The space of embedded minimal surfaces of fixed genus in a 3-manifold V; Fixed genus

By TOBIAS H. COLDING and WILLIAM P. MINICOZZI II

Abstract

This paper is the fifth and final in a series on embedded minimal surfaces. Following our