Introduction

Motivation

In order to illustrate briefly what these lectures are about, I'd like to give an informal sketch of two closely related theorems from the early days of symplectic topology. The first is a beautiful application of the theory of closed pseudoholomorphic curves as introduced by Gromov in [Gro85], and its proof requires only a few basic facts from this theory, plus some knowledge of the standard homological intersection product from algebraic topology. The second theorem admits a closely analogous proof, but we will see that the intersection-theoretic portion of the argument is difficult to make precise, because it is no longer homological – it requires some generalization of the intersection product in which "cycles" need not be closed. One of the main objectives of the subsequent lectures will be to make this idea precise and demonstrate what else it can be used for.

The statements of these theorems assume familiarity with the notions of minimal symplectic 4-manifolds, symplectomorphisms, symplectic submanifolds, the standard symplectic structure on $\mathbb{R}^4$, the sign of a transverse intersection, and the homological intersection product – some background on all of these topics is covered in Lectures 1 and 2.

Theorem 1 *Suppose (M, ω) is a closed, connected, minimal symplectic 4-manifold containing a pair of symplectic submanifolds $S_1, S_2 \subset M$ with the following properties:*

- *Both are homeomorphic to S^2.*
- *Both have vanishing homological self-intersection number:*

$$[S_1] \cdot [S_1] = [S_2] \cdot [S_2] = 0.$$

- *The set $S_1 \cap S_2 \subset M$ consists of a single transverse and positive intersection.*

Then there exists a symplectomorphism identifying (M, ω) *with* $(S^2 \times S^2, \omega_0)$ *such that* S_1 *and* S_2 *are identified with* $S^2 \times \{const\}$ *and* $\{const\} \times S^2$*, respectively, and* ω_0 *is a product of two area forms on* S^2.

This result says in effect that if we are given a certain type of "local" information about submanifolds of a closed symplectic 4-manifold, then this is enough to recover its global structure. From an alternative perspective, it says that the vast majority of closed symplectic 4-manifolds do not contain certain types of symplectic submanifolds. The second result says something similar, but now the symplectic manifold is noncompact and the "local" information we are given is its structure outside of some compact subset – the theorem is typically summarized by saying that there do not exist any exotic symplectic 4-manifolds that look "standard at infinity."

Theorem 2 *Suppose* (M, ω) *is an open, connected, minimal symplectic 4-manifold with a compact subset* $K \subset M$ *such that* $(M \setminus K, \omega)$ *is symplectomorphic to the complement of a compact subset in the standard symplectic* $\mathbb{R}^4$. *Then* (M, ω) *is globally symplectomorphic to the standard symplectic* $\mathbb{R}^4$.

Remark 3 Both of these theorems appeared in less general forms in Gromov's paper [Gro85] (see §2.4.A_1' and §3.C, respectively). The statements given above are attributed to both Gromov and McDuff, as they rely on the slightly more sophisticated intersection theory of closed holomorphic curves that was developed by McDuff within a few years after Gromov's paper (see, in particular, [McD90]). Theorem 2 can also be rephrased as the statement that S^3 with its standard contact structure admits a unique minimal symplectic filling, and we will discuss this version of the result in Lecture 5 (see, in particular, Corollary 5.7).

Let's sketch a proof of Theorem 1. The starting point is the observation that since S_1 and S_2 are both *symplectic* submanifolds and their intersection is transverse and positive, one can choose a compatible almost complex structure $J : TM \to TM$ on (M, ω) that preserves the tangent spaces of S_1 and S_2 (see §1.1 for more on almost complex structures). This makes S_1 and S_2 into images of embedded *J-holomorphic spheres*, i.e., smooth maps $u: S^2 \to M$ that satisfy the *nonlinear Cauchy–Riemann equation*

$$Tu \circ i = J \circ Tu,$$

where $i: TS^2 \to TS^2$ is the almost complex structure on S^2 resulting from its standard identification with the extended complex plane $\mathbb{C} \cup \{\infty\}$. The advantage of replacing symplectic submanifolds by J-holomorphic spheres is a matter of rigidity: the condition of being a symplectic submanifold is open and

