Minimal genus in $S^1 \times M^3$

P. B. Kronheimer[1]

Harvard University, Cambridge MA 02138

1. Introduction

1.1 Statement of the result

Given a smooth four-manifold X and a class σ in $H_2(X;\mathbb{Z})$, one can ask what is the least possible genus for a smoothly embedded, oriented surface Σ in X whose fundamental class is σ. Although the issue does not arise in the most familiar examples, one should allow Σ to be disconnected, in which case it is more natural to try and minimize the quantity

$$\chi_-(\Sigma) = \sum_{g_i > 0} 2g_i - 2,$$

where g_i is the genus of the component Σ_i (*cf.* [16]). Our purpose here is to address this question in the case that X is a product, $S^1 \times M^3$.

Four-manifolds such as $S^1 \times M^3$ seem rather different from the usual run of Kähler surfaces, which form the building bricks of the examples that are most often studied when gauge theory is a tool. In particular, the established techniques for attacking the minimal genus problem (primarily the 'adjunction inequality' involving the basic classes) are ineffective in these cases, because the Seiberg-Witten invariants of such products may contain little or no information. (It is easy to find examples where the invariants are identically zero.) Nevertheless, by using the gauge theory tools rather less directly, we shall obtain sharp results for a large proportion of these manifolds. The following is our main result.

Theorem 1. *Let M be a closed, irreducible, oriented three-manifold carrying a smooth, taut foliation $\mathcal{F}$ by oriented two-dimensional leaves. On the four-manifold $S^1 \times M$, let ϵ be the pull-back of the euler class $e(T\mathcal{F})$. Then for any smoothly embedded oriented surface Σ in $S^1 \times M$ representing a homology class σ, we have*

$$\chi_-(\Sigma) \geq \sigma \cdot \sigma + \epsilon \cdot \sigma.$$

[1] Partially supported by NSF grant number DMS-9531964

Some remarks may clarify the nature of this theorem. First, the hypothesis of irreducibility (every two-sphere bounds a ball) is there only to exclude $S^1 \times S^2$. Other reducible three-manifolds cannot carry taut foliations. Second, if the surface Σ is contained in the three-manifold $\{\mathrm{pt}\} \times M$, the inequality $\chi_-(\Sigma) \geq \epsilon \cdot \sigma$ is proved by Thurston in [16]. More generally, if Σ lies in $S^1 \times M$ but is homologous to a surface in M, then the self-intersection term is zero, and the inequality of the theorem follows from a result of Gabai [5], which states that a singular surface, such as the one obtained by projecting Σ into M, is always homologous to an embedded surface of the same genus. Finally, if it happens that ϵ is a basic class for the four-manifold in the sense of [9] or [13], then the inequality of the theorem coincides with the adjunction inequality. (A simple proof of the adjunction inequality is contained in section 2.1.) The interesting aspect of the theorem is precisely that ϵ is *not* in general a basic class.

1.2 An example

The strength of Theorem 1 is due to Gabai's theorem [5] on the existence of taut foliations, as a consequence of which one can often obtain complete information. Consider a general homology class σ in $H_2(S^1 \times M)$ written in the form

$$\sigma = [S^1] \times \gamma + \sigma_0,$$

where σ_0 is a class arising from the inclusion of $H_2(M)$ and γ is an element of $H_1(M)$. The self-intersection number $\sigma \cdot \sigma$ in the four-manifold is twice the intersection number $k = \gamma \cdot \sigma_0$ in the three-manifold, which we may assume to be non-negative. Suppose that the classes σ_0 and γ can be represented by a surface Σ_0 and curve Γ with the properties:

(a) $\chi_-(\Sigma_0)$ is minimal amongst representatives of its homology class in M, and no component of Σ_0 is a sphere;

(b) the geometric intersection of Γ and Σ_0 is transverse and its cardinality coincides with the algebraic intersection number, k.

(These conditions can always be met if $b_1(M) = 1$.) The class σ in the four-manifold can then be represented by the singular surface

$$\Sigma' = (S^1 \times \Gamma) \cup \Sigma_0, \tag{1}$$

which has k double points. At each double point we may replace a pair of transverse disks with an embedded annulus having the same oriented boundary, to obtain a smooth surface Σ representing the class. This surface has

$$\begin{aligned}\chi_-(\Sigma) &= 2k + \chi_-(\Sigma_0)\\ &= \sigma\cdot\sigma + \chi_-(\Sigma_0).\end{aligned}$$

If M is irreducible, then by the main theorem of [5], condition (a) implies that there exists a taut foliation $\mathcal{F}$ on M whose euler class ϵ satisfies

$$\epsilon\cdot\sigma_0 = \chi_-(\Sigma_0).$$

If Σ_0 is a union of tori, then it may be that $\mathcal{F}$ is only C^0. In all other cases, [5] provides a smooth foliation. When $\mathcal{F}$ is smooth, Theorem 1 tells us that the surface Σ that we have described achieves the smallest possible value for χ_- amongst representatives of the class σ. To summarize, we have the following corollary of Theorem 1 and [5].

