . This type of observation is already familiar from the beginning of Section 6, where it was employed (in dimension 2) to introduce the invariant $\rho(\Sigma, F)$. Indeed, if $\hat{\Sigma} = \partial C$, then $\hat{\Sigma} \circ \gamma = 0$ for any loop $\gamma \subset \tilde{M}$. This allows us to define u_j as $\hat{\Sigma} \circ \beta$, where β is a positively oriented path which connects the appropriate negative end of \tilde{M} with a generic point $x \in U_j$. It follows that no component of $\hat{\Sigma}$ has a domain U_j on both sides.

¹¹When f is primitive, \tilde{M} is connected and has a single positive and a single negative end.

For $\hat{\Sigma} \in \mathcal{B}(\Sigma)$, let

$$(6.1) \quad \rho(\hat{\Sigma}, \tilde{M}) = \text{osc}\{u_l\} := \max_l \{u_l\} - \min_l \{u_l\}$$

We introduce a subset $\mathcal{B}_k(\Sigma) \subset \mathcal{B}(\Sigma)$ based on liftings $\hat{\Sigma}$, subject to $\rho(\hat{\Sigma}, \tilde{M}) = k$.

For a lifting $\hat{\Sigma}$ which separates those components of \tilde{M} where it resides, $\rho(\hat{\Sigma}, \tilde{M}) = 1$. Furthermore, we have the following lemma, which improves on Lemma 6.5.

Lemma 6.6. *Let $\Sigma \subset M$ realize $k[F]$, the k -multiple of the homology class of a fiber F . Then there exists a special lifting $\hat{\Sigma} \in \mathcal{B}(\Sigma)$ so that $\rho(\hat{\Sigma}, \tilde{M}) = k$.*

Conversely, if a surface $\Sigma \subset M$ admits a lifting $\hat{\Sigma} \in \mathcal{B}(\Sigma)$, then $[\Sigma]$ is proportional to $[F]$ with the coefficient $\rho(\hat{\Sigma}, \tilde{M})$.

Proof. We start with the case $k = 1$. Applying the Thom-Pontryagin construction to any oriented surface $\Sigma \subset M$ produces a map $f_\Sigma : M \rightarrow S^1$. Here f_Σ stands for an approximation of the Thom-Pontryagin map P_Σ by a Morse function which coincides with P_Σ in the vicinity of Σ . Denote by $\tilde{M}_\Sigma \rightarrow M$ the cyclic cover induced by f_Σ , and by $\tilde{f}_\Sigma : \tilde{M}_\Sigma \rightarrow \mathbb{R}$ an appropriate lift of f_Σ . We denote by $\hat{\Sigma}$ the surface of constant level $\tilde{f}_\Sigma^{-1}(f_\Sigma(\Sigma)) \subset \tilde{M}_\Sigma$.

Since F and Σ are homologous in M , the Thom-Pontryagin maps $f_\Sigma : M \rightarrow S^1$ and $f_F : M \rightarrow S^1$, produced by Σ and F , are homotopic. Thus, they induce equivalent cyclic coverings of M . The spaces of these coverings are \mathbb{Z} -equivariantly homeomorphic with the homeomorphism $\phi : \tilde{M}_\Sigma \rightarrow \tilde{M}_F$ covering the identity map. We employ ϕ to identify the two spaces, and use the notation \tilde{M} for both of them. We also identify $\hat{\Sigma}$ with $\phi(\hat{\Sigma})$.

We denote by t the upward deck translation, a generator of the cyclic group \mathbb{Z} acting on \tilde{M} . Then \tilde{M} can be represented as a union $\bigcup_{n \in \mathbb{Z}} t^n(\tilde{M}_\Sigma)$. Here the fundamental region \tilde{M}_Σ is bounded by $t(\hat{\Sigma})$ and $\hat{\Sigma}$. Therefore, the 2-cycle $\hat{\Sigma}$ is a boundary of the 3-chain $\tilde{M}^{+\infty}(\hat{\Sigma}) := \bigcup_{n \geq 0} t^n(\tilde{M}_\Sigma)$. Hence, $\rho(\hat{\Sigma}, \tilde{M}) = 1$.

When $[\Sigma] = k[F]$, a similar argument applies to Σ and k parallel copies of a fiber F . Consider the covering $\tilde{M}_{kF} \rightarrow M$ induced by the Thom-Pontryagin map f_{kF} . As before, there exists a lifting $\hat{\Sigma} \subset \tilde{M}_{kF}$ which separates the positive and negative ends of \tilde{M}_{kF} . The covering $\tilde{M}_{kF} \rightarrow M$ factors through $\tilde{M} \rightarrow M$. The fiber of $p : \tilde{M}_{kF} \rightarrow \tilde{M}$ consists of k points; furthermore, \tilde{M}_{kF} is homeomorphic to k copies of \tilde{M} (this point will be explained later). Since $\hat{\Sigma} \subset \tilde{M}_{kF}$ separates the positive and negative ends of \tilde{M}_{kF} , the portion of $\hat{\Sigma}$ residing in each of the k copies of \tilde{M} separates the ends of the relevant copy. Applying the transfer p_* to the 3-chain $\tilde{M}_{kF}^{+\infty}(\hat{\Sigma})$ produces a 3-chain which bounds $p(\hat{\Sigma}) \subset \tilde{M}$. We notice that $p(\hat{\Sigma})$ consists of k surfaces, each of which separates the ends of \tilde{M} . This proves the first claim.

Assume that $[\Sigma]$ is not proportional to a fiber, but still admits a lifting from $\mathcal{B}(\Sigma)$. Then there exists a loop γ so that $[\gamma] \in \text{Ker } f_*$ and $[\Sigma] \circ [\gamma] \neq 0$ (cf. the proof of Lemma 6.5). Denote by $\hat{\gamma} \subset \tilde{M}$ a lift of γ . It is a loop. When $\hat{\Sigma} \in \mathcal{B}(\Sigma)$, then, for all n , $t^n([\hat{\gamma}]) \circ [\hat{\Sigma}] = 0$. Hence, $[\Sigma] \circ [\gamma] = 0$. This contradiction proves the second claim of the lemma. \square

For a given $\Sigma \subset M$ comprised of several components, a lifting $\hat{\Sigma} \in \mathcal{B}(\Sigma)$ is not unique, even up to deck translations (although for each component it is unique). For example, if a union $\hat{\Sigma}_0$ of a few components of $\hat{\Sigma} \in \mathcal{B}(\Sigma)$ is a boundary of a 3-chain in \tilde{M} , then we can apply any deck transformation to $\hat{\Sigma}_0$, while leaving

$\hat{\Sigma} \setminus \hat{\Sigma}_0$ intact, to produce a new lifting from $\mathcal{B}(\Sigma)$. Evidently, some restrictions on $\Sigma \subset M$ must be in place in order to claim the uniqueness, up to deck translations, of the lifting $\hat{\Sigma} \in \mathcal{B}(\Sigma)$.

We need to spell out the argument which we already used in the proof of Lemma 6.6. Let $z^k : S^1 \rightarrow S^1$ be the canonical map of degree k . For the time being, we choose the circle of radius $1/2\pi$. We denote by f_k the composition of $f : M \rightarrow S^1$ with z^k (so that $f_1 = f$), and let $\tilde{M}_k \rightarrow M$ be the cyclic covering induced by the f_k . One can view \tilde{M}_k as a balanced product $\tilde{M} \times_{\{T^n\}} \mathbb{Z}$. The cyclic T -action on the product $\tilde{M} \times \mathbb{Z}$ is defined by the formula $T(x, q) = (\tau(x), q - k)$, where $x \in \tilde{M}$, $q \in \mathbb{Z}$, and τ is the preferred generator of the cyclic action on \tilde{M} . The transformation $t : (x, q) \rightarrow (x, q + 1)$ commutes with T , and thus gives rise to a cyclic t -action on \tilde{M}_k . The obvious t -equivariant map $\tilde{M}_k \rightarrow \mathbb{Z}/k\mathbb{Z}$ is onto and divides \tilde{M}_k into k disjoint copies of \tilde{M} . Therefore, the natural k -to-1 map $p_k : \tilde{M}_k \rightarrow \tilde{M}$ splits.

Any t -equivariant function on \tilde{M}_k is generated by a function $\tilde{h} : \tilde{M} \times \mathbb{Z} \rightarrow \mathbb{R}$, subject to two properties: 1) $\tilde{h}(x, q + 1) = \tilde{h}(x, q) + 1$ (equivariance), and 2) $\tilde{h}(\tau(x), q - k) = \tilde{h}(x, q)$ (being a well-defined function on the balanced product). In particular, the t -equivariant function $\tilde{f} : \tilde{M} \rightarrow \mathbb{R}$ (which covers f) produces a function \tilde{h}_k with the properties 1) and 2) above: just put $\tilde{h}_k(x, q) = k \cdot \tilde{f}(x) + q$. We denote by \tilde{f}_k the function on \tilde{M}_k generated by \tilde{h}_k .

Now, at least for the important case when $[\Sigma]$ is k -proportional to the homology class $[F]$ of a fiber, we will give a more conceptual interpretation of the twist numbers $\rho(\Sigma, \sim)$ in terms of the cyclic cover $\tilde{M}_k \rightarrow M$ induced by the map f_k .

For each surface Σ representing a homology class $k[F]$, we can measure the number of times it is “wrapped” by f around S^1 : let

$$(6.2) \quad h(\Sigma; f) := 1 + \min_{\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)} \lceil \text{osc}(\tilde{f}_k|_{\hat{\Sigma}_k}) \rceil.$$

Here $\lceil r \rceil$ stands for the integral part of a real number r . Abusing previous notation, $\mathcal{B}_k(\Sigma)$ in (6.2) denotes the set of Σ -liftings which separate the positive and negative ends of \tilde{M}_k .

The following definitions aim to introduce notions of *breadth* and *height* of a given surface $\Sigma \subset M$ relative to a given map $f : M \rightarrow S^1$. They rely on Lemma 6.6. We always assume that surfaces Σ and F are in general position.

Definition 6.7. For a given map $f : M \rightarrow S^1$, let $\Sigma \subset M$ be a surface representing the k -multiple of the homology class of a fiber. Let F be a union of a few fiber components (not necessarily belonging to the same fiber). Denote by $\mathcal{A}(F)$ the set of all liftings $\{\hat{F}\}$ of F to the space \tilde{M}_k . Each surface \hat{F} intersects with only finitely many copies $\{t^n(\hat{\Sigma})\}_n$ of the surface $\hat{\Sigma}_k \subset \tilde{M}_k$, where $\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)$. We minimize the number of such copies over the set $\mathcal{A}(F)$, denote the minimum by $b(F, \Sigma)$, and call it the *breadth* of F relative to Σ .

Note that, for any f -fiber F' in general position with F , we have $F \cap F' = \emptyset$. Thus, $b(F, F') = 0$.

At this point, it is not clear why $b(F, \Sigma)$ does not depend on the choice of the special lifting $\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)$. Definition 6.7 will be justified by linking $b(F, \Sigma)$ directly with the quantity $\rho(\Sigma, F)$, which is independent on the liftings of Σ .

Proposition 6.8. *Assume that $\Sigma \subset M$ represents a k -multiple of the homology class of a fiber, and let F be any union of fiber components. Then $\rho(\Sigma, F) = b(F, \Sigma)$.*

Proof. For a lifting $\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)$, let $u_{\hat{\Sigma}_k}$ denote a step function which takes value 0 at the points of the domain $\tilde{M}_k^{-\infty}(\hat{\Sigma}_k)$ and value 1 at the points of $\tilde{M}_k^{+\infty}(\hat{\Sigma}_k)$. Recall that $\hat{\Sigma}_k$ is the boundary of the 3-chain $\tilde{M}_k^{+\infty}(\hat{\Sigma}_k)$.

Consider all the surfaces $\{t^n(\hat{\Sigma}_k)\}_n$ having a non-empty intersection with a particular lifting $\hat{F}_k \subset \tilde{M}_k$ of F . Restrict the potential function

$$u := \sum_{\{n \mid t^n(\hat{\Sigma}_k) \cap \hat{F}_k \neq \emptyset\}} u_{t^n(\hat{\Sigma}_k)}$$

to \hat{F}_k . Crossing in \hat{F}_k an oriented curve from $\hat{C}_n = \hat{F}_k \cap t^n(\hat{\Sigma}_k)$ in the positive normal direction is the same as crossing in \tilde{M}_k the corresponding component of $\hat{\Sigma}_k$ in the preferred normal direction: both have the effect of increasing the potential u by 1. Therefore, u gives rise to a 0-cochain on the oriented graph dual to the pattern $\bigcup_n \hat{C}_n$ in \hat{F}_k .