Cambridge University Press
978-1-108-49740-4 — Lectures on Contact 3-Manifolds, Holomorphic Curves and Intersection Theory
Chris Wendl
Excerpt
More Information

thus quite flexible, i.e., the space of all symplectic submanifolds is unmanageably large, whereas J-holomorphic spheres are solutions to an elliptic PDE, and thus tend to come in finite-dimensional moduli spaces, which are sometimes (if we're lucky!) even compact. For this reason, we now consider for each $k = 1, 2$ the *moduli spaces*

$$\mathcal{M}_k(J) := \Big\{ u\colon S^2 \to M \;\Big|\; Tu \circ i = J \circ Tu \text{ and } [u] := u_*[S^2] = [S_k] \in H_2(M) \Big\} \Big/ \operatorname{Aut}(S^2, i),$$

where $\operatorname{Aut}(S^2, i)$ is the group of holomorphic automorphisms $\varphi\colon S^2 \to S^2$ of the extended complex plane (i.e., the Möbius transformations), acting on the space of J-holomorphic maps $u\colon S^2 \to M$ by $\varphi \cdot u := u \circ \varphi$. We assign to this space the natural topology arising from C^∞-convergence of maps. Both $\mathcal{M}_1(J)$ and $\mathcal{M}_2(J)$ are clearly nonempty, since they contain equivalence classes of parametrizations of the submanifolds S_1 and S_2, respectively. One can now apply general results from the theory of J-holomorphic curves to prove that for generic choices of the almost complex structure J, $\mathcal{M}_1(J)$ and $\mathcal{M}_2(J)$ both are compact smooth 2-dimensional manifolds. A quick survey of the analytical results behind this is given in Appendix A.1, and we will sketch the proof in a somewhat more general setting in Lectures 1 (see Lemmas 1.17 and 1.18) and 2, though we do not plan to go too deeply into such analytical details in this book.

What we will discuss in more detail is the intersection-theoretic properties of the J-holomorphic spheres in $\mathcal{M}_1(J)$ and $\mathcal{M}_2(J)$. We observe first that the hypotheses of Theorem 1 clearly imply

$$[S_1] \cdot [S_2] = 1,$$

as this intersection number can be computed as a signed count of transverse intersections between S_1 and S_2, for which there is only one intersection to count, and it is positive. In Lecture 2 and Appendix B, we will discuss a standard result known as *positivity of intersections*, which implies that whenever $u\colon \Sigma \to M$ and $v\colon \Sigma' \to M$ are two closed J-holomorphic curves with nonidentical images in an almost complex 4-manifold M, their intersections are all isolated and count positively toward the homological intersection number $[u] \cdot [v] \in \mathbb{Z}$; moreover, the contribution of each isolated intersection is exactly $+1$ if and only if that intersection is transverse. This is very strong information, from which one can deduce the following:

(1) For each $k = 1, 2$ and every pair of distinct elements $u, v \in \mathcal{M}_k(J)$, the images of $u\colon S^2 \to M$ and $v\colon S^2 \to M$ are disjoint. (This follows from the condition $[S_k] \cdot [S_k] = 0$.)

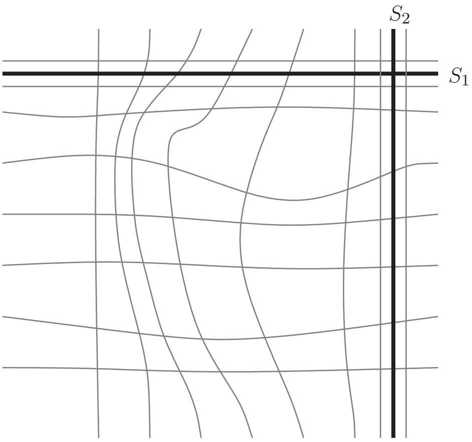

Figure 1 The two symplectic submanifolds $S_1, S_2 \subset M$ generate two transverse foliations by holomorphic spheres in the proof of Theorem 1. The two families can be regarded as a "coordinate grid" that identifies M with $S^2 \times S^2$.

(2) For every $u \in \mathcal{M}_1(J)$ and $v \in \mathcal{M}_2(J)$, the maps $u\colon S^2 \to M$ and $v\colon S^2 \to M$ have exactly one intersection point, which is transverse and positive.

A related result discussed in §2.1, called the *adjunction formula*, makes it possible to characterize in homological terms which J-holomorphic curves in an almost complex 4-manifold are embedded, and in this case it implies

(3) Every element of $\mathcal{M}_1(J)$ or $\mathcal{M}_2(J)$ is embedded.