Corollary 2. *Let Σ_0 and Γ be a surface and a curve in an irreducible three-manifold M, satisfying (a) and (b). Suppose that not all components of Σ_0 are tori. Then the surface Σ obtained by smoothing* (1) *achieves the smallest possible value of χ_- in its homology class.* □

An interesting case is a three-manifold M obtained by zero-surgery on a knot:

Corollary 3. *Let M be obtained by zero-surgery on a knot K whose genus g is at least two. Let $\Sigma \subset S^1 \times M$ be a smoothly embedded surface representing the homology class*

$$\sigma = m([S^1]\times\delta) + n\tau \tag{2}$$

where δ and τ are generators for $H_1(M)$ and $H_2(M)$, and m, n are non-negative integers. Then any surface Σ representing σ has

$$\chi_-(\Sigma) \geq 2mn + (2g-2)n.$$

This inequality is sharp.

Proof. There is no difficulty in achieving (a) and (b). By the results of [5], the genus-minimizing surface representing the class τ in M has the same genus as the knot. Taking n parallel copies of this surface gives a representative for $n\tau$ with $\chi_- = (2g-2)n$, and this is least possible. The corollary therefore follows from the previous one. □

It is worth noting that Gabai's foliation is, in practice, smooth even for many genus-one knots, including all those with fewer than eleven crossings [6]. The above corollary can therefore be applied also to these knots. It seems most likely that it can in fact be applied to *all* knots.

For zero-surgery on K, the basic classes[1] of $S^1 \times M$ are various multiples of the Poincaré dual of $[S^1] \times \delta$, and the largest multiple which occurs is $2r - 2$, where r is the degree of the symmetrized Alexander polynomial A_K, i.e. the largest r for which the coefficient of $(t^r + t^{-r})$ in $A_K(t)$ is non-zero [12]. The adjunction inequality therefore gives only the lower bound

$$\chi_-(\Sigma) \geq 2mn + (2r-2)n$$

in the situation of Corollary 3. The genus of a knot is an upper bound for the degree r of the Alexander polynomial, and the two are equal if K is a fibered knot (as well as for all knots with fewer than eleven crossings, by [6] again). But by taking a connected sum of Whitehead doubles, one readily constructs knots with trivial Alexander polynomial and arbitrary genus.

1.3 Symplectic structures

The discrepancy between the lower bound provided by the adjunction inequality and the lower bound provided by our theorem leads to a non-existence result for symplectic structures. Let M again be obtained by zero-surgery on a knot K. It is pointed out in [4] that a necessary condition for $S^1 \times M$ to be symplectic is that the symmetrized Alexander polynomial $A_K(t)$ is *monic*; that is, the coefficient of the leading term $(t^r + t^{-r})$ should be ± 1. We can add the following constraint:

Proposition 4. *Let M be obtained by zero-surgery on a knot K of genus two or more. A necessary condition for $S^1 \times M$ to admit a symplectic structure is that its genus be equal to the degree of its symmetrized Alexander polynomial. This result also holds for knots of genus one, as long as the foliation whose existence is given by [5] can be made smooth.*

Proof. Suppose that $S^1 \times M$ has a symplectic form ω. We may assume that ω is Poincaré dual to an integer homology class of the form (2) with n strictly positive. By a theorem of Donaldson [2], after passing to a multiple of this class if necessary, we are assured of the existence of a symplectic submanifold H, Poincaré dual to ω. The genus of H is determined by the adjunction formula, so

$$\chi_-(H) = \sigma \cdot \sigma - c_1 \cdot \sigma,$$

where c_1 is the first Chern class of the symplectic structure. From the results of [14, 15], and [12], one knows that the first Chern class is dual to

$$-c_1 = (2r-2)([S^1] \times \delta)$$

[1] The four-manifold we consider has $b^+ = 1$, so we must be a little careful in defining the basic classes.