We observe that, since $p : \hat{\Sigma}_k \rightarrow \Sigma$ and $p : \hat{F}_k \rightarrow F$ are 1-to-1 maps, the covering map $p : \tilde{M}_k \rightarrow M$ defines a diffeomorphism of pairs $\coprod_n \hat{C}_n \subset \hat{F}_k$ and $\Sigma \cap F \subset F$, where the disjoint union employs all the non-empty \hat{C}_n 's. In particular, the images $\mathcal{C}_n := p(\hat{C}_n)$ of distinct intersections \hat{C}_n are disjoint in M .

Thus, u also produces a 0-cochain on the oriented graph, dual to the pattern $\Sigma \cap F$ in F (equivalently, an integral 2-chain on F whose boundary is $\Sigma \cap F$). It takes exactly as many consecutive values as the number of patterns $\{\hat{C}_n \neq \emptyset\}$. Hence, $\text{osc}(u|_F) \geq \rho(\Sigma, F) + 1$.

For each component F^β of F , the potential u is determined, up to a constant, by the oriented intersection $F^\beta \cap \Sigma \subset F^\beta$. Therefore, $\text{osc}(u|_{F^\beta}) = \rho(\Sigma, F^\beta) + 1$.

Next, we minimize $\text{osc}(u|_F)$ by independently applying deck translations in \tilde{M}_k to various components \hat{F}_k^β to achieve the equality $\text{osc}(u|_F) - 1 = \min_\beta \{\rho(\Sigma, F^\beta)\} := \rho(\Sigma, F)$. This can be done by moving each component \hat{F}_k^β above $\hat{\Sigma}_k$ and so that it has a non-empty intersection with the fundamental domain bounded by $\hat{\Sigma}_k$ and $t(\hat{\Sigma}_k)$. As a result, $b(F, \Sigma) = \rho(\Sigma, F)$. \square

Let $\mathcal{S}_{\chi_-}^k$ be the set of surfaces $\Sigma \subset M$ which realize the Thurston norm of $k[F]$, where $[F]$ stands for the homology class of an f -fiber. Among the $\Sigma \in \mathcal{S}_{\chi_-}^k$, consider surfaces with the minimal breadth $b(F_R, \Sigma)$. We denote this optimal breadth by $b_{\chi_-}(F_R, k[F])$. When $k = 1$, we also use the abbreviation $b_{\chi_-}(f)$ for $b_{\chi_-}(F_R; [F])$.

Combining Proposition 6.8 with Lemma 6.2, we get

Corollary 6.9. *Let R be the repeller set for a map $f : M \rightarrow S^1$, and $[F]$ the homology class of a fiber. Then*

$$b_{\chi_-}(F_R, k[F]) = \rho_{\chi_-}(k[F], F_R) = \min_{\Sigma \in \mathcal{S}_{\chi_-}^k} \{\rho_{\chi_-}^\circ(\Sigma, F_R)\}.$$

Lemma 6.10. *Let $\Sigma \subset M$ represent a k -multiple of the homology class of a fiber and let F be any finite union of fiber components. Then $\rho(\Sigma, F) \leq h(\Sigma, f) - \epsilon$, where $\epsilon = 0, 1$, depending on the particular location of F in M .*

When F is a fiber (hence, $k = 1$) and Σ is connected, then $\rho(\Sigma, F) = h(\Sigma; f) - \epsilon$, which implies a very weak dependence of $\rho(\Sigma, F)$ on the fiber F .

Proof. Take $\hat{\Sigma}_k$ which separates the positive and negative ends of \tilde{M}_k and delivers $h(\Sigma; f)$. Let $\hat{F}_k \subset \tilde{M}_k$ be a lifting of F which, together with $\hat{\Sigma}_k$, delivers $b(F, \Sigma) = \rho(\Sigma, F)$ (as described in the proof of Proposition 6.8). Denote by F^β a typical connected component of F .

Using the cyclic t -action on \tilde{M}_k , the set of n 's for which $t^n(\hat{\Sigma}_k) \cap \hat{F}_k^\beta \neq \emptyset$ is a reflection with respect to 0 of the set of n 's for which $\hat{\Sigma}_k \cap t^n(\hat{F}_k^\beta) \neq \emptyset$. Therefore, the cardinality of such n 's does not exceed $\text{osc}\{n : \hat{\Sigma}_k \cap t^n(\hat{F}_k^\beta) \neq \emptyset\}$. Since, for each n , $t^n(\hat{F}_k^\beta)$ belongs to a constant level set of the function f_k , and for distinct n these levels are integrally spaced, it follows that

$$\text{osc}\{n : \hat{\Sigma}_k \cap t^n(\hat{F}_k^\beta) \neq \emptyset\} \leq \epsilon + \lceil \text{osc}(\tilde{f}_k|_{\hat{\Sigma}_k}) \rceil := h(\Sigma, f) - \epsilon.$$

Therefore, $\#\{n : t^n(\hat{\Sigma}_k) \cap \hat{F}_k^\beta \neq \emptyset\} \leq h(\Sigma, f)$. Due to Proposition 6.8 and its proof, $\rho(\Sigma, F)$ is the maximum over all β of the LHS of the previous inequality. Hence, $\rho(\Sigma, F) \leq h(\Sigma, f) - \epsilon$.

When Σ is connected and F is a fiber, the same arguments show that $\rho(\Sigma, F) = h(\Sigma, f) - \epsilon$. Indeed, the connectivity of Σ forces it to cross all the “intermediate floors” $\{t^n(\hat{F}_k)\}$ between the top and the bottom one. \square

Definition 6.11. Employing (6.2), put $h_{\chi_-}(k[F], f) = \min_{\Sigma \in \mathcal{S}_{\chi_-}^k} \{h(\Sigma, f)\}$. We also use the abbreviation $h_{\chi_-}(f)$ for $h_{\chi_-}([F], f)$.

Crudely, the difference between height and breadth is like the difference between the degree and the number of non-zero monomials in a Laurent polynomial from the ring $\mathbb{R}[t, t^{-1}]$.




FIGURE 10.

When Σ admits an f -positive loop γ which hits it only once, one can deform f in such a way that Σ will have arbitrary big height and breadth with respect to the deformed map (cf. Example 8.12).

The following corollary is a tautology worth mentioning because its converse will be a focus of our efforts in Section 8. It follows from the observation that, if one can find a χ_- -optimal surface Σ among fibers or fiber components of a given map f , then you can find it disjoint from any other fixed union F of fibers or fiber components. Thus, $\rho(\Sigma, F) = 0$. When $b(F, \Sigma)$ makes sense, it vanishes as well.

Corollary 6.12. *Let $f : M \rightarrow S^1$ be any Morse map as in Lemma 4.2.*

- *If a finite union of fiber components delivers the minimal value of $\chi_-(\sim)$ in a vertical homology class $[\Sigma]$, then $\rho_{\chi_-}([\Sigma], F_R) = 0$.*
- *If a finite union of fiber components delivers the minimal value of $\chi_-(\sim)$ in the homology class $k[F]$, then $b_{\chi_-}(F_R, k[F]) = 0$.*
- *If a finite union of fiber components delivers the minimal value of $\chi_-(\sim)$ in the homology class $k[F]$, then $h_{\chi_-}(k[F], f) = 1$.* \square

By [T], compact leaves of taut (that is, no generalized Reeb components) foliations deliver the Thurston norm of their homology classes. Therefore,

Corollary 6.13. *Let a compact oriented surface $\Sigma \subset M$ be a union of leaves of a taut smooth foliation \mathcal{F} whose leaves are transversal to ∂M . Consider a Morse map $f : M \rightarrow S^1$ homotopic to the Thom-Pontryagin map $f_\Sigma : M \rightarrow S^1$, the homotopy being an identity on Σ . Then $\rho_{\chi_-}(f) = b_{\chi_-}(f) = 0$ and $h_{\chi_-}(f) = 1$.* \square

Although an effective computation of the invariants $b_{\chi_-}(F_R, k[F])$, $b_{\chi_-}(f)$, and $h_{\chi_-}(f)$ seems to be as difficult as the computation of the norm $\| [F] \|_T$, one has a good grip on how these invariants might change under an S^1 -controlled homotopy of the map f .

Lemma 6.14. *Let S^1 be a circle with circumference 1. Assume that a homotopy $\{f_\tau : M \rightarrow S^1\}_{0 \leq \tau \leq 1}$ is such that each τ -path $f_\tau(x)$, $x \in M$, winds less than q times around the circle. If $h_{\chi_-}(k[F], f_0) \leq l$, then $h_{\chi_-}(k[F], f_1) \leq l + 2kq$.*

When the image $f_\tau(x)$ of any point $x \in M$ moves clockwise, then $h_{\chi_-}(k[F], f_1) \leq l + kq$. Hence, $\rho_{\chi_-}(k[F], F_{R_1}) \leq l + kq$ as well.

Proof. Let $\Sigma^* \subset M$ be a surface which delivers the minimal value $\chi_-(\Sigma^*)$ of $\chi_-(\sim)$ in the k -multiple of the homology class of a fiber. Due to Lemma 6.6, Σ^* admits a special lifting $\hat{\Sigma}_k^* \in \mathcal{B}_k(\Sigma^*)$. In addition, assume that $\hat{\Sigma}_k^*$ is chosen so that $1 + [\text{osc}(\tilde{f}_{0,k}|_{\hat{\Sigma}_k^*})] = h(k[F], f_0) \leq l$.

For any Σ in the homology class $k[F]$, consider the function

$$\Phi(\tau, \hat{\Sigma}_k) := \text{osc}(\tilde{f}_{\tau,k}|_{\hat{\Sigma}_k}),$$

where $\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)$ and $\tilde{f}_{\tau,k} : \tilde{M}_k \rightarrow \mathbb{R}$ covers the homotopy f_τ . Note that, for any $x \in \tilde{M}_k$, $|\tilde{f}_{1,k}(x) - \tilde{f}_{0,k}(x)| < kq$. Hence, for any lifting $\hat{\Sigma}_k \in \mathcal{B}_k(\Sigma)$ we have $|\Phi(1, \hat{\Sigma}_k) - \Phi(0, \hat{\Sigma}_k)| < 2kq$, implying $\Phi(1, \hat{\Sigma}_k) < \Phi(0, \hat{\Sigma}_k) + 2kq$. Thus,

$$h(\Sigma^*, f_1) := 1 + \min_{\hat{\Sigma}_k^* \in \mathcal{B}_k(\Sigma^*)} [\Phi(1, \hat{\Sigma}_k^*)] \leq 1 + [\Phi(0, \hat{\Sigma}_k^*)] + 2kq = h(\Sigma^*, f_0) + 2kq.$$

Therefore, $h(k[F], f_1) \leq h(\Sigma^*, f_0) + 2kq := h(k[F], f_0) + 2kq$.

A similar argument is valid for a “clockwise” homotopy. \square

Corollary 6.15. *The S^1 -controlled size of a homotopy which links a given map f to a map with a χ_- -minimizing fiber or a union of fiber components gives an upper bound on $h_{\chi_-}(f)$, and thus on $\rho_{\chi_-}(f)$.*

Proof. Let $\Sigma \subset M$ be a surface minimizing $\chi_-(\sim)$ in the homology class of an f -fiber. Consider any Morse approximation f_1 of the Thom-Pontryagin map f_Σ (it is homotopic to f) such that $f_1|_\Sigma = f_\Sigma|_\Sigma = pt$. By Corollary 6.12, $h_{\chi_-}(f_1) = 1$. By a compactness argument, there exists a minimal natural number q so that the image of any point in M under the homotopy linking f with f_1 winds less than q times around the circle. By Lemma 6.14, $\rho_{\chi_-}(f) \leq h_{\chi_-}(f) \leq 1 + 2q$. \square

7. RESOLVING INTERSECTIONS WITH FIBERS

The main thrust of the arguments below is influenced by the proofs of Theorem 1 in [T] and of Theorem 2.3 in [H].

Given a Morse map $f : M \rightarrow S^1$, consider an embedded oriented surface $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ representing a vertical homology class. Let F be a finite union of fiber components in general position with respect to Σ . Denote by $\{C_i\}$ the components of the intersection $\mathcal{C} = \Sigma \cap F$. As before, each C_i is given an orientation by its preferred normal in F , which coincides with the preferred normal of $\Sigma \subset M$.