Finally, we will see in §1.3 that whenever $u \in \mathcal{M}_k(J)$ is an embedded J-holomorphic sphere in one of these moduli spaces, the 2-parameter family of nearby J-holomorphic spheres in $\mathcal{M}_k(J)$ forms a smooth foliation of the neighborhood of $u(S^2)$ in M. Combining this with the compactness of $\mathcal{M}_k(J)$, it follows that the set of points in M that are contained in the images of any of the spheres in $\mathcal{M}_k(J)$ is both open and closed, and thus it is everything: the holomorphic spheres of $\mathcal{M}_k(J)$ foliate M. The result is the "coordinate grid" depicted in Figure 1: starting from the two symplectically embedded spheres $S_1, S_2 \subset M$, we obtain two smooth families of embedded J-holomorphic spheres that each foliate M such that each sphere in $\mathcal{M}_1(J)$ has a unique transverse intersection with each sphere in $\mathcal{M}_2(J)$. It follows that there is a diffeomorphism

$$M \xrightarrow{\cong} \mathcal{M}_1(J) \times \mathcal{M}_2(J), \tag{1}$$

Cambridge University Press
978-1-108-49740-4 — Lectures on Contact 3-Manifolds, Holomorphic Curves and Intersection Theory
Chris Wendl
Excerpt
More Information

assigning to each point $p \in M$ the unique pair of holomorphic spheres $(u, v) \in \mathcal{M}_1(J) \times \mathcal{M}_2(J)$ such that both have p in their images. Moreover, for each individual element of $\mathcal{M}_1(J)$ parametrized by a map $u : S^2 \to M$, there is a diffeomorphism

$$S^2 \xrightarrow{\cong} \mathcal{M}_2(J)$$

sending each $z \in S^2$ to the unique holomorphic sphere $v \in \mathcal{M}_2(J)$ that has $u(z)$ in its image; this proves that $\mathcal{M}_2(J)$ has the topology of S^2, and, in the same manner, one shows $\mathcal{M}_1(J) \cong S^2$. In summary, (1) can now be interpreted as a diffeomorphism from M to $S^2 \times S^2$. There is still a bit of work to be done in identifying the symplectic structure ω with a product of two area forms, but the techniques needed for this are not hard – they involve geometric tools such as the Moser stability theorem for deformations of symplectic forms (see, e.g., [MS17]), but no serious analysis is required.

The original proof of Theorem 2 used a clever "capping" trick to derive it from Theorem 1. For this motivational discussion, I would like to sketch a different proof that is conceptually simpler, but trickier in the technical details.

By the hypotheses of Theorem 2, we can decompose the open symplectic manifold (M, ω) into two regions: one is the compact (but otherwise completely unknown) subset $K \subset M$, and the other is a region that we can identify with $(\mathbb{R}^4 \backslash K', \omega_{\mathrm{std}})$ for some compact set $K' \subset \mathbb{R}^4$, where ω_{std} denotes the standard symplectic form on $\mathbb{R}^4$. We would like to argue as we did in Theorem 1; that is, find a nice pair of "seed curves" to generate two well-behaved moduli spaces of J-holomorphic curves that can then be used to form a coordinate grid identifying M with $\mathbb{R}^4$. One easy way to find such seed curves is by observing that $\mathbb{R}^4$ has a natural identification with $\mathbb{C}^2$ such that the natural multiplication by i on $\mathbb{C}^2$ defines a compatible almost complex structure on $(\mathbb{R}^4, \omega_{\mathrm{std}})$. This is useful for the following reason: $\mathbb{C}^2$ contains two obvious families of holomorphic planes

$$f_w : \mathbb{C} \to \mathbb{C}^2 : z \mapsto (z, w), \qquad \text{for } w \in \mathbb{C},$$
$$g_w : \mathbb{C} \to \mathbb{C}^2 : z \mapsto (w, z), \qquad \text{for } w \in \mathbb{C},$$

all of which are properly embedded maps, with two distinct types of asymptotic behavior. To describe the latter, choose a large constant $R > 0$, let $\mathbb{D}^4_R \subset \mathbb{C}^2$ denote the disk of radius R and identify $\mathbb{C}^2 \backslash \mathbb{D}^4_R$ with $(R, \infty) \times S^3$ by viewing S^3 as the unit sphere in $\mathbb{C}^2$ and applying the diffeomorphism

$$(R, \infty) \times S^3 \xrightarrow{\cong} \mathbb{C}^2 \backslash \mathbb{D}^4_R : (r, x) \mapsto rx.$$