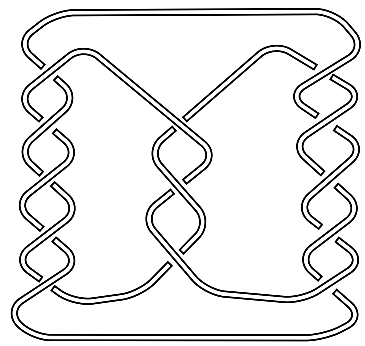

Figure 1: The pretzel knot $(5, -3, 5)$.

where r is the degree of the symmetrized Alexander polynomial. We therefore have

$$\chi_-(H) \leq 2mn + (2r - 2)n.$$

Since n is positive, this equality is inconsistent with the lower bound contained in Corollary 3 if r is smaller than the genus of the knot. □

If K is fibered, then $S^1 \times M$ is symplectic. (This is essentially an observation of Thurston's.) The proposition above, and the result already mentioned from [4], say that if $S^1 \times M$ is symplectic, then the knot looks fibered as far as its Alexander polynomial and genus are concerned. However, there are non-fibered knots whose Alexander polynomial is monic with degree equal to the genus of the knot. An infinite family of pretzel knots with these properties is exhibited in [1]. The first of these is the pretzel knot $(5, -3, 5)$ shown in Figure 1. It would be interesting to know whether the corresponding four-manifold $S^1 \times M$ has a symplectic structure. Any such symplectic structure would have $c_1 = 0$.

2. Proof of the theorem

2.1 Scalar curvature and monopole classes

As a first step to proving the Theorem 1, we reformulate the proof of the adjunction inequality for embedded surfaces. We shall do this by following the approach of [11]. The result is a proof which seems more elementary than the line taken in [8], and is certainly simpler than the proof contained in [13].

Let X be a closed manifold of any dimension, and let $\alpha \in H^2(X; \mathbb{R})$ be a 2-dimensional cohomology class. For each Riemannian metric h on X, one can define a norm $\|\alpha\|_h$ using the L^2 norm of the harmonic representative. Consider now the quantity

$$|\alpha| = 4\pi \sup_h \big(\|\alpha\|_h / \|s_h\|_h\big), \tag{3}$$

where s_h is the scalar curvature, and the supremum is taken over all Riemannian metrics. The supremum may be infinite, but with this understood, it is clear that the above construction defines a norm on the linear subspace on which it is finite. This quantity was consider in [11] for the case that X is three-dimensional. In the case of an oriented four-manifold, one can vary this construction by decomposing the harmonic representative for a metric h into its self-dual and anti-self-dual parts, $\alpha^+ + \alpha^-$, and defining

$$|\alpha|^+ = 4\pi\sqrt{2} \sup_h \big(\|\alpha^+\|_h / \|s_h\|_h\big).$$

Note that in general one has

$$\begin{aligned} \|\alpha^+\|_h^2 + \|\alpha^-\|_h^2 &= \|\alpha\|_h^2 \\ \|\alpha^+\|_h^2 - \|\alpha^-\|_h^2 &= (\alpha \smile \alpha)[X], \end{aligned}$$

so if $\alpha \smile \alpha = 0$ then $|\alpha| = |\alpha|^+$. This is the reason for the additional factor of $\sqrt{2}$ in the second definition.

On an oriented four-manifold X equipped with a Spin^c structure $\mathfrak{c}$ and Riemannian metric h, the Seiberg-Witten monopole equations [17, 10] are the following pair of equations for a section Φ of the associated spin bundle $W^+ = W_{\mathfrak{c}}^+$ and a spin connection A:

$$\begin{aligned} \rho(F_{\hat{A}}^+) - \{\Phi \otimes \Phi^*\} &= 0 \\ D_A \Phi &= 0. \end{aligned} \tag{4}$$

Here D_A is the Dirac operator, ρ is Clifford multiplication, $F_{\hat{A}}$ denotes the curvature of the induced connection in the line bundle $\Lambda^2 W^+$, and the curly brackets denote the trace-free part of the endomorphism. Our conventions here follow [10]. We write $c_1(\mathfrak{c})$ for the first Chern class of W^+.

A straightforward calculation, contained in [17], is the basis of many of the analytic properties of these equations: an application of the Weitzenbock formula for $D_A^* D_A$ and an integration by parts leads to the inequality

$$2\sqrt{2}\|F_{\hat{A}}^+\|_h \le \|s\|_h \tag{5}$$

for any solution, and hence the inequality

$$4\pi\sqrt{2}\|\alpha^+\|_h \le \|s\|_h$$

for the class $\alpha = c_1(\mathfrak{c})$ (reduced to real coefficients). The following definition and lemma merely rewrites this result.

Definition 5. A class $\alpha \in H^2(X;\mathbb{R})$ is a *monopole class* if it arises as $c_1(\mathfrak{c})$ for some Spin^c structure $\mathfrak{c}$ with the property that the corresponding equations (4) have a solution for *all* Riemannian metrics h on X.