In a small neighborhood of each component C_i , the surfaces Σ and F divide M into *four* regions. The orientations of F and Σ pick a unique pair of non-adjacent quadrants, say I and III, along a typical intersection curve C_0 as shown in Figures 11 and 12.

Along C_i , we can *resolve* the intersection of F and Σ in a unique way, as shown in Figure 12, diagram A. The resolution $F \bowtie_i \Sigma$ will occupy a pair of non-adjacent quadrants. The resolved surface $F \bowtie_i \Sigma$ inherits the normal frames of Σ and F . These local resolutions $F \bowtie_i \Sigma$ can be pasted into a well-defined oriented surface $F \bowtie \Sigma$, homologous to $[\Sigma] + [F]$ (cf. Figure 11).




FIGURE 11. Resolving the intersection of Σ and F . Note the new intersection—the bold line—of the resolution $F \bowtie \Sigma$ with F .

As we are trying to paste $F \bowtie_i \Sigma$ together, we eliminate all the old intersection curves from $F \cap \Sigma$ and often are forced to introduce new intersections of $F \bowtie \Sigma$ with F . This happens because, over each connected component F_j° of $F \setminus (\Sigma \cap F)$, some of the “germs” of $F \bowtie_i \Sigma$ will reside *above* F , and some *below* it.

Consider, for instance resolving m coherently oriented meridians on a torus $T^2 = F$. The new surface $F \bowtie_i \Sigma$ still will have m intersection loops with the torus. In this example, the resolution does not help to simplify the intersection pattern. However, if at least two meridians have opposite orientations, the simplification becomes possible.

We intend to show that the new intersections are simpler than the original ones, and that, through iterations of resolutions, they can be eventually eliminated, thus producing a new embedded surface which does not intersect F at all.

In addition to the canonical resolution (diagram A in Figure 12), we also will use its modifications, shown in diagrams B, C.




FIGURE 12. Three ways of resolving the intersection.

Each time we resolve the intersection $\mathcal{C} = F \cap \Sigma$ along a particular component C_i of \mathcal{C} , we delete a regular neighborhood $U_i \approx C_i \times [-\epsilon, +\epsilon]$ of C_i from Σ and a regular neighborhood $V_i \approx C_i \times [-\delta, +\delta]$ of C_i from F and repaste the four components $\{C_i \times \pm\epsilon, C_i \times \pm\delta\}$ in a new way. Comparing the Euler numbers of $\Sigma \setminus \coprod_i U_i$ and $F \setminus \coprod_i V_i$ with those of $F \bowtie \Sigma$ and $F \coprod \Sigma$, we see that $\chi(F \bowtie \Sigma) = \chi(F) + \chi(\Sigma)$.

To describe the relation between $\chi_-(F \bowtie \Sigma)$ and $\chi_-(F) + \chi_-(\Sigma)$, we will use a few ideas from [T], pages 103-104. In general, the desired additivity $\chi_-(F \bowtie \Sigma) = \chi_-(F) + \chi_-(\Sigma)$ is upset by the new spherical or disk components generated as a result of the resolution. However, there are situations where the emergence of new spherical components can be prevented.

Recall that it is possible to perform 2-surgery on Σ using the 2-disks bounding in F the loops and arcs from $\Sigma \cap F$. If we perform the surgery starting with the most “inner” disks in F and gradually moving “outwards”, the resulting surface Σ^\odot is, up to the obvious isotopies, unique. Its intersection with F is free of loops which bound disks in F .

Definition 7.1. We say that $\Sigma \subset M$ is *well-positioned with respect to* $F \subset M$, if the transversal intersection $\Sigma^\odot \cap F$ has no components which bound a disk in Σ^\odot .

In particular, if $F \subset M$ is an *incompressible* surface, then any Σ is well-positioned with respect to it.

Note that if $\Sigma \cap F$ contains a loop which bounds a disk in Σ but not in F , then Σ is not well-positioned with respect to F . However, even if no such loop exists, it is still possible that the intersection $\Sigma^\odot \cap F$ will contain a loop which bounds a disk in Σ^\odot but not in F .

Lemma 7.2. *Let Σ represent an f -vertical homology class, and let F be a union of fiber components. If Σ is well-positioned with respect to F , then the surface $\Sigma^\odot \subset M$ has the following properties:*

- 1) Σ^\odot is cobordant to Σ ;
- 2) $\chi_-(\Sigma^\odot) \leq \chi_-(\Sigma)$;
- 3) $\chi_-(F \bowtie \Sigma^\odot) = \chi_-(F) + \chi_-(\Sigma^\odot)$.

In particular, if $\rho^\odot(\Sigma, F) = 0$,¹² then Σ is well-positioned with respect to F .

Proof. The new surface Σ^\odot has properties 1) and 2) (cf. Lemma 3.1). Its intersection $\Sigma^\odot \cap F$ is free of components bounding a disk in F . We claim that no *new* spherical or disk component S is present in $F \bowtie \Sigma^\odot$. Indeed, such an S would be a union $S_{\Sigma^\odot} \cup S_F$ of two surfaces whose common boundary is a subset of $\Sigma^\odot \cap F$. Here S_F is homeomorphic to a union of some domains in which $\Sigma^\odot \cap F$ divides F ,

¹²That is, $\Sigma^\odot \cap F = \emptyset$.

and S_{Σ^\odot} is homeomorphic to a union of some domains in which $\Sigma^\odot \cap F$ divides Σ^\odot . Note that, if a sphere is divided into two complementary domains, at least one of them must contain a boundary component which bounds in that domain a disk—a connected domain in the plane has a non-positive Euler number unless it is a disk. Since $\Sigma^\odot \cap F$ does not bound a disk in F (disk-bounding components have been eliminated by the 2-surgery) and in Σ^\odot (by being well-positioned), it is impossible to generate a new spherical or disk component S . \square

Remark. Here is the only point where a parallel program for the genus invariants faces similar but more serious difficulties. Unless the number of components in $F \bowtie \Sigma$ is less than or equal to the number of components in $F \sqcup \Sigma$, the desired equality $g(F \bowtie \Sigma) \leq g(F) + g(\Sigma)$ is not valid. Unfortunately, in general, we do not know how to control the number of components in $F \bowtie \Sigma$. However, if $F \cap \Sigma$ is connected, then $g(F \bowtie \Sigma) = g(F) + g(\Sigma)$. \square

Now, we use the dual graph $K_{\mathcal{C}}$ of $\mathcal{C} \subset F$ and its modifications as bookkeeping devices for recording the resolutions of the intersection $\Sigma \cap F$. We will subject $K_{\mathcal{C}}$ to elementary modifications, which will mimic particular ways (see Figure 12) of resolving the intersection \mathcal{C} . Some modifications will erase a curve from \mathcal{C} , will eliminate the corresponding edge in $K_{\mathcal{C}}$, and will merge the two vertices it connects into a single one.

First, we eliminate the edges which correspond to loops bounding a disk in F , or to arcs bounding a disk in $(F, \partial F)$. This elimination corresponds to 2-surgery, as described in the beginning of the proof of Theorem 5.2, where we eliminated a totally reducible pattern \mathcal{C} , while keeping the invariants $g(\Sigma)$, $\chi_-(\Sigma)$ on a decline.

After that, we consider the curves from the modified $\mathcal{C} = \mathcal{C}^\circ$ which correspond to the edges emanating from the vertices with the maximal value m of the potential u . Then we can perform the A-type resolutions along them. Next, we perform the B-type resolutions along the rest of the intersection curves. This has an effect on $K_{\mathcal{C}}$ of eliminating the edges of maximal type, merging all vertices of level $u = (m - 1)$ with the appropriate vertices of the m -level, and keeping the rest of the graph unchanged. The modified graph $K_{\tilde{\mathcal{C}}}$ will have a “truncated” level function \tilde{u} of its own. Moreover, u , being restricted to the portion $K_{\mathcal{C}}^{<m}$ of $K_{\mathcal{C}}$ below m , is the pull-back of \tilde{u} under the obvious map $K_{\mathcal{C}}^{<m} \rightarrow K_{\tilde{\mathcal{C}}}$. Hence, $\rho(\tilde{\mathcal{C}}) = \rho(\mathcal{C}) - 1$. The resulting surface $\Sigma^\odot \bowtie F$ is homologous to $[\Sigma] + [F]$ and has an intersection with F described by the graph $K_{\tilde{\mathcal{C}}}$.

Since the intersection $(\Sigma^\odot \bowtie F) \cap F$ consists of curves which are isotopic to the original curves from $\Sigma \cap F$ (cf. Figures 11, 12), and since we have excluded all the curves which bound a disk from our original intersection \mathcal{C} , the new intersection $(\Sigma^\odot \bowtie F) \cap F$ is free of disk-bounding curves as well.

This recipe can be repeated again and again until, after $\rho^\circ(\mathcal{C})$ iterations, we eliminate the intersection with F completely. In this algorithm, the potential u helps to paste the local resolutions (of A, B and C-types) together. The final surface $\Sigma_{\tilde{F}}$ resides in the homology class of $[\Sigma] + \rho^\circ(\mathcal{C}) \cdot [F]$.

After $\rho^\circ(\Sigma, F)$ resolutions, $\chi_-(\Sigma_{\tilde{F}}) \leq \chi_-(\Sigma) + \rho^\circ(\Sigma, F) \cdot \chi_-(F)$, provided that Σ was well-positioned with respect to F .

When Σ is not well-positioned, we need to add a correction term to the RHS of the inequality above. This correction term $\mu^\circ(\Sigma, F)$ equals twice the number of “new” spheres plus the number of “new” disks present in $\Sigma_{\tilde{F}} = \Sigma^\odot \bowtie \{\rho^\circ(\Sigma, F) \cdot F\}$,

but not in Σ^\odot or in $\rho^\circ(\Sigma, F) \cdot F$. In other words,

$$(7.1) \quad \mu^\circ(\Sigma, F) = \chi_+(\Sigma^\odot \bowtie \{\rho^\circ(\Sigma, F) \cdot F\}) - \chi_+(\Sigma^\odot) - \chi_+(\rho^\circ(\Sigma, F) \cdot F).$$

Each new sphere will consume at least two disks bounding a curve from $\Sigma^\odot \cap F$ which bounds a disk in Σ^\odot but does not bound a disk in F . Each new relative disk will require at least one such curve. Let us denote by $\nu^\circ(\Sigma, F)$ the number of such curves. Then $\mu^\circ(\Sigma, F) \leq 2\nu^\circ(\Sigma, F)$.

In special cases we can rule out the emergence of new spheres just from observing the intersection pattern $\Sigma \cap F$ in F . Let U be one of the connected domains into which $\Sigma^\odot \cap F$ divides F . Imagine that U contains a handle or, what is the same, that $d < 2 - \chi(U)$, where d is the number of components in ∂U . Evidently, such a U cannot contribute to a sphere in $\Sigma^\odot \bowtie F$. Moreover, further resolutions can only enlarge U , thus preserving the handle inside U . Such a case is described in a model Example 8.12 and depicted in Figure 15 in the next section.

We assemble these observations in

Lemma 7.3. *For any finite union F of fiber components and any oriented surface $\Sigma \subset M$ representing a vertical 2-homology class, there exists an embedded surface Σ' with the properties:*

- Σ' is homologous to the cycle $[\Sigma] + \rho^\circ(\Sigma, F)[F]$.
- $\chi_-(\Sigma') \leq \chi_-(\Sigma) + \rho^\circ(\Sigma, F) \cdot \chi_-(F) + \mu^\circ(\Sigma, F)$.
- $\Sigma' \cap F = \emptyset$.

If Σ is well-positioned with respect to F , then $\mu^\circ(\Sigma, F) = 0$. This is the case when $\rho^\circ(\Sigma, F) = 0$ (i.e. when $\Sigma^\odot \cap F = \emptyset$), or when F is an incompressible surface, or when $\Sigma^\odot \cap F$ divides F into domains, each of which contains a handle. \square

8. TWIST, VARIATION, AND THE χ_- -OPTIMIZATION

We are in position to derive our main results. Basically, we follow the train of thought presented in Section 5, but now we will bring to the game the twist and height invariants from Section 6 and the graph-theoretical considerations from Sections 3 and 4.

Consider an attractor set A and a repeller set R in Γ_f .

Given a formal combination $\sum_i \kappa_i F_i$ ($\kappa_i \in \mathbb{R}$), of f -oriented fiber components $\{F_i\}$, define its norm $\|\sum_i \kappa_i F_i\|$ by the formula $\sum_i |\kappa_i| \cdot \chi_-(F_i)$.