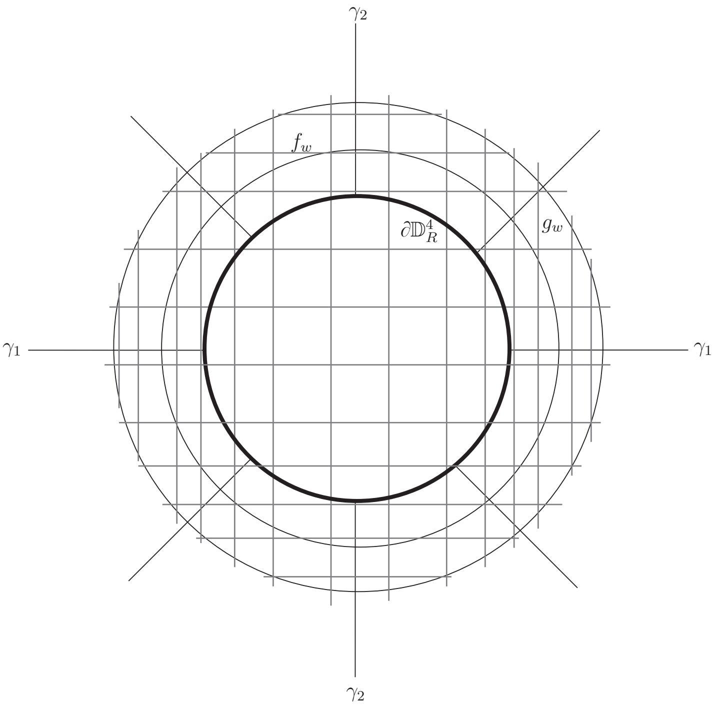

Figure 2 The two families of properly embedded holomorphic planes f_w and g_w form a coordinate grid for $\mathbb{C}^2$ and are each asymptotic on the cylindrical end $\mathbb{C}^2\backslash\mathbb{D}^4_R \cong (R,\infty)\times S^3$ to one of two specific loops $\gamma_1, \gamma_2 \subset S^3$.

Then each f_w or g_w maps a neighborhood of infinity into an arbitrarily small neighborhood of the cylinder $(R,\infty)\times\gamma_1$ or $(R,\infty)\times\gamma_2$, respectively, where we define

$$\gamma_1 := S^1 \times \{0\} \subset S^3 \subset \mathbb{C}^2, \qquad \gamma_2 := \{0\} \times S^1 \subset S^3 \subset \mathbb{C}^2.$$

A schematic picture of this asymptotic behavior and the resulting transverse pair of holomorphic foliations of $\mathbb{C}^2$ is shown in Figure 2. Informally, we will say that the planes f_w are asymptotic to γ_1 and the planes g_w are asymptotic to γ_2; more precise definitions of this terminology will appear in §2.4 when we discuss asymptotically cylindrical maps.

Now since $K' \subset \mathbb{C}^2 = \mathbb{R}^4$ is compact, $\mathbb{D}^4_R$ will contain K' for any $R > 0$ sufficiently large, so that we can also regard (M,ω) as containing a copy of the region identified above with $(R,\infty)\times S^3$. Let us fix such a radius and choose a compatible almost complex structure J on (M,ω) that matches the standard multiplication by i on $\mathbb{C}^2 \setminus \mathbb{D}^4_R \cong (R, \infty) \times S^3$. The curves f_w and g_w can then be regarded as J-holomorphic planes in M for every $w \in \mathbb{C}$ with $|w| > R$, and just as in Theorem 1, these two families define elements in a pair of connected moduli spaces $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ of J-holomorphic planes in M, where we can use the loops γ_1 and γ_2 to prescribe the asymptotic behavior of the curves in the moduli spaces. There exists a well-developed theory of moduli spaces of J-holomorphic curves with this type of asymptotic behavior, a survey of which is given in Appendix A.2. In the present context, it can be applied to prove that $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ are both smooth 2-dimensional manifolds, and they are also compact except for the obvious way in which they are not: a sequence $u_j \in \mathcal{M}_k(J; \gamma_k)$ for $k \in \{1, 2\}$ will fail to have a convergent subsequence if and only if for large j it is of the form $u_j = f_{w_j} \in \mathcal{M}_1(J; \gamma_1)$ or $u_j = g_{w_j} \in \mathcal{M}_2(J; \gamma_2)$ for a sequence $w_j \in \mathbb{C}$ with $|w_j| \to \infty$. This gives each of $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ the topology of a compact surface with one boundary component attached to a *cylindrical end* of the form $\mathbb{C} \setminus \mathbb{D}_R \cong (R, \infty) \times S^1$.