Lemma 6. *If $\alpha \in H^2(X;\mathbb{R})$ is a monopole class on a closed, oriented four-manifold X, then $|\alpha|^+ \leq 1$.* □

The next lemma, like the one above, was proved in a three-dimensional version in [11]. Also like the previous lemma, it is entirely elementary.

Lemma 7. *If X is as in the previous lemma and $\alpha \in H^2(X;\mathbb{R})$ is a class with $|\alpha|^+ \leq 1$, then*

$$\chi_-(\Sigma) \geq \alpha \cdot \sigma$$

for every embedded surface $\Sigma \subset X$ representing a class σ with $\sigma \cdot \sigma = 0$.

Proof. It is enough to consider the case that Σ is connected and is not a sphere. Because Σ has trivial normal bundle, the boundary of the tubular neighborhood of Σ is $S^1 \times \Sigma$, and we can find a metric h_1 on X which contains an isometric copy of $[0,1] \times S^1 \times \Sigma$ carrying a product metric in which S^1 has unit length, and Σ has unit area and constant scalar curvature. Let h be a metric obtained from this one by replacing the factor $[0,1]$ by $[0,r]$. This metric on X contains a long cylindrical piece when r is large. We have

$$\|s_h\|_h = 4\pi r^{1/2}(2g(\Sigma) - 2) + O(1)$$

as $r \to \infty$, while any 2-form ω representing the class α must satisfy

$$\|\omega\|_h \geq r^{1/2}(\alpha \cdot \sigma).$$

The last inequality gives

$$2\|\omega^+\|_h^2 \geq r(\alpha \cdot \sigma)^2 + (\alpha \smile \alpha)[X],$$

and hence

$$\sqrt{2}\|\omega^+\|_h \geq r^{1/2}(\alpha \cdot \sigma) + O(1).$$

The hypothesis $|\alpha|^+ \leq 1$ simply means that

$$\sqrt{2}\|\omega^+\|_h \leq (1/4\pi)\|s_h\|_h$$

for all h, and the inequality $2g - 2 \geq \alpha \cdot \sigma$ is a consequence of the estimates we have made. □

Combining the two lemmas above, we deduce that if α is a monopole class and Σ is an embedded surface representing a class σ with zero square, then

$$\chi_{-}(\Sigma) \geq \alpha \cdot \sigma.$$

This is a raw form of the adjunction inequality. It is usually made more useful as follows. First, let $X_k = X \# k\bar{\mathbb{CP}}^2$ be the connected sum of X and k copies of $\mathbb{CP}^2$ with reversed orientation. Let

$$\alpha_k = \alpha + e_1 + \cdots + e_k, \tag{6}$$

where e_i is the generator of H^2 in the ith copy of $\bar{\mathbb{CP}}^2$. If α_k is a monopole class on X_k for all $k \gg 0$, then one has the inequality

$$\chi_{-}(\Sigma) \geq \sigma \cdot \sigma + \alpha \cdot \sigma \tag{7}$$

for any embedded surface Σ in X representing a class σ of non-negative square. This statement can be reduced to the raw version by passing from Σ to the surface $\tilde{\Sigma}$ with trivial normal bundle obtained by forming an internal connected sum of Σ with a suitable number of spheres dual to the generators of the summands $\bar{\mathbb{CP}}^2$ in X_k. Finally, if $b^+(X) > 1$, then a class α is a monopole class if it is a basic class; that is, if it is $c_1(\mathfrak{c})$ for a Spin^c structure $\mathfrak{c}$ for which the moduli space of solutions to the equations (4), after perturbation, represents a non-trivial homology class in the configuration space. The blow-up formula ([7], Proposition 2) implies that if α is a basic class on X, then α_k is a basic class on X_k. The inequality (7) therefore holds for basic classes on X.

2.2 Using finite covers

We now focus on the case that X is $S^1 \times M$, for some closed oriented three-manifold M. We continue to write X_k for the connected sum $X \# k\bar{\mathbb{CP}}^2$, and for a given class $\alpha \in H^2(X; \mathbb{R})$, we write α_k for the class (6). On account of Lemma 7 and the remarks above, the following proposition implies Theorem 1.

Proposition 8. *Let $\mathcal{F}$ be an oriented, taut foliation of M, let ϵ be the pull-back of $e(T\mathcal{F})$ to $X = S^1 \times M$, and let ϵ_k denote the class*

$$\epsilon_k = \epsilon + e_1 + \cdots + e_k,$$

on $X_k = X \# k\bar{\mathbb{CP}}^2$. Then $|\epsilon_k|^+ \leq 1$, for all $k \geq 0$.