In a similar way, we introduce a χ_- -weighted l_1 -seminorm $\|\sim\|_A$ on $\mathbb{R}[A]$ by the formula

$$(8.1) \quad \|\kappa\|_A = \sum_{a \in A} |\kappa(a)| \cdot \chi_-(F_a).$$

Here $\kappa \in \mathbb{R}[A]$ and F_a denotes the fiber component corresponding to a point $a \in A \subset \Gamma_f$. When all $\chi_-(F_a) \neq 0$, the unit ball in this norm is a convex hull spanned by the vectors $\{\pm \chi_-(F_a)^{-1} a\}_{a \in A}$.

When Γ_f admits a tree decomposition $\bigsqcup_{r \in R} T_r^\pm$ as in Lemma 4.2, thanks to Lemma 4.6, we have an epimorphism $P : \mathbb{R}[A] \rightarrow H_2^f \otimes \mathbb{R}$. One can combine (8.1) with P to define a “vertical” seminorm $\|\sim\|_{H^f}$ on $H_2^f \subset H_2(M, \partial M; \mathbb{R})$ by the formula:¹³

$$(8.2) \quad \|[\Sigma]\|_{H^f} = \min_{\{\kappa \in \mathbb{R}[A] \mid P(\kappa) = [\Sigma]\}} \{\|\kappa\|_A\}$$

¹³It is easy to verify that the LHS of (8.2) satisfies the triangle inequality.

Alternatively, consider

Definition 8.1. The seminorm $\|[\Sigma]\|_{H^f}$ of a vertical homology class $[\Sigma]$ can be defined as the minimum of $\chi_-(\sim)$, taken over all finite oriented unions of distinct fiber components which represent $[\Sigma]$.

The equivalence of this less technical definition with the one given by (8.2) follows from Lemma 4.2 and 4.6 coupled with a familiar observation: replacing any fiber component by a union of components indexed by A can be accomplished via 2-surgery — an operation which decreases the value of $\chi_-(\sim)$.

Fix an f -vertical homology class $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$ represented by a \mathbb{Z} -linear combination of the cycles $\{[F_a]\}_a$ with coefficients $\{\alpha_a\}$. Let α_Σ stand for a function $A \rightarrow \mathbb{Z}$ defined by the formulas $\{\alpha_\Sigma(a) = \alpha_a\}$.

For any probe surface Σ representing $[\Sigma]$, denote by $\rho_\Sigma : R \rightarrow \mathbb{Z}_+$ a function which assigns to each element $r \in R$ the value $\rho^\circ(\Sigma, F_r)$, where F_r stands for the fiber component corresponding to the point $r \in \Gamma_f$.

We pick a basis $\{C_k\}$ of 1-cycles in $H_1(\Gamma_f; \mathbb{Z})$. Given a function $\delta : \Gamma_f \rightarrow \mathbb{Z}$ with finite support located in the complement to the vertices of Γ_f , denote by $\int_{C_k} \delta$ the sum $\sum_{y \in \text{supp}(\delta)} \delta(y) \cdot [F_y \circ C_k]$, where $[F_y \circ C_k]$ stands for the (algebraic) intersection number between the loop C_k and the fiber component F_y indexed by $y \in \Gamma_f$.




FIGURE 13.

Let $\mathcal{V}_\Sigma \subset \mathbb{Z}[A]$ be an affine sublattice defined by the system of linear equations:

$$(8.3) \quad \mathcal{V}_\Sigma := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = \int_{C_k} \alpha_\Sigma + \int_{C_k} \rho_\Sigma \right\},$$

where the loops $\{C_k\}$ form a basis in $H_1(\Gamma_f; \mathbb{Z})$. We notice that the first integral on the RHS of (8.3) depends only on the homology class $[\Sigma]$, while the second integral depends on its particular representative Σ , or rather on the function ρ_Σ which Σ defines on R .

Similarly, put

$$(8.4) \quad \mathcal{W}_\Sigma := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = \int_{C_k} \alpha_\Sigma \right\}.$$

By Lemmas 4.5 and 4.6, the subgroup $H_2^f \subset H_2(M, \partial M; \mathbb{Z})$ of vertical homology classes is isomorphic to the quotient $\mathbb{Z}[A]/\mathcal{T}_A$, where

$$\mathcal{T}_A := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = 0 \right\}.$$

We will be interested in special elements $\kappa^{\mathcal{V}_\Sigma} \in \mathcal{V}_\Sigma$ and $\kappa^{\mathcal{W}_\Sigma} \in \mathcal{W}_\Sigma$ which minimize the $\|\sim\|_A$ -norm. By its very definition, $\|\kappa^{\mathcal{W}_\Sigma}\|_A = \sum_a |\kappa_a^{\mathcal{W}_\Sigma}| \cdot \chi_-(F_a)$ is the minimal χ_- -characteristic among all \mathbb{Z} -combinations of fiber components $\{F_a\}$ which realize the given homology class $[\Sigma]$, that is, $\|\kappa^{\mathcal{W}_\Sigma}\|_A = \|P(\kappa^{\mathcal{W}_\Sigma})\|_{H^f}$. Recall that each fiber component F is cobordant to a union of a few F_a 's. This union is produced by performing 2-surgery on F . Hence, $\chi_-(F) \geq \sum_a \chi_-(F_a)$. Therefore, the surface $F_{best}^{[\Sigma]} := \coprod_{a \in A} \kappa_a^{\mathcal{W}_\Sigma} \cdot F_a$ also delivers $\|[\Sigma]\|_{H^f}$ — the minimum of the χ_- -characteristic among *all* combinations of fiber components which realize the homology class $[\Sigma]$. Here $\kappa_a^{\mathcal{W}_\Sigma} \cdot F_a$ stands for a disjoint union of $|\kappa_a^{\mathcal{W}_\Sigma}|$ fiber components residing in a regular neighborhood of the f -oriented surface F_a ; their orientations are prescribed by the sign of $\kappa_a^{\mathcal{W}_\Sigma}$.

Theorem 8.2. *Let $f : M \rightarrow S^1$ be a Morse map with no local extrema, no bubbling singularities¹⁴ and with $f : \partial M \rightarrow S^1$ being a fibration. Then, for any oriented surface Σ representing a vertical 2-homology class $[\Sigma] = \sum_a \alpha_a [F_a]$,*

$$(8.5) \quad \chi_-(\Sigma) \geq \sum_{a \in A} |\kappa_a^{\mathcal{V}_\Sigma}| \cdot \chi_-(F_a) - \sum_{r \in R} \rho^\circ(\Sigma, F_r) \cdot \chi_-(F_r) - \mu^\circ(\Sigma, F_R).$$

Here $\kappa^{\mathcal{V}_\Sigma} \in \mathbb{Z}[A]$ is the vector of the affine lattice (8.3)¹⁵ which minimizes the norm (8.1). Its norm depends only on the homology class $[\Sigma]$ and the twist function $\rho_\Sigma : R \rightarrow \mathbb{Z}_+$. The number $\mu^\circ(\Sigma, F_R)$, counting the new spherical and disk components in the resolution $\Sigma^\circ \bowtie F_R$, is defined by the formula (7.1) (with $F = F_R$).

Formula (8.5) can be expressed in terms of the vertical norms:

$$(8.6) \quad \chi_-(\Sigma) \geq \|[\Sigma] + \sum_{r \in R} \rho^\circ(\Sigma, F_r)[F_r]\|_{H^f} - \|\sum_{r \in R} \rho^\circ(\Sigma, F_r) \cdot F_r\| - \mu^\circ(\Sigma, F_R).$$

Proof. As in Section 6, we consider transversal intersections $\Sigma \cap F_r$ giving rise to the twist invariants $\rho_r(\Sigma) := \rho(\Sigma, F_r)$. As before, special attention is paid to the curves from $\Sigma \cap F_r$ which bound disks in F_r . They help us to perform 2-surgery on Σ , which can only diminish the value of $\chi_-(\sim)$ and $g(\sim)$. The surgery produces a surface Σ° .

Note that the resolutions along distinct $\Sigma^\circ \cap F_r$'s are completely independent and can be performed in any order. It will require $\rho_r^\circ(\Sigma) := \rho(\Sigma^\circ, F_r) = \rho^\circ(\Sigma, F_r)$ resolutions to separate Σ and F_r . Therefore, after at most $\rho_R^\circ(\Sigma) := \max_{r \in R} \{\rho_r^\circ(\Sigma)\}$ resolutions, the resolved surface $\Sigma^\circ \bowtie F_R$ will be separated from the fiber union $F_R := \bigcup_r F_r$. It will reside in the homology class $[\Sigma] + \sum_r \rho_r^\circ(\Sigma)[F_r]$.

By Lemma 7.3,

$$(8.7) \quad \chi_-(\Sigma^\circ \bowtie F_R) \leq \chi_-(\Sigma) + \sum_{r \in R} \rho_r^\circ(\Sigma) \cdot \chi_-(F_r) + \mu^\circ(\Sigma, F_R).$$

Now, as in the proof of Theorem 5.2, using the gradient and minus the gradient flows, we can push $\Sigma^\circ \bowtie F_R$ away from F_R and towards the union $F_A := \coprod_{a \in A} F_a$ (here we rely on Lemma 4.2). This push becomes possible only after a number of 2-surgeries using the descending and the ascending 2-disks of critical points as their cores. The 2-disks are facing the fibers F_r (see Figure 9 and a similar argument in the proof of Theorem 5.2). The surgery will produce a new surface Σ^* , homologous

¹⁴For example, when no homology class in H_2^f admits a spherical or disk representative, any intrinsically harmonic f will do.

¹⁵For a harmonic f , one can pick a basis of f -positive loops $\{C_k\}$ in the formulas (8.3), (8.4).

to $\Sigma^\odot \bowtie F_R$. By Lemma 3.1, $\chi_-(\Sigma^*) \leq \chi_-(\Sigma^\odot \bowtie F_R)$. Therefore, Σ^* resides in a regular neighborhood of F_A and its χ_- -characteristic is less than or equal to the RHS of (8.7). Denote by Σ_a^* the union of Σ^* -components $\Sigma_{a,i}^*$ residing in a regular neighborhood U_a of F_a . The retraction $U_a \rightarrow F_a$ induces maps $\Sigma_{a,i}^* \rightarrow F_a$ of degrees $\kappa_{a,i}$. Put $\kappa_a = \sum_i \kappa_{a,i}$. By an argument as in Lemma 5.3, $\chi_-(\Sigma_a^*) \geq \sum_a |\kappa_a| \cdot \chi_-(F_a)$. Therefore, we get

$$(8.8) \quad \chi_-(\Sigma) + \sum_{r \in R} \rho_r^\circ(\Sigma) \cdot \chi_-(F_r) + \mu^\circ(\Sigma, F_R) \geq \sum_{a \in A} |\kappa_a| \cdot \chi_-(F_a).$$

On the other hand, since Σ^* and $\Sigma^\odot \bowtie F_R$ are cobordant,

$$(8.9) \quad \sum_{a \in A} \kappa_a \cdot [F_a] = [\Sigma] + \sum_{r \in R} \rho_r^\circ(\Sigma) \cdot [F_r].$$

This equation implies that $\kappa \in \mathbb{Z}[A]$ defined by $\{\kappa(a) = \kappa_a\}$ belongs to the affine sublattice \mathcal{V}_Σ (cf. (8.3)). Indeed, just consider the intersection numbers of the basic 1-cycles $\{C_k\}$ with the LHS and RHS of (8.9).

Since we have little control over the κ , we safely minimize the RHS of (8.8) (the $\sim \| \cdot \|_A$ -norm of κ) over the set \mathcal{V}_Σ to get the desired inequality (8.5). \square

When all the twists $\{\rho_r^\circ(\Sigma) = 0\}_{r \in R}$, the spaces (8.3) and (8.4) coincide. Furthermore, for such a Σ , $\mu^\circ(\Sigma, F_R) = 0$. Therefore, for any Σ in the homology class $[\Sigma]$, $\chi_-(\Sigma) \geq \|[\Sigma]\|_{H^f}$. On the other hand, by definition, $\|[\Sigma]\|_T \leq \|[\Sigma]\|_{H^f}$. Hence, the vanishing $\{\rho_r^\circ(\Sigma) = 0\}$ implies $\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}$ —a union of fiber components delivers $\|[\Sigma]\|_T$.

Let $\Sigma \subset M$ be the best combination of f -fiber components which delivers the Thurston norm of $[\Sigma]$. If we deform f in such a way that the new map f_1 and the old f share the same set of the “repelling” fiber components, then vanishing of $\{\rho_r^\circ(\Sigma)\}$ for f implies vanishing of $\{\rho_r^\circ(\Sigma)\}$ for f_1 .