If we want to apply these two moduli spaces the same way they were used in Theorem 1, then we need to establish the following:

Lemma 4 *The moduli spaces $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ described above have the following properties:*

(1) For each $k = 1, 2$ and every pair of distinct elements $u, v \in \mathcal{M}_k(J; \gamma_k)$, the images of $u : \mathbb{C} \to M$ and $v : \mathbb{C} \to M$ are disjoint.
(2) For every $u \in \mathcal{M}_1(J; \gamma_1)$ and $v \in \mathcal{M}_2(J; \gamma_2)$, the maps $u : \mathbb{C} \to M$ and $v : \mathbb{C} \to M$ have exactly one intersection point, which is transverse and positive.
(3) Every element of $\mathcal{M}_1(J; \gamma_1)$ or $\mathcal{M}_2(J; \gamma_2)$ is embedded.

Indeed, one can then argue exactly as in the proof of Theorem 1 that the two moduli spaces $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ form two transverse smooth foliations of M by planes, producing a coordinate grid (see Figure 3) that identifies M with $\mathbb{C} \times \mathbb{C} \cong \mathbb{R}^4$. The question I would now like to focus on is this: why is Lemma 4 true?

The answer does not come from homological intersection theory, as the curves in $\mathcal{M}_1(J; \gamma_1)$ and $\mathcal{M}_2(J; \gamma_2)$ are noncompact and do not represent homology classes. One can, however, use differential topological arguments to verify the second claim in the lemma: the fact that each f_w intersects each $g_{w'}$ exactly once transversely implies via a homotopy argument that the same will be true for any pair $u \in \mathcal{M}_1(J; \gamma_1)$ and $v \in \mathcal{M}_2(J; \gamma_2)$. Indeed, $\mathcal{M}_1(J; \gamma_1)$

Cambridge University Press
978-1-108-49740-4 — Lectures on Contact 3-Manifolds, Holomorphic Curves and Intersection Theory
Chris Wendl
Excerpt
More Information

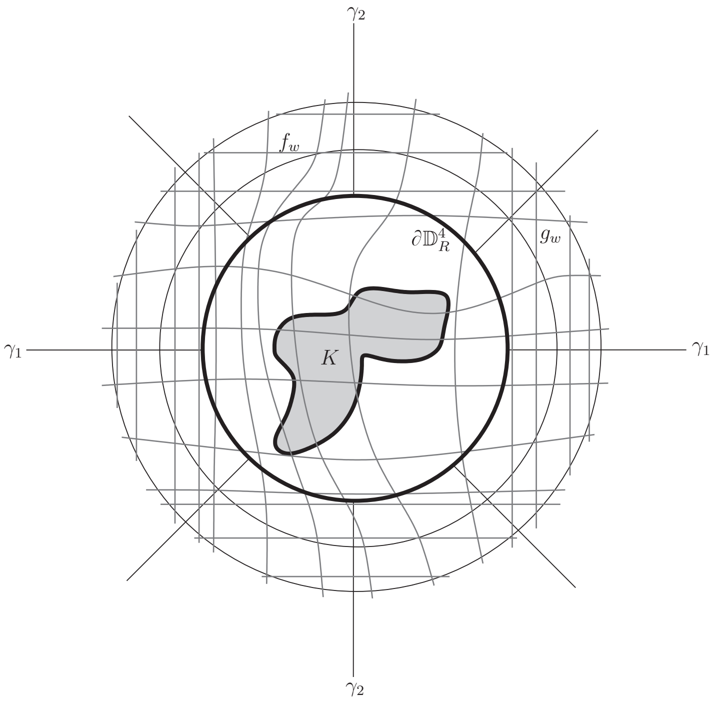

Figure 3 The moduli spaces $\mathcal{M}_1(J;\gamma_1)$ and $\mathcal{M}_2(J;\gamma_2)$ of proper J-holomorphic planes asymptotic to the loops $\gamma_1, \gamma_2 \subset S^3$ form two transverse foliations of M in Theorem 2, building a coordinate grid that proves $M \cong \mathbb{C} \times \mathbb{C} = \mathbb{R}^4$.

and $\mathcal{M}_2(J;\gamma_2)$ are each connected spaces of properly embedded planes that are asymptotic to disjoint loops in S^3, and thus they map neighborhoods of infinity to completely disjoint regions near infinity in M. This ensures that there exist homotopies of properly embedded maps

$$u_\tau : \mathbb{C} \to M, \qquad v_\tau : \mathbb{C} \to M, \qquad \tau \in [0,1]$$

with $u_0 = u$, $u_1 = f_w$, $v_0 = v$ and $v_1 = g_{w'}$ such that the intersections of u_τ with v_τ for every $\tau \in [0,1]$ are confined to compact subsets of both domains. Standard arguments as in [Mil97] then imply that u and v must have the same algebraic intersection count as f_w and $g_{w'}$, which is 1, so in light of positivity of intersections, u and v can only have one intersection point, and it must be transverse.