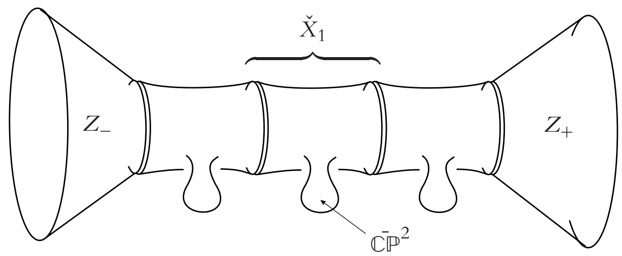

Figure 2: The geometry of Z_1^3

Proof. Let any metric h_1 on X_k be given. We must show that

$$4\pi\sqrt{2}\|\epsilon_k^+\|_{h_1} \leq \|s_{h_1}\|_{h_1}.$$

Let X_k^N be the N-fold cyclic cover of X_k arising from the S^1 factor in X. (This manifold is diffeomorphic to X_{Nk}.) Let $\check{X}_k$ be the manifold with two boundary components obtained by cutting X_k along a standard copy of M, and let $\check{X}_k^N$ be obtained similarly from X_k^N, so that $\check{X}_k^N$ is the concatenation of N copies of $\check{X}_k$. The metric h_1 gives rise to a metric h_N on X_k^N and $\check{X}_k^N$.

We now follow the line of [10] (see also [11]). According to [3], there exist contact structures given by two-plane fields ξ_+, ξ_- on M, which approximate the tangent distribution to $\mathscr{F}$ and which are compatible with the two different orientations of M. These contact structures give rise to symplectic structures on the cone

$$Z_+ = [1, \infty) \times M$$

and its opposite, Z_-. Let Z_0 be the manifold (diffeomorphic to $\mathbb{R} \times M$) formed by attaching three pieces,

$$Z_0 = Z_- \cup ([-1, 1] \times M) \cup Z_+.$$

As in [10], the symplectic structures on the cones can be combined with a symplectic structure on the central piece to form a symplectic structure Ω on Z_0. A compatible metric gives Z_0 a geometry in which both ends are expanding cones. In the terminology of [10], this geometry on the ends of Z_0 gives Z_0 an AFAK (asymptotically flat almost Kähler) structure.

For each N and k, we construct a new AFAK manifold Z_k^N by removing the central piece $[-1, 1] \times M$ from Z_0 and replacing it with $\check{X}_k^N$ (or equivalently with N copies of $\check{X}_k$). This manifold is equipped with a metric $h(Z_k^N)$, which coincides with the metric h_N on the piece $\check{X}_k^N$, and is equal to the asymptotically conical metric of the AFAK structure outside a compact subset of the two conical pieces Z_+ and Z_-. The geometry of Z_1^3 is illustrated in Figure 2.

The manifold Z_k^N carries a symplectic form Ω, which is asymptotic to the cone-like symplectic structure on the ends. On the ends of the manifold, Ω is compatible with the Riemannian metric $h(Z_k^N)$, but in a compact region of the interior (containing $\check{X}_k^N$) it is not, because the metric h_1 which we chose at the start of this construction was arbitrary.

Following [10], we now introduce the Seiberg-Witten equations on Z_k^N, with a perturbing term on the ends of the manifold. The symplectic structure Ω determines a Spin^c structure $\mathfrak{c}$ on Z_k^N, whose first Chern class agrees with the pull-back of ϵ_k on the central piece, $\check{X}_k^N$. On the two ends of the manifold, the symplectic structure and compatible metric determine a preferred spin connection A_0 and a unit-length spinor Φ_0 satisfying the Dirac equation $D_{A_0}\Phi_0 = 0$. Extend these to the interior of the manifold, in such a way that they are periodic on the central piece $\check{X}_k^N$. Consider now the equations

$$\begin{gathered} \rho(F_{\hat{A}}^+) - \{\Phi \otimes \Phi^*\} = \rho(\pi) \\ D_A \Phi = 0 \end{gathered} \tag{8}$$

on Z_k^N, where π is an imaginary-valued self-dual two-form which vanishes on $\check{X}_k^N$ and is given by

$$\rho(\pi) = \rho(F_{\hat{A}_0}^+) - \{\Phi_0 \otimes \Phi_0^*\}$$

outside compact subsets of Z_+ and Z_-. The pair (A_0, Φ_0) is a solution of these equations on the ends of the manifold, but not in the interior. On the central piece, the equations coincide with the original Seiberg-Witten equations (4).