In combination with Corollary 6.12, these observations lead to

Corollary 8.3 (The Best Fiber Component Criterion). *Let $f : M \rightarrow S^1$ be a Morse map as in Theorem 8.2. Then, for any oriented surface Σ representing a vertical 2-homology class $[\Sigma] = \sum_a \alpha_a [F_a]$,*

- *If $\rho^\circ(\Sigma, F_R) = 0$, then Σ delivers the Thurston norm $\|[\Sigma]\|_T$.*
- *For such maps f , the Best Fiber Component Theorem $\{\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}\}$ is equivalent to the property $\rho_{\chi_-}([\Sigma], F_R) = 0$.*
- *Let f_1 be a map as in Theorem 8.2 and homotopic to f . Assume that the two maps share the same set of fiber components indexed by their repeller sets R and R_1 . Then $\{\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}\}$ implies $\{\|[\Sigma]\|_T = \|[\Sigma]\|_{H^{f_1}}\}$.*
- *When $[\Sigma] = [F]$ —the homology class of a fiber—, then $b_{\chi_-}(F_R, [F]) = 0$ implies $\| [F] \|_T = \| [F] \|_{H^f}$.* \square

Example 8.4. Often the best union of fiber components and the best fiber are quite different. The relation between them could be non-trivial, but it can be described in pure combinatorial terms involving $\pi_f : \Gamma_f \rightarrow S^1$, the 1-cochain $\tau_{\chi_-}(f)$ from Section 3 and a marking of edges corresponding to the spherical components.

Let us examine Figure 1. The fiber F corresponding to a ray from the center which intersects with the slanted edge of Γ_f (in the dark shaded sector) is comprised of two spherical components and a surface F_0 of genus 3. Note that F_0 is homologically trivial in M . Therefore, the two spheres $F_1 := F \setminus F_0$ are homologous to $[F]$.

Since $\chi_-(F_1) = 0$, F_1 is the best combination of fiber components. On the other hand, $\{\chi_-(f^{-1}(\theta))\}$ take values 4 and 2 only. We notice that the map f violates the Calabi positive loop property and, hence, is not intrinsically harmonic. \square

In general, the relation between solutions $\kappa^{\mathcal{V}_\Sigma}$ and $\kappa^{\mathcal{W}_\Sigma}$ of the two optimization problems (cf. (8.3) and (8.4)) is subtle. However, for special, so-called, *f-balanced* $[\Sigma]$'s, we can get a handle on the relation between the optimal norms $\|\kappa^{\mathcal{V}_\Sigma}\|_A$ and $\|\kappa^{\mathcal{W}_\Sigma}\|_A$.

Definition 8.5. We say that a homology class $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$ is *f-balanced*, if it is proportional, over the positive rationals, to the vertical class $[F_R]$.

In combinatorial terms, the proportion between the numbers of (\pm) -weighted repellers and the $[\Sigma]$ -supporting attractors along any loop C in Γ_f is positive and C -independent. The sign attached to each singleton depends on the orientation of the loop and the orientation of the singleton induced by the map $\pi_f : \Gamma_f \rightarrow S^1$.

For example, when all the f -fibers are connected, the homology class of a fiber or its multiples is *f-balanced*. Also, the class $[F_A] = [F_R]$, and hence is *f-balanced*.

We notice that the inequalities (8.7) can be relaxed by replacing each twist $\rho_r^\circ(\Sigma)$ by their maximum $\rho^\circ(\Sigma, F_R) := \max_{r \in R} \{\rho_r^\circ(\Sigma)\}$. In other words, one can employ the same number $\rho^\circ(\Sigma, F_R)$ of resolutions at each F_r to create a “less optimal” surface $\Sigma^\odot \bowtie F_R$ which might contain a few extra copies of some F_r 's. Hence, (8.8) and (8.9) will be modified:

$$(8.10) \quad \chi_-(\Sigma) + \rho^\circ(\Sigma, F_R) \cdot \sum_{r \in R} \chi_-(F_r) + \mu^\circ(\Sigma, F_R) \geq \sum_{a \in A} |\kappa_a| \cdot \chi_-(F_a),$$

$$(8.11) \quad \sum_{a \in A} \kappa_a \cdot [F_a] = [\Sigma] + \rho^\circ(\Sigma, F_R) \cdot \sum_{r \in R} [F_r].$$

Employing (8.11), we define affine sublattices in $\mathbb{Z}[A]$ — modified versions of (8.3), (8.4), — by prescribing the intersection numbers of the 2-cycle $\sum_{a \in A} \kappa_a \cdot [F_a]$ with a basis of loops $\{C_k \subset M\}_k$ in $H_1(\Gamma_f, \mathbb{Z})$:

$$(8.12) \quad \tilde{\mathcal{V}}_\Sigma := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = ([\Sigma] \circ C_k) + \rho^\circ(\Sigma, F_R) \cdot ([F_R] \circ C_k) \right\},$$

$$(8.13) \quad \mathcal{W}_\Sigma := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = ([\Sigma] \circ C_k) \right\},$$

$$(8.14) \quad \mathcal{U} := \left\{ \kappa \in \mathbb{Z}[A] : \int_{C_k} \kappa = ([F_R] \circ C_k) \right\}.$$

As before, we are interested in vectors $\kappa^{\tilde{\mathcal{V}}_\Sigma} \in \tilde{\mathcal{V}}_\Sigma$, $\kappa^{\mathcal{W}_\Sigma} \in \mathcal{W}_\Sigma$, $\kappa^{\mathcal{U}} \in \mathcal{U}$ which will minimize the $\|\sim\|_A$ -norm. When $[\Sigma]$ is *f-balanced* (i.e. proportional to $[F_R]$), the vectors $\{[\Sigma] \circ C_k\}_k$ and $\{[F_R] \circ C_k\}_k$ are *proportional* with a positive coefficient of proportionality. As a result, we can assume that $\kappa^{\tilde{\mathcal{V}}_\Sigma} = \kappa^{\mathcal{W}_\Sigma} + \rho^\circ(\Sigma, F_R) \cdot \kappa^{\mathcal{U}}$, where $\kappa^{\mathcal{W}_\Sigma}$ and $\kappa^{\mathcal{U}}$ are proportional with a positive proportionality coefficient. By definition, $\|\kappa^{\mathcal{W}_\Sigma}\|_A = \|[\Sigma]\|_{H^f}$. We notice that $\|\kappa^{\mathcal{U}}\|_A = \|[F_A]\|_{H^f}$. Indeed, use Lemma 4.3 to replace $C_k \circ F_R$ with $C_k \circ F_A$ in (8.12) and (8.14).

Minimizing the RHS of (8.10), subject to (8.11), and using that

$$\begin{aligned} \|\kappa^{\tilde{\mathcal{V}}_\Sigma}\|_A &= \|\kappa^{\mathcal{W}_\Sigma}\|_A + \rho^\circ(\Sigma, F_R) \cdot \|\kappa^{\mathcal{U}}\|_A = \|[\Sigma]\|_{H^f} + \rho^\circ(\Sigma, F_R) \cdot \|[F_A]\|_{H^f} \\ &= \|[\Sigma] + \rho^\circ(\Sigma, F_R) \cdot [F_A]\|_{H^f}, \end{aligned}$$

we get our main result:

Theorem 8.6. *Let $f : M \rightarrow S^1$ be a Morse map with $f : \partial M \rightarrow S^1$ being a fibration. Assume that f has no local extrema and no bubbling singularities. Let $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$ be an f -balanced class. Put $\|F_R\| := \chi_-(F_R)$.*

Then, for any oriented surface $(\Sigma, \partial\Sigma) \subset (M, \partial M)$ representing $[\Sigma]$,

$$(8.15) \quad \chi_-(\Sigma) \geq \|[\Sigma]\|_{H^f} - \rho^\circ(\Sigma, F_R) \cdot \text{Var}_{\chi_-}(f) - \mu^\circ(\Sigma, F_R),$$

where $\text{Var}_{\chi_-}(f) := \|F_R\| - \|F_A\|_{H^f} = \|F_R\| - \|F_R\|_{H^f}$.

When Σ is well-positioned with respect to F_R , then $\mu^\circ(\Sigma, F_R) = 0$.

For a self-indexing map f , the variation $\|F_R\| - \|F_R\|_{H^f}$ is equal to the number of (non-bubbling) f -singularities. \square

Corollary 8.7. *The statements of Theorem 8.6 are valid when f is intrinsically harmonic and no non-trivial class in H_2^f admits a representation by spheres and disks. For instance, this is the case if the Hurewicz homomorphism $\pi_2(M, \partial M) \rightarrow H_2(M, \partial M; \mathbb{Z})$ is trivial. In particular, Theorem 8.6 is valid for harmonic Morse maps of irreducible 3-manifolds.* \square

Let $\mu^\circ([\Sigma], F_R)$ denote the minimum of (7.1) taken over all surfaces Σ which deliver the Thurston norm of $[\Sigma]$ and the value $\rho^\circ([\Sigma], F_R)$.

Minimizing the RHS of (8.15) over the set of all surfaces which deliver the Thurston norm of $[\Sigma]$ and employing Definition 6.3, we get

Corollary 8.8. *Let $f : M \rightarrow S^1$ be as in Theorem 8.6, and let $[\Sigma] \in H_2^f$ be a balanced class. Then its Thurston seminorm $\|[\Sigma]\|_T$ can be compared to the f -vertical seminorm $\|[\Sigma]\|_{H^f}$ on H_2^f :*

$$(8.16) \quad \|[\Sigma]\|_{H^f} \geq \|[\Sigma]\|_T \geq \|[\Sigma]\|_{H^f} - \rho_{\chi_-}([\Sigma], F_R) \cdot \text{Var}_{\chi_-}(f) - \mu^\circ([\Sigma], F_R)$$

In particular, if the twist $\rho_{\chi_-}([\Sigma], F_R)$ or the variation $\text{Var}_{\chi_-}(f)$ vanishes, then $\mu^\circ([\Sigma], F_R) = 0$, and we get $\|[\Sigma]\|_T = \|[\Sigma]\|_{H^f}$. \square

Vanishing of the variation $\text{Var}_{\chi_-}(f) \geq \chi_-(F_R) - \chi_-(F_A)$ is a rare event: it can only happen when $\chi_-(F_R) = \chi_-(F_A)$. In such a case, by Lemma 3.5, all the singularities of f must be of the bubbling type. If all the singularities are bubbling, then all the fiber components are incompressible. Indeed, let S_i^2 be a spherical fiber component residing in the vicinity of a critical point x_i . We cut M open along $\bigsqcup_i S_i^2$ so that it decomposes into a number of spherical rings $S_i^2 \times [0, 1]$ and the rest, which we denote by M° . We can attach a disk D_i^3 to each $S_i^2 \subset M^\circ$ and extend f in an obvious way across the disk. For each i , the new map $f' : M' \rightarrow S^1$ will have exactly one new critical point y_i —a local extremum—located at the center of D_i^3 . A standard deformation f'' of f' will cancel each pair (x_i, y_i) , thus producing a fibration f'' . Its fibers are incompressible. Now, if a disk $D^2 \subset M$ bounds a loop γ residing in an f -fiber $F \subset M^\circ$, then there is disk $D^2 \subset M'$ which also bounds γ . Since fibers of f'' are incompressible, γ must bound a disk in F . Therefore, if $\chi_-(F_R) = \chi_-(F_A)$, then for any Σ we have $\mu^\circ(\Sigma, F_R) = 0$.

Although the hypotheses of Theorem 8.6 exclude the bubbling singularities, its generalization—Theorem 8.13,—allows them. A model example can be produced by attaching 1-handles to a disjoint union of fibrations over the circle and extending the fibering maps across the handles as depicted in Figure 16.

Corollary 8.9. *Let $f : M \rightarrow S^1$ and Σ be as in Corollary 8.8. Assume that a surface Σ which delivers the Thurston norm $\|[\Sigma]\|_T$ together with the twist $\rho_{\chi_-}([\Sigma], F_R)$ is well-positioned with respect to F_R . If $\text{Var}_{\chi_-}(f) \neq 0$, then*

$$(8.17) \quad \rho_{\chi_-}([\Sigma], F_R) \geq \frac{\|[\Sigma]\|_{H^f} - \|[\Sigma]\|_T}{\|F_R\| - \|F_R\|_{H^f}}.$$

When Σ represents the balanced homology class $[F]$ of a fiber, then the RHS of (8.17) also gives a lower bound on the size of the breadth $b_{\chi_-}(f)$ and height $h_{\chi_-}(f)$.