This type of argument does not suffice to prove the other two claims in Lemma 4. For example, suppose we would like to prove that two distinct curves $u, v \in \mathcal{M}_1(J;\gamma_1)$ must always be disjoint. It is easy to believe this in light of

Cambridge University Press
978-1-108-49740-4 — Lectures on Contact 3-Manifolds, Holomorphic Curves and Intersection Theory
Chris Wendl
Excerpt
More Information

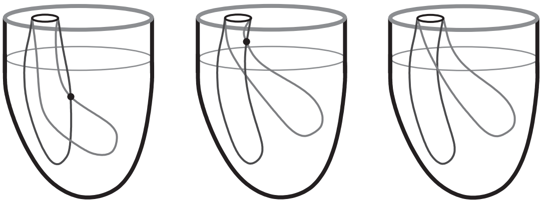

Figure 4 The algebraic intersection count $u \cdot v \in \mathbb{Z}$ between two proper maps of noncompact domains can change under homotopies if the two maps have matching asymptotic behavior.

the curves that we can explicitly see; i.e., f_w and $f_{w'}$ both belong to $\mathcal{M}_1(J;\gamma_1)$ for any $w, w' \in \mathbb{C}$ sufficiently large, and they are clearly disjoint if $w \neq w'$. To extend this to the curves that we cannot explicitly see because they do not live entirely in the region $(R,\infty) \times S^3 \subset M$, we would ideally like to argue via homotopy invariance, namely that if u_τ and v_τ are two continuous families of curves in $\mathcal{M}_1(J;\gamma_1)$ with u_0 and v_0 disjoint, then u_1 and v_1 must also be disjoint. But here we have a problem that did not arise in the previous paragraph: the curves u_τ and v_τ in this homotopy are always asymptotic to *the same* loop $\gamma_1 \subset S^3$, so their images in M always become arbitrarily close to each other in the cylindrical end $(R,\infty) \times S^3$. In this situation, there is no way to make sure that intersections are confined to compact subsets, and we can imagine, in fact, that under a homotopy, some intersections might just escape to infinity and disappear (see Figure 4)!

It is a remarkable fact that, in the situation under consideration, this nightmare scenario cannot happen, and Lemma 4 is indeed true. To understand why, we will have to explore the asymptotic behavior of noncompact J-holomorphic curves much more deeply. Still more interesting, perhaps, is that in more general situations, the nightmare scenario of Figure 4 really can happen, but it can also be *controlled*: one can define an *asymptotic contribution* that measures the possibility for "hidden" intersections to emerge from infinity under small perturbations. It turns out that just like the contribution of an isolated intersection between two J-holomorphic curves, this asymptotic contribution is always nonnegative, and adding it to the algebraic count of actual intersections produces a meaningful homotopy-invariant intersection product. Once this product and the corresponding generalization of the adjunction formula have been understood, proving results like Lemma 4 becomes quite easy.

The first hints of a systematic intersection theory for noncompact holomorphic curves appeared in Hutchings's work on embedded contact homology

[Hut02], and the theory was developed in earnest a few years later in the Ph.D. thesis of Richard Siefring [Sie05] and his two papers [Sie08, Sie11]. Our primary objectives in these lectures will be to explain where this theory comes from, demonstrate how to use it, and give some examples of what it can be used for. We'll start in Lectures 1 and 2 by reviewing the intersection theory for closed holomorphic curves and discussing one of its most famous applications, McDuff's theorem [McD90] on symplectic ruled surfaces (which is a variation on Theorem 1). The asymptotic analysis required for Siefring's theory is then surveyed in Lecture 3 (mostly without the proofs since these are analytically somewhat intense), and Lecture 4 uses these asymptotic results to define the precise generalizations of the homological intersection product and the adjunction formula that are needed for results such as Lemma 4. In Lecture 5, we will demonstrate how to use the theory via a generalization of Theorem 2, framed in the language of contact 3-manifolds and their symplectic fillings.