The main result of [10] says that, because Ω extends to a symplectic form on the whole manifold, the equations (8) have a solution (A_N, Φ_N) for the metric $h(Z_k^N)$. The solution is asymptotic to (A_0, Φ_0) on the ends. The usual C^0 estimate for solutions to the Seiberg-Witten equations (see Lemma 3.14 in [10], for example) gives a bound on $|\Phi_N|$; and this bound is independent of N, because the scalar curvature of $h(Z_k^N)$ and the size of π are both uniformly bounded, independent of N. The equations then yield uniform bounds on all gauge-invariant quantities derived from (A_N, Φ_N), such as the covariant derivatives of Φ_N and the curvature.

Lemma 9. *Let Y_+ and Y_- be collar neighborhoods of the two boundary components of the central piece $\check{X}_k^N$ in Z_k^N. We can find a gauge transformation u_N, so that after applying this gauge transformation to (A_N, Φ_N) we obtain a solution to the equations satisfying*

$$|A_N - A_0| \leq C$$

pointwise on Y, where C is a constant which is independent of N. In the same gauge, similar estimates hold for all covariant derivatives of $A_N - A_0$.

Proof of the lemma. Let Γ_+ be the family of generators of the cone Z_+. Extend these rays a little into the central piece so as to have a family of rays covering $Y_+ \cup Z_+$, with a similar family Γ_-. By applying a suitable gauge transformation to (A_N, Φ_N), we can arrange that it approaches (A_0, Φ_0) exponentially fast along the rays of $\Gamma_\pm$; more precisely, there is a compact set K containing the central piece, such that outside K we have

$$|A_N - A_0| \leq Ce^{\epsilon r} \tag{9}$$

where r is the distance along the rays, with similar estimates for the covariant derivatives of $A_N - A_0$. Furthermore, as shown in [10], the constants C and ϵ can be taken to be independent of N. (This is Proposition 3.15 of [10], with some care taken over the dependence of the constants on the geometry of the manifold.) Now use integration along the rays to define a gauge transformation ν_+ on $Z_+ \cup Y_+$ which is asymptotic to 1, so that after applying ν_+, the component of $A_N - A_0$ along the rays is zero. In this 'radial gauge', the bound (9) holds on the larger set $Z_+ \cup Y_+$, because of the pointwise bound we have on curvature. Define a similar gauge transformation ν_- on the other end. Because they are both asymptotic to 1, these two gauge transformations are homotopic to the identity, and there is therefore a gauge transformation ν defined on the whole manifold which agrees with these two on the ends. $\square$

After applying the gauge transformation which the lemma provides, we restrict (A_N, Φ_N) to the central piece $\check{X}_k^N$ and modify it using a cut-off function so that it is equal to (A_0, Φ_0) near the two boundary components. That is, we choose a cut-off function β on $\check{X}_k^N$ which is 0 on the two boundary components and 1 outside $Y_+ \cup Y_-$, and we set

$$\begin{aligned} A'_N &= A_N + (1 - \beta)(A_0 - A_N) \\ \Phi'_N &= \Phi_N + (1 - \beta)(\Phi_0 - \Phi_N) \end{aligned}$$

The modified configuration (A'_N, Φ'_N) gives rise to a configuration on the closed manifold X_k^N, after the two boundary components are identified, because the original configuration (A_0, Φ_0) is periodic on the central piece.

The configuration (A'_N, Φ'_N) on the closed manifold X_k^N will satisfy the (unmodified) Seiberg-Witten equations everywhere except on $Y_+ \cup Y_-$, which forms a two-sided collar of a copy of M in the closed manifold. Lemma 9 tells us that A'_N, Φ'_N and all their derivatives are bounded in that region, with constants that are independent of N. If we now apply the argument that leads to the inequality (5), the failure of (A'_N, Φ'_N) to be a solution leads to

$$2\sqrt{2}\|F_N^+\|_{h_N}^2 \le \|s_{h_N}\|_{h_N}\|F_N^+\|_{h_N} + C,$$

where C is a (new) constant which is independent of N. (We have used F_N to denote the curvature of A'_N.) From this we can obtain, for example,

$$2\sqrt{2}\|F_N^+\|_{h_N} \le \max\{\, \|s_{h_N}\|_{h_N} + 1,\ 2\sqrt{2}C\,\}.$$