Therefore, if the variation $\|F_R\| - \|F_R\|_{H^f}$ is relatively small and $\|F\|_{H^f} \gg \|F\|_T$, then any well-positioned surface Σ which delivers the Thurston norm of the homology class $[F]$ must be \hat{f} -tall. \square

Now, we consider maps f whose graphs Γ_f are very special. For them, the homology class of a fiber is balanced. The self-indexing maps evidently fall in that category.

Theorem 8.6 and Corollary 8.8, in combination with Proposition 6.8 and Lemma 6.10, imply a relaxing progression of inequalities:

Corollary 8.10. *Let $f : M \rightarrow S^1$ be as in Corollary 8.8. Assume that the homology class $[F]$ of a fiber is balanced. Let $\Sigma \subset M$ represent $k[F]$ and be well-positioned with respect to F_R . Then (8.15) implies*

$$(8.18) \quad \chi_-(\Sigma) \geq \|F\|_{H^f} - b(F_R, \Sigma) \cdot \text{Var}_{\chi_-}(f),$$

$$(8.19) \quad \chi_-(\Sigma) \geq \|F\|_{H^f} - h(\Sigma; f) \cdot \text{Var}_{\chi_-}(f).$$

Similarly, (8.16) implies

$$(8.20) \quad \|F\|_T \geq \|F\|_{H^f} - b_{\chi_-}(F_R, k[F]) \cdot \text{Var}_{\chi_-}(f),$$

$$(8.21) \quad \|F\|_T \geq \|F\|_{H^f} - h_{\chi_-}(k[F], f) \cdot \text{Var}_{\chi_-}(f). \quad \square$$

The corollary below and Figure 14 depict a special case of the balanced fiber.

Corollary 8.11. *Let $f : M \rightarrow S^1$ be a Morse map with no local extrema and whose restriction to the boundary ∂M is a fibration. Assume that all the fiber components $\{F_r\}_{r \in R}$, actually, are fibers (cf. Figure 14). Let Σ be an oriented surface which represents the homology class $[F]$ of a fiber and is well-positioned with respect to F_R . Let F_{best} denote an oriented union of fiber components which delivers $\|F\|_{H^f}$. Then*

$$(8.22) \quad \chi_-(\Sigma) \geq \chi_-(F_{\text{best}}) - \rho^\circ(\Sigma, F_R) \cdot \sum_{r \in R} [\chi_-(F_r) - \chi_-(F_{\text{best}})].$$

As a result,

$$(8.23) \quad \|F\|_T \geq \|F\|_{H^f} - \rho_{\chi_-}([F], F_R) \cdot \sum_{r \in R} [\chi_-(F_r) - \|F\|_{H^f}].$$

In particular, (8.22) and (8.23) are valid for any map f with all its fibers connected, provided that a well-positioned Σ delivering $\|F\|_T$ and $\rho_{\chi_-}([F], F_R)$ exists. \square

Example 8.12. We would like to recycle the algorithm from the harmonic twister example. We will see that this algorithm produces maps f with few singularities and a small variation, but arbitrarily big twists and heights.

Let us start with the fibration $f_0 : T \rightarrow S^1$ of the solid torus $T = D^2 \times S^1$ over the circle. The χ_- -number of its fiber D^2 is zero. By deforming f_0 slightly inside T ,





FIGURE 14.

FIGURE 15. A lift of the worst f'_k -fiber to the space of the cyclic cover $D^2 \times \mathbb{R} \rightarrow D^2 \times S^1$. The product structure is defined by f_0 .

we can introduce a pair $\{a, b\}$ of index 1 and index 2 critical points, while keeping all the fibers *connected*: for homological reasons, the pattern depicted in the left hand side of diagram E of Figure 5 cannot be realized by a trivial cobordism.

Now, as in the harmonic twister example, we deform f'_0 into a new map f'_k while keeping the deformation fixed on the boundary ∂T . This is done by applying move A from Figure 5 k times. In the process, b overtakes a exactly k times in the race around the circle. By Lemma 6.14, $h_{\chi_-}(f'_k) \leq 1 + k$. As Figure 5, A, testifies, all the fibers of f'_k still must be connected. Also, because of the same diagram A, $\chi_-(F'_{k,best}) = 2k$. Evidently, $Var_{\chi_-}(f'_k) = 2$.

In fact, the original fiber D^2 is well-positioned with respect to the surface $F_{r,k}$ — the worst fiber of f'_k . To validate this fact requires a more careful analysis of the geometry of $F_{r,k}$ relative to f_0 , as depicted in Figure 15 for $k = 0, 1, 2, 3$. The figure shows a lift $\hat{F}_{r,k}$ of $F_{r,k}$ to the space $D^2 \times \mathbb{R}$ of the cyclic cover. The surface $\hat{F}_{r,k}$ is comprised of a number of left and right 1-handles attached to a disk marked with a bold circle. Deforming f'_{k-1} into f'_k results in attaching a new right ($k \equiv 0(2)$) or left ($k \equiv 1(2)$) handle to $\hat{F}_{r,k-1}$ and stretching the old handles. The dots with dark (light) centers show the locations of the \tilde{f}'_k -critical points of index 1 (2).

As we resolve intersections of $\hat{F}_{r,k}$ with multiple translates of $\{t^n(\hat{D}^2)\}$ of the \tilde{f}_0 -fiber \hat{D}^2 (they are marked with flags in Figure 15), we see that each translate cuts through $\hat{F}_{r,k}$ in a way that leaves at least one handle of $\hat{F}_{r,k}$ to the left and one handle to the right of $t^n(\hat{D}^2)$. Hence, no new 2-disks or spheres are produced as a result of the resolutions, i.e. $\mu^\circ(D^2, F_{r,k}) = 0$. By Theorem 8.6, $h_{\chi_-}(f'_k) \geq \rho_{\chi_-}(f'_k) \geq k$. Furthermore, Figure 15 testifies that $\rho_{\chi_-}(f'_k) = k$.

Since the deformed maps f'_k coincide with f_0 on the boundary ∂T , we can use them for twisting any given map $f : M \rightarrow S^1$: just take a regular neighborhood of an f -positive loop $\gamma \subset M$ in the role of T . \square

Let maps $\{f_\alpha : M_\alpha \rightarrow S^1\}_\alpha$ be as in Theorem 8.2. As we attach 1-handles $\{T_\beta \approx S^2_\beta \times D^1_\beta\}_\beta$ to $\coprod_\alpha M_\alpha$ and extend the maps across the handles as shown in Figure 16, we form a new manifold M and a new map $f : M \rightarrow S^1$. Its graph Γ_f is obtained from $\coprod_\alpha \Gamma_{f_\alpha}$ by attaching to it a few new edges, each of which contributes a new pair of trivalent vertices of indices 1 and 2. Moreover, the orientation of these new edges is such that each of them must contain an attractor. In other words, 1-surgery does not change the repeller set! New branches will be added to the original trees $\{T_r^\pm\}$ covering $\coprod_\alpha \Gamma_{f_\alpha}$, and the new trees with the old roots will cover the new graph Γ_f (as in Lemma 4.2).




FIGURE 16. Attaching a slanted 1-handle.

From our point of view, the “simplest” case of fibrations f_α is a bit subtle: formally, fibrations do not satisfy the basic assumptions of Lemma 4.2. Recall that our techniques rely on the decomposition of Γ_{f_α} into the trees $\{T_{\alpha,r}^\pm\}$. Attaching handles to the space of a fibration might produce a loop in the new graph only with bubbling vertices of the same index (cf. Figure 6). In such a case, we need first to deform the fibering map f_α slightly so that a pair of canceling critical points of indices 1 and 2 is introduced. We can choose a deformation which will not disturb the majority of fibers. So, the best fiber of the deformed map remains the winner in its homology class. Effectively, the deformation introduces one repeller along the loop which represents the graph Γ_{f_α} . After such conditioning of f_α , we are ready to attach 1-handles any way we want.

Consider the natural epimorphism

$$(8.24) \quad \begin{aligned} \Phi_2 : H_2(M, \partial M; \mathbb{Z}) &\xrightarrow{j_2} H_2(M, \partial M \sqcup (\bigsqcup_\beta S^2_\beta); \mathbb{Z}) \\ &\approx \bigoplus_\alpha H_2(M_\alpha, \partial M_\alpha \sqcup (\bigsqcup_{\beta \in B_\alpha} D^3_{\alpha\beta}); \mathbb{Z}) \approx \bigoplus_\alpha H_2(M_\alpha, \partial M_\alpha; \mathbb{Z}), \end{aligned}$$

where $\{D_{\alpha\beta}^3\}_{\beta \in B_\alpha}$ denote 3-disks in M_α bounding the bases of the 1-handles attached to M_α . For every class $[\Sigma] \in H_2(M, \partial M; \mathbb{Z})$, let $\Phi_2([\Sigma]) = \bigoplus_\alpha [\Sigma]_\alpha$. The epimorphism Φ_2 splits. We pick a splitting homomorphism

$$\Psi_2 : \bigoplus_\alpha H_2(M_\alpha, \partial M_\alpha; \mathbb{Z}) \rightarrow H_2(M, \partial M; \mathbb{Z}).$$

Geometrically, Ψ_2 can be defined by picking a basis of 2-cycles in

$$\bigoplus_\alpha H_2(M_\alpha, \partial M_\alpha; \mathbb{Z})$$

and isotoping them away from the 3-disks $\{D_{\alpha\beta}^3\}$. There is nothing canonical about our choice of Ψ_2 .

Note that $\Phi_2([S_\beta^2]) = 0$. Therefore, vertical classes are mapped by Φ_2 to vertical classes: $\Phi_2(H_2^f) = \bigoplus_\alpha H_2^{f_\alpha}$.

Our goal is to describe the relation between the norms $\|[\Sigma]\|_T$ and $\{\|[\Sigma]_\alpha\|_T\}_\alpha$. In the process, we also will describe the relation between the twists $\rho_{\chi_-}([\Sigma], F_R)$ and $\{\rho_{\chi_-}([\Sigma]_\alpha, F_{R_\alpha})\}$ of the maps f and $\{f_\alpha\}$.

Let $\{\Sigma_\alpha \subset M_\alpha\}$ be surfaces which deliver the norms $\|[\Sigma_\alpha]\|_T$ of their f_α -vertical homology classes $[\Sigma_\alpha]$ and the values of the twist invariants $\rho([\Sigma_\alpha], F_{R_\alpha})$ for the maps $\{f_\alpha\}$. Let $[\Sigma]$ be a vertical class, so that $\Phi_2([\Sigma]) = \bigoplus_\alpha [\Sigma]_\alpha$. Consider the surface $\bigsqcup_\alpha \Sigma_\alpha \subset M$. Since the kernel of Φ_2 is spanned by the spheres $\{S_\beta^2\}$, for some integers $\{\kappa_\beta\}$, we have $[\Sigma] = \sum_\beta \kappa_\beta [S_\beta^2] + \sum_\alpha \Psi_2([\Sigma_\alpha])$. Therefore, $[\Sigma]$ admits a representative Σ which is a disjoint union of $\bigsqcup_\alpha \Sigma_\alpha \subset M \setminus (\bigsqcup_{\alpha,\beta} D_{\alpha,\beta}^3)$ with a bunch of vertical spheres. Recall, that the repeller set R for Γ_f is the disjoint union of the repeller sets $\{R_\alpha \subset \Gamma_f\}$. The spheres do not contribute to $\rho(\Sigma, F_R)$: they are disjoint from the surface F_R . Therefore, $\rho_{\chi_-}([\Sigma], F_R) \leq \max_\alpha \rho_{\chi_-}([\Sigma_\alpha], F_{R_\alpha})$. Also, $\chi_-(\Sigma) = \sum_\alpha \chi_-(\Sigma_\alpha) = \sum_\alpha \|[\Sigma_\alpha]\|_T$. So, $\|[\Sigma]\|_T \leq \sum_\alpha \|[\Sigma_\alpha]\|_T$.

A similar argument is applicable to an estimation of $h_{\chi_-}(f)$. Consider surfaces $\{\Sigma_\alpha \subset M_\alpha \setminus (\bigsqcup_\beta D_{\alpha,\beta}^3)\}$ which deliver the norm $\|F_\alpha\|_T$ of the f_α -fiber homology class $[F_\alpha]$. Among them pick Σ_α with the minimal value of $h(\Sigma_\alpha, f_\alpha)$ (cf. (6.2)). One can align their special liftings $\{\hat{\Sigma}_\alpha \subset \tilde{M}_\alpha\}$ so that they will be located above $\tilde{f}_\alpha^{-1}(0)$ and will have a non-empty intersection with $\tilde{f}_\alpha^{-1}((0, 1])$. The space \tilde{M} is built by performing an equivariant 1-surgery on $\bigsqcup_\alpha \tilde{M}_\alpha$. Although $\{\hat{\Sigma}_\alpha \in \mathcal{B}_1(\Sigma_\alpha)\}$, $\bigsqcup_\alpha \hat{\Sigma}_\alpha$ is not automatically in $\mathcal{B}_1(\bigsqcup_\alpha \Sigma_\alpha)$ in \tilde{M} : one needs to add to $\bigsqcup_\alpha \hat{\Sigma}_\alpha$ a few spheres $\{\hat{S}_\beta^2\}$ from $\tilde{f}^{-1}(1) \cap (\bigsqcup_\beta \tilde{T}_\beta)$ to get a surface $(\bigsqcup_\alpha \Sigma_\alpha) \sqcup (\bigsqcup_\beta S_\beta^2)$ which admits a special lifting bounding a 3-chain in \tilde{M} . However, this addition will not change the value of $\chi_-(\bigsqcup_\alpha \Sigma_\alpha)$. Therefore, $h_{\chi_-}(f) \leq \max_\alpha \{h_{\chi_-}(f_\alpha)\}$.

These considerations lead to

Theorem 8.13. *Let maps $\{f_\alpha : M_\alpha \rightarrow S^1\}_\alpha$ be as in Theorem 8.2. Let $f : M \rightarrow S^1$ be constructed by performing 1-surgery on the map $\bigsqcup_\alpha f_\alpha$. Then, for any f -vertical class $[\Sigma]$,*

$$(8.25) \quad \|[\Sigma]\|_T \leq \sum_\alpha \|[\Sigma_\alpha]\|_T,$$

where $\Phi_2([\Sigma]) = \bigoplus_\alpha [\Sigma]_\alpha$. Therefore, 1-surgery on a map which satisfies the hypotheses of Theorem 8.2 does not increase the Thurston norm of its vertical classes (under any splitting homomorphism Ψ_2).

Also, this 1-surgery does not increase the twist and height invariants:

$$(8.26) \quad \rho_{\chi_-}([\Sigma], F_R) \leq \max_{\alpha} \{\rho_{\chi_-}([\Sigma_{\alpha}], F_{R_{\alpha}})\},$$

$$(8.27) \quad h_{\chi_-}(f) \leq \max_{\alpha} \{h_{\chi_-}(f_{\alpha})\}. \quad \square$$

Although the new map f has bubbling singularities, the conclusions of Lemma 4.2 are still valid for its graph Γ_f : it admits a cover by trees rooted at $R = \bigsqcup_{\alpha} R_{\alpha}$ — a fact central to our previous arguments. Indeed, each new edge contains an attractor. Therefore,

Corollary 8.14. *If, for each f_{α} as in Theorem 8.2, a union $F_{\alpha} \subset M_{\alpha} \setminus (\bigsqcup_{\beta} D_{\alpha,\beta}^3)$ of fiber components delivers the Thurston norm, that is, if $\|F_{\alpha}\|_{Hf_{\alpha}} = \|F_{\alpha}\|_T$, then $\bigsqcup_{\alpha} F_{\alpha} \subset M$ delivers the Thurston norm of any class $[\Sigma]$ whose Φ_2 -image is $\bigoplus_{\alpha} [F_{\alpha}]$.*

In particular, performing 1-surgery on a disjoint union of fibrations $\{f_{\alpha}\}$ results in a map f whose fiber F delivers $\|F\|_T$.

Proof. By Corollary 8.3, $\|F_{\alpha}\|_{Hf_{\alpha}} = \|F_{\alpha}\|_T$ is equivalent to $\rho([F_{\alpha}], F_{R_{\alpha}}) = 0$. Using (8.25) from Theorem 8.13 and again Corollary 8.3, the first claim follows.

In the case when f_{α} is a fibration, first we slightly deform the fibering map, so that a pair of canceling critical points of indices 1 and 2 is introduced and at least one fiber is untouched by the deformation. After this, we attach 1-handles. \square

Our results can be applied to links in 3-manifolds.

We define the Thurston norm $\chi_-(L, M)$ of a (framed) link L in a closed 3-manifold M to be the minimum of χ_- -invariants of all oriented embedded surfaces which bound L . Similarly, denote by $g(L, M)$ the genus of the link L .

In many interesting cases the Thurston norm $\chi_-(L, M)$ coincides with the Alexander norm $\|(L, M)\|_{Al}$ of $L \subset M$ (see [Mc] for the definition) and is delivered by Seifert's algorithm [Cr], [Mur]. In general, $\|(L, M)\|_{Al} \leq \chi_-(L, M)$, provided that the first Betti number $b_1(M \setminus L) \geq 2$ [Mc].

At least for connected sums of fibered links, we will compute $\chi_-(L, M)$ in terms of the vertical norms.

Let L be a framed link in an oriented, closed 3-manifold M , and U its open tubular neighborhood. We denote by M° the manifold $M^{\circ} = M \setminus U$. The framing defines an embedding $L \subset \partial M^{\circ}$.

In view of Theorem 8.13 and Corollary 8.14, we get the following proposition.

Corollary 8.15. *Let L_{α} be a framed fibered link in an oriented closed 3-manifold M_{α} . Let $(F_{\alpha}, \partial F_{\alpha}) \subset (M_{\alpha}^{\circ}, \partial M_{\alpha}^{\circ})$ be an oriented surface which bounds L_{α} and delivers*

$\chi_-(L_{\alpha}, M_{\alpha})$. Denote by $f_{\alpha} : M_{\alpha}^{\circ} \rightarrow S^1$ a Morse map which satisfies the hypotheses of Theorem 8.2 and with the surface F_{α} as one of its fibers.¹⁶

Let M be a manifold constructed by performing 1-surgery on $\bigsqcup_{\alpha} M_{\alpha}$, the 1-handles being attached to $\bigsqcup_{\alpha} M_{\alpha}^{\circ}$. A new link $L = \bigsqcup_{\alpha} L_{\alpha} \subset M$ is formed. The maps $\{f_{\alpha}\}$ are extended across the 1-handles as depicted in Figure 16 to give rise to a new Morse map $f : M^{\circ} \rightarrow S^1$. Then, $\chi_-(L, M) = \sum_{\alpha} \chi_-(F_{\alpha})$.

In particular, this additivity formula holds when all the L_{α} 's are fibered links with fibers F_{α} . \square

¹⁶A suitable approximation to the Thom-Pontryagin map $f_{F_{\alpha}} : M_{\alpha}^{\circ} \rightarrow S^1$ will do.

Example 8.16. We illustrate this construction in the context of structures induced by holomorphic functions $h : \mathbb{C}^2 \rightarrow \mathbb{C}$. Let $h(\mathbf{z}) := h(z_1, z_2)$ be a complex polynomial function with isolated critical points. Some of these points $\{\mathbf{z}_\alpha^*\}$ might reside on the complex curve $V = h^{-1}(0)$, where they manifest themselves as isolated singularities of V .

Consider the intersection $L_{r,\alpha}$ of the r -sphere $S_{r,\alpha}^3 \subset \mathbb{C}^2$ centered on \mathbf{z}_α^* with the real surface V . For a generic r , $L_{r,\alpha}$ is a smooth link in S_r^3 .

The function $f(\mathbf{z}) = h(\mathbf{z})/|h(\mathbf{z})|$ defines a smooth map $f : \mathbb{C}^2 \setminus V \rightarrow S^1$. In terms of the natural parameter θ along the circle, f is given by a (real multivalued) formula

$$\theta(\mathbf{z}) = -i \log(h(\mathbf{z})) + i \log(|h(\mathbf{z})|) = \text{Im}\{\log(h(\mathbf{z}))\}.$$

The map f induces maps $f_{r,\alpha} : S_{r,\alpha}^3 \setminus L_{r,\alpha} \rightarrow S^1$. According to Milnor [M1], for small r 's, the $f_{r,\alpha}$ are fibrations with a fiber $F_{r,\alpha}$.

In $\mathbb{C}^2 \setminus V$ one can link each pair $(S_{r,\alpha}^3, S_{r,\alpha+1}^3)$ by a path γ_α which is mapped by f in a monotone fashion into S^1 . We can pick all $\{\gamma_\alpha\}$ in such a way that they do not intersect each other and do not penetrate into the interiors of the spheres. Attaching to the spheres thin 1-handles $\{D^3 \times \gamma_\alpha\}$ with the γ_α 's as cores, we produce a new sphere S^3 with a link $L = V \cap S^3 = \coprod_\alpha L_{r,\alpha}$ and a map $f : (S^3)^\circ \rightarrow S^1$ of the sort described in Corollary 8.15. \square

9. ON TANGENCIES OF SURFACES TO THE FIBERS

In this section we are going to link the singularities of a Morse map $f : M \rightarrow S^1$ with the singularities of its restriction $f| : \Sigma \rightarrow S^1$ to a probe surface Σ residing in the complement to the f -singularities and realizing a particular 2-homology class of M .

Denote by \mathcal{F}_f the singular foliation in M defined by the fibers of a Morse map $f : M \rightarrow S^1$ as in Theorem 8.2. The map $f : \partial M \rightarrow S^1$ is assumed to be a fibration over the circle. Denote by M° the complement in M to the f -critical points, and by M^\bullet the compactification of M° by the 2-spheres $\{S_\alpha^2\}$ “surrounding” the singularities $\{x_\alpha\}$ of \mathcal{F}_f . Let $\mathcal{S}^2 := \bigsqcup_\alpha S_\alpha^2$.

The 2-plane field formed by the tangent planes to \mathcal{F}_f defines an oriented 2-bundle ξ_f over M^\bullet (the orientation being induced by the orientations of M and S^1). Since $f : \partial M \rightarrow S^1$ is a fibration, the bundle $\xi_f|_{\partial M}$ is trivial with a preferred trivialization defined by the intersections of the 2-planes with the boundary ∂M . Its Euler class $\chi(\xi_f) \in H^2(M^\bullet; \mathbb{Z})$ is Poincaré dual in M^\bullet to a 1-cycle $\beta_f \in H_1(M^\bullet, \mathcal{S}^2; \mathbb{Z})$. This 1-cycle can be viewed as the zero set of a generic section σ of the bundle ξ_f , subject to a number of boundary conditions over ∂M and \mathcal{S}^2 . Specifically, we require the restriction $\sigma|_{\partial M}$ to be a non-vanishing section of $\xi_f|_{\partial M}$ pointing in the direction of the fibers of $f : \partial M \rightarrow S^1$. Also, we assume $\sigma|_{S_\alpha^2}$ to be a generic section of $\xi_f|_{S_\alpha^2}$ formed by vectors tangent to the trace of \mathcal{F}_f on S_α^2 .

For any *immersed* surface $i : (\Sigma, \partial\Sigma) \rightarrow (M^\circ, \partial M)$, the integer $\langle \chi(\xi_f|_\Sigma), [\Sigma] \rangle$ is equal to the algebraic intersection $\Sigma \circ \beta_f$. This number is an invariant of the homology class of $i_*[\Sigma] \in H_2(M^\bullet, \partial M; \mathbb{Z})$.

Isolated points, where Σ is tangent to the non-singular foliation \mathcal{F}_f on M° , come in two flavors: *positive*, when the preferred orientations of Σ and \mathcal{F}_f agree, and *negative*, when they disagree. By intersecting the surface with the fibers, the non-singular foliation \mathcal{F} induces a singular oriented foliation \mathcal{F}_Σ on Σ . Isolated

points of tangency also can be of *elliptic* and *hyperbolic* types. For the elliptic type, the index of the vector field X_Σ on Σ , normal to \mathcal{F}_Σ , is positive, and for the hyperbolic ones, it is negative.

We denote by h_+ and h_- the number of positive and negative hyperbolic points. Similarly, let e_+ and e_- stand for the number of positive and negative elliptic tangencies.

Let \mathcal{I}_+ denote the sum of indices of the vector field X_Σ at all positive tangent points, and let \mathcal{I}_- denote the sum of indices of X_Σ at all negative tangent points. When all the tangencies are of the Morse type, then

$$(9.1) \quad \mathcal{I}_+ = e_+ - h_+, \quad \mathcal{I}_- = e_- - h_-.$$

As in [T], one can prove that the Euler characteristic of Σ and its intersection with the relative 1-cycle β_f can be calculated in terms of the tangencies:

$$(9.2) \quad \chi(\Sigma) = \mathcal{I}_+ + \mathcal{I}_-, \quad \langle \chi(\xi_f|_\Sigma), [\Sigma] \rangle = \mathcal{I}_+ - \mathcal{I}_-.$$

If Σ has no spherical or disk components, then $\chi_-(\Sigma) = |\mathcal{I}_+ + \mathcal{I}_-|$.