The last inequality holds on X_k^N, on which the form $(i/2\pi)F_N$ represents the pull-back of the class ϵ_k. By averaging the form F_N over the action of the cyclic group of order N, we obtain an invariant form which descends to a form f the quotient space X_k, where it represents $(2\pi/i)\epsilon_k$. The inequality becomes

$$2\sqrt{2}\|f^+\|_{h_1} \le \max\{\, \|s_{h_1}\|_{h_1} + (1/N),\ 2\sqrt{2}C/N\,\},$$

on X_k. The harmonic representative of the class satisfies at least as strong a bound, so by taking the limit as N goes to infinity we obtain

$$4\pi\sqrt{2}\|\epsilon_k^+\|_{h_1} \le \|s_{h_1}\|_{h_1}.$$

This is the inequality we set out to establish. □

3. Further remarks

Reducible three-manifolds. For a closed three-manifold M, the dual Thurston norm can be defined on $H^2(M;\mathbb{R})$ by the formula

$$|\alpha|_T = \sup_\Sigma \big(\alpha \cdot [\Sigma]/\chi_-(\Sigma)\big),$$

the supremum being taken over all oriented embedded surfaces Σ in M. The norm of α is taken to be infinite if α has non-zero pairing with an embedded torus (or, a fortiori, with an embedded sphere). The corollary of the main theorem of [5] which we used in the proof of Corollary 2 can be rephrased as saying that, if M is irreducible, then the unit ball for this norm is the convex hull of the classes $e(T\mathcal{F})$, as $\mathcal{F}$ runs over all C^0 taut foliations. The same statement is true with *smooth* taut foliations, provided that M admits at least one such foliation, and this is the case if the dual Thurston norm is finite on at least one non-zero class.

With this understood, our main theorem can be rephrased as saying that, if α is a class with $|\alpha|_T = 1$, then for any embedded surface Σ in $S^1 \times M$ representing a homology class σ, we have

$$\chi_-(\Sigma) \ge \sigma \cdot \sigma + \alpha \cdot \sigma, \tag{10}$$

provided that M is irreducible and carries a smooth taut foliation. In this form, the result continues to hold for many reducible three-manifolds:

Theorem 10. *Let M be a connected sum of irreducible, closed oriented three-manifolds, in which each summand carries at least one smooth taut foliation. Let α be a class with dual Thurston norm 1 on M. Use the same letter to denote the pull-back of this class to $S^1 \times M$. Then for any embedded surface Σ in $S^1 \times M$ representing a homology class σ, we have*

$$\chi_-(\Sigma) \geq \sigma \cdot \sigma + \alpha \cdot \sigma. \tag{11}$$

Proof. The necessary modification to the argument is similar to the one that was used in [11]. □

As a particular example, let K_1 be the untwisted Whitehead double of the trefoil, let K_2 be the sum $K_1 \# K_1$, a genus-two knot, let M_2 be obtained by zero-surgery on K_2, and let M be the connected sum $M_2 \# M_2$. The four-manifold $X = S^1 \times M$ has $b^+ = 2$ and $b^1 = 3$. The Seiberg-Witten invariants of X are all zero, but our results imply that embedded surfaces in X satisfy the inequality (11), where α is the class from M which is equal to twice the generator on each of the two summands.

The norm $|\alpha|$ in general. The norm (3) can be defined on $H^2(X;\mathbb{R})$ for a manifold X of any dimension. It was shown in [11] that in dimension three, the norm coincides with the dual Thurston norm, at least for the class of three-manifolds M considered in Theorem 10. Proposition 8 and Lemma 7 completely determine this norm on the four-manifold $X = S^1 \times M$, for the same class of three-manifolds: the norm $|\alpha|$ is equal to the dual Thurston norm on classes pulled back from M, and is infinite on all other classes. Putting these two observations together, we can say that the pull-back map $H^2(M) \hookrightarrow H^2(S^1 \times M)$ is an isometry for these norms. It would be interesting to have a more direct verification of this.

In dimensions 5 and greater, the norm $|\alpha|$ is infinite for all non-zero classes α in H^2. This is not hard to verify, by an adaptation of the proof of Lemma 7.

Other four-manifolds. One motivation for considering the four-manifolds $S^1 \times M$ in connection with embedded surfaces came from [4], in which interesting examples of four-manifolds are constructed using such a product manifold as a building block. In particular, for each knot K, the authors of [4] construct a simply connected manifold X_K with the homotopy type of a $K3$ surface by removing the tubular neighborhood of a suitable torus T in a genuine $K3$ and replacing it with $S^1 \times M$, where M is the knot complement. The Seiberg-Witten invariants of X_K are shown to be determined by (and determine) the Alexander polynomial of K. In particular, if K has the same Alexander invariant as the unknot, then X_K cannot be distinguished from $K3$ by its Seiberg-Witten invariants, or by any other technique at present. Nevertheless, if K is knotted, it is natural to conjecture that X_K and $K3$ are not diffeomorphic.

More specifically, if g is the genus of K and τ is the class in X_K carried by a torus parallel to the torus T that was removed, then one might conjecture that an embedded surface Σ representing a class σ in $H_2(X_K)$ would satisfy the constraint

$$\chi_-(\Sigma) \geq \sigma \cdot \sigma + 2g\tau \cdot \sigma. \tag{12}$$

The results of this paper do not seem to be a large step towards establishing this conjecture, but they do show that inequalities of this sort may hold even in cases where the Seiberg-Witten invariants are trivial. Another positive note is that our results can be applied to the four-manifold obtained by applying the construction of [4] to a four-torus rather than to a $K3$. Although the resulting 'fake' four-torus is easily distinguished by its larger fundamental group, it is nevertheless interesting to have the lower bound (12) in this case.

In a much more speculative vein, one might look at these results as slight evidence that taut codimension-2 foliations of four-manifolds impose constraints on embedded surfaces. We say that a foliation of a four-manifold X is *taut* if there is a closed 2-form on X which is positive on the leaves. A foliation $\mathcal{F}$ determines an almost complex structure and hence a canonical class $\kappa(\mathcal{F}) \in H^2(X)$. It may be that if $\mathcal{F}$ is a taut foliation of a suitable four-manifold, then embedded surfaces Σ in X satisfy a constraint of the form (11) for the class $\kappa(\mathcal{F})$. The word 'suitable' would need to be defined so as to exclude some counterexamples such as $\Sigma \times S^2$, just as $S^1 \times S^2$ is excluded from Theorem 1.

One could also allow $\mathcal{F}$ to have isolated singularities of the sort that arise in holomorphic foliations of complex surfaces. A suitable extension of our speculation to this situation would encompass the question of the fake $K3$ surfaces X_K.

Acknowledgment. The author thanks Tom Mrowka for the stimulating discussions which prompted this paper, and for helpful comments on earlier drafts.

References

[1] R. H. Crowell and H. F. Trotter, *A class of pretzel knots*, Duke Math. J. **30** (1963), 373–377.

[2] S. K. Donaldson, *Symplectic submanifolds and almost-complex geometry*, J. Differential Geom. **44** (1996), no. 4, 666–705.

[3] Y. M. Eliashberg and W. P. Thurston, *Contact structures and foliations on 3-manifolds*, Turkish J. Math. **20** (1996), no. 1, 19–35.

[4] R. Fintushel and R. Stern, *Knots, links and four-manifolds*, preprint.

[5] D. Gabai, *Foliations and the topology of 3-manifolds*, J. Differential Geom **18** (1983), no. 3, 445–503.

[6] ———, *Foliations and genera of links*, Topology **23** (1984), no. 4, 381–394.

[7] D. Kotschick, J. W. Morgan, and C. H. Taubes, *Four-manifolds without symplectic structures but with nontrivial Seiberg-Witten invariants*, Math. Res. Lett. **2** (1995), no. 2, 119–124.

[8] P. B. Kronheimer and T. S. Mrowka, *The genus of embedded surfaces in the projective plane*, Math. Res. Lett. **1** (1994), no. 6, 797–808.

[9] ———, *Embedded surfaces and the structure of Donaldson's polynomial invariants*, J. Differential Geom. **41** (1995), no. 3, 573–734.

[10] ———, *Monopoles and contact structures*, Inventiones Mathematicae (1997), to appear.

[11] ———, *Scalar curvature and the Thurston norm*, preprint, 1997.

[12] G. Meng and C. H. Taubes, SW = *Milnor torsion*, Math. Res. Lett. **3** (1996), no. 5, 661–674.

[13] J. W. Morgan, Z. Szabó, and C. H. Taubes, *A product formula for the Seiberg-Witten invariants and the generalized Thom conjecture*, J. Differential Geom. **44** (1996), no. 4, 706–788.

[14] C. H. Taubes, *The Seiberg-Witten invariants and symplectic forms*, Math. Res. Lett. **1** (1994), no. 6, 809–822.

[15] ———, *More constraints on symplectic forms from Seiberg-Witten invariants*, Math. Res. Lett. **2** (1995), no. 1, 9–13.

[16] W. P. Thurston, *A norm for the homology of 3-manifolds*, Mem. Amer. Math. Soc. **59** (1986), no. 339, i–vi and 99–130.

[17] E. Witten, *Monopoles and four-manifolds*, Math. Res. Lett. **1** (1994), no. 6, 769–796.

This version: September 9, 1997