Consider an immersed surface $(\Sigma, \partial\Sigma) \propto (M^\bullet, \partial M)$ which is homologous in $(M^\bullet, \partial M)$ to a combination F_{best} of fiber components which delivers $\|[\Sigma]\|_{Hf}$. For such a Σ , $\langle \chi(\xi_f|_\Sigma), [\Sigma] \rangle = \langle \chi(\xi_f|_{F_{best}}), [F_{best}] \rangle$.

At the same time, $\langle \chi(\xi_f|_{F_{best}}), [F_{best}] \rangle = \chi(F_{best})$ — the bundle ξ_f is formed by the tangent planes to the fibers. Therefore, we have

Lemma 9.1. *If an immersed surface Σ is homologous to F_{best} in $(M^\bullet, \partial M)$ and both Σ and F_{best} have no spherical and disk components, then $\chi_-(\Sigma) = |\mathcal{I}_+ + \mathcal{I}_-|$ and $\chi_-(F_{best}) = |\mathcal{I}_+ - \mathcal{I}_-|$. As a result, $\chi_-(\Sigma) \geq \chi_-(F_{best})$ if and only if $\mathcal{I}_+ \cdot \mathcal{I}_- \geq 0$. \square*

In combination with Theorem 8.6, this leads to our next theorem, which makes it possible to estimate the twist $\rho^\circ(\Sigma, F_R)$ of a probe surface Σ in terms of the Morse data of f and of its restriction to Σ .

We do not know if any probe surface $\Sigma \subset M$ homologous to F_{best} in M can be replaced by an *embedded* surface $\Sigma' \subset M^\bullet$ with $\chi_-(\Sigma') \leq \chi_-(\Sigma)$ and homologous to F_{best} in M^\bullet . It is easy to verify that such a replacement Σ' exists among immersed surfaces.

Theorem 9.2. *Let $f : M \rightarrow S^1$ be a map as in Theorem 8.2. Let a surface $\Sigma \subset M^\bullet$ be well-positioned with respect to F_R and homologous in $(M^\bullet, \partial M)$ to a best combination F_{best} of fiber components (i.e. $\chi_-(F_{best}) = \|[F_{best}]\|_{Hf}$). Assume that both Σ and F_{best} have no spherical and disk components. Then*

$$Var_{\chi_-}(f) \cdot \rho^\circ(\Sigma, F_R) \geq |\mathcal{I}_+ - \mathcal{I}_-| - |\mathcal{I}_+ + \mathcal{I}_-|.$$

Furthermore,

$$\frac{1}{2} Var_{\chi_-}(f) \cdot \rho^\circ(\Sigma, F_R) \geq \mathcal{I}_- \quad \text{and} \quad \mathcal{I}_+ \leq 0.$$

For generic tangencies, the last two inequalities can be re-written as

$$(9.3) \quad \frac{1}{2} Var_{\chi_-}(f) \cdot \rho^\circ(\Sigma, F_R) \geq e_- - h_- \quad \text{and} \quad e_+ \leq h_+.$$

The first statement is non-trivial only when the signs of \mathcal{I}_+ and \mathcal{I}_- are *opposite*. The last two inequalities imply that:

- a surface Σ , as above, with many negative elliptic tangent points and few negative hyperbolic ones must have a sizable twist;

- the number of positive elliptic tangent points does not exceed the number of positive hyperbolic ones.




FIGURE 17.

Proof. Employing that $\mathcal{I}_+ - \mathcal{I}_- = \langle \chi(\xi_f|_\Sigma), [\Sigma] \rangle = \chi(F_{best}) \leq 0$, and that $\mathcal{I}_+ + \mathcal{I}_- = \chi(\Sigma) \leq 0$, all the statements follow from the diagram in Figure 17. The grey area in the diagram depicts the solution set for the inequality $|\mathcal{I}_+ - \mathcal{I}_-| - |\mathcal{I}_+ + \mathcal{I}_-| \leq 2a$, where $2a = \text{Var}_{\chi_-}(f) \cdot \rho^\circ(\Sigma, F_R)$. The solutions of all the three inequalities above is the domain shaded with dark grey. The bold diagonal lines represent surfaces Σ with a fixed negative value of $\langle \chi(\xi_f|_\Sigma), [\Sigma] \rangle$. \square

We conclude this section with a few remarks about minimal surfaces in the homology class of a fiber. We observe that the singularities of the map f can act as *attractors* for families of minimal surfaces. Below we describe such a behavior in general terms. However, our grasp of this interesting phenomenon is poor.

The basic fact is that any two minimal (connected) surfaces have only isolated tangencies of *hyperbolic* type, unless they are identical. Although such tangencies are not necessarily of the Morse type, they still are canonical: their smooth type is modeled after the tangency at the origin of the surface $\{t = \text{Re}(z^n)\}$ and the surface $\{t = 0\}$ in the 3-dimensional space $\mathbb{R} \times \mathbb{C}$ (cf. Lemma 1.4 in [FHS]).

In [K], for a given intrinsically harmonic 1-form ω , and thus for any harmonic map into the circle, we have constructed (two-parametric) families of riemannian metrics g_μ on M^d with the following properties:

- ω is harmonic with respect to g_μ ,
- outside the disks $\{D_\mu(x_\alpha)\}_\alpha$ of radius μ surrounding the singularities $\{x_\alpha\}_\alpha$ of ω , the foliation \mathcal{F}_ω is comprised of minimal hypersurfaces,
- the deviation of the leaves inside the disks from minimality is $\sim \mu^{d-1}$ -small.

Therefore, for a given intrinsically harmonic map $f : M^3 \rightarrow S^1$ and any collection of *non-singular* fibers $\{F\}$ whose closure does not contain the f -singularities, there is a family of metrics g_μ on M^3 such that f is harmonic and the surfaces F are minimal. Furthermore, the argument in [K] shows that the fibers of $f : \partial M^3 \rightarrow S^1$

are geodesic loops. The boundary ∂M is *sufficiently convex* in the sense of Meeks and Yau [MY] (in fact, it is flat).

According to Theorem 5.1 in [FHS], for a large class of 3-manifolds M (so-called P^2 -irreducible ones), any two-sided *incompressible* oriented embedded surface $\alpha : \Sigma' \subset M$, distinct from S^2 , can be isotoped to 1) a *minimal embedded* surface $\Sigma \subset M$, or to 2) a surface $\Sigma \subset M$ which is a boundary of a regular neighborhood of a *one-sided minimal* embedded surface.

In the first case, the minimal surface minimizes the area in the homotopy class of α . In the second case, the minimal surface realizes *half* of the minimal area in the homotopy class of α .

Combining this result with the special properties of the metrics $\{g_\mu\}$, we conclude that any embedded incompressible surface in the homology class of a fiber has a minimal or a “near-minimal” representative $\Sigma \subset M$, having only *hyperbolic* tangencies with the f -fibers *outside* of the disks $D_\mu(x_\alpha)$. In other words, all the elliptic tangent points must be located *inside* the disks $\{D_\mu(x_\alpha)\}$. As $\mu \rightarrow 0$, they are attracted towards the singularities of f .

10. NOTATION LIST

- $var_{\chi_-}(f)$ cf. formulas (3.1) and (4.1)
- $Var_{\chi_-}(f)$ cf. formula (4.2)
- $\rho(\Sigma, F)$ cf. the 6th and 7th paragraphs in Section 6
- $\rho^\circ(\Sigma, F)$ cf. the 2nd paragraph before Definition 6.1
- $\rho_{\chi_-}(\Sigma, F)$ cf. Definition 6.1
- $\rho_{\chi_-}([\Sigma], F_R)$ cf. Definition 6.3
- $\rho_{\chi_-}(f) := \rho_{\chi_-}([F], F_R)$, F being a fiber
- $\rho(\hat{\Sigma}, \tilde{M})$ cf. formula (6.1)
- $h(\Sigma, f)$ cf. formula (6.2)
- $b(F, \Sigma)$ cf. Definition 6.7
- $\mu^\circ(\Sigma, F)$ cf. formula (7.1)
- $b_{\chi_-}(F_R, k[F])$ cf. the paragraph before Corollary 6.9
- $h_{\chi_-}(k[F], f)$ cf. Definition 6.11
- $h_{\chi_-}(f) := h_{\chi_-}([F], f)$, F being a fiber
- $\mu^\circ([\Sigma], F_R)$ cf. the 1st paragraph after Corollary 8.7
- $\|[\Sigma]\|_T$ cf. the 3rd paragraph in the Introduction
- $\|[\Sigma]\|_{H^f}$ cf. formula (8.2) and Definition 8.1
- $\|\kappa\|_A$ cf. formula (8.1)

ACKNOWLEDGMENT

This work was shaped by numerous, thought-provoking conversations with Jerry Levine. I am thankful for his generosity and support. I am also grateful to the referee, whose suggestions significantly improved the clarity of the original presentation.

REFERENCES

- [C] Calabi, E., *An Intrinsic Characterization of Harmonic 1-Forms*, in *Global Analysis, Papers in Honor of K. Kodaira*, D.C. Spencer and S. Ianga, Eds., (1969), pp. 101-117. MR40:6585
- [Cr] Crowell R. H., *Genus of alternating link types*, *Annals of Math.*, 69 (1959), pp. 258-275. MR20:6103b

- [FKL] Farber, M., Katz, G., and Levine J., *Morse Theory of Harmonic Forms*, Topology, vol. 37, No. 3 (1998), pp. 469-483. MR99i:58026
- [FHS] Freedman, M., Hass, H., Scott, P., *Least Area Incompressible Surfaces in 3-Manifolds*, Inventiones Mathematicae, No. 71 (1983), pp. 609-642. MR85e:57012
- [G] Gabai, D., *Foliations and the Topology of 3-Manifolds*, J. Differential Geometry, No. 18 (1983), pp. 445-503. MR86a:57009
- [H] Hass, J., *Surfaces Minimizing Area in Their Homology Class and Group Actions on 3-Manifolds*, Math. Zeitschrift, Vol. 199, (1988), pp. 501-509. MR90d:57017
- [HS] Hass, J., Scott, P., *The Existence of Least Area Surfaces in 3-Manifolds*, Trans. of AMS, vol. 310, No. 1, (1988), pp. 87-114. MR90c:53022
- [Ho] Honda, K., *A Note on Morse Theory of Harmonic 1-Forms*, Topology No 38 (1) (1999), pp. 223-233. MR99h:58030
- [K] Katz, G., *Harmonic Forms and Near-minimal Singular Foliations*, Comment. Math. Helvetici 77 (2002), 39-77. MR2003b:57042
- [Mc] McMullen, C. T., *The Alexander polynomial of a 3-manifold and the Thurston norm on cohomology*, Ann. Sci. École Norm. Sup. (4) 35 (2002), 153-171. MR2003d:57044
- [MY] Meeks, W.H., Yau, S.T., *The Existence of Embedded Minimal Surfaces and the Problem of Uniqueness*, Math. Z., No 179 (1982), pp. 151-168. MR83j:53060
- [M] Milnor, J., *Lectures on the h-cobordism Theorem*, Princeton University Press, 1965. MR32:8352
- [M1] Milnor, J., *Singular points of complex hypersurfaces*, Princeton University Press, 1968. MR39:969
- [Mur] Murasugi K., *On the genus of the alternating knot, I, II*, J. Math. Soc. Japan 10 (1958), pp. 94-105, 235-248. MR20:6103a
- [Su] Sullivan, D., *A Cohomological Characterization of Foliations Consisting of Minimal Surfaces*, Comment. Math. Helvetici, No.54 (1979), pp. 218-223. MR80m:57022
- [T] Thurston, W.P., *A Norm for the Homology of 3-Manifolds*, Memoirs of AMS, Vol. 59, No. 339 (1986), pp.100-130. MR88h:57014
- [W] Wall, C.T.C., *Surgery on Compact Manifolds*, Academic Press, London & New York, 1970. MR55:4217

DEPARTMENT OF MATHEMATICS, BENNINGTON COLLEGE, BENNINGTON, VERMONT 05201-6001

E-mail address: gabrielkatz@rcn.com

E-mail address: gkatz@bennington.edu

Current address: Department of Mathematics, Brandeis University, Waltham, Massachusetts 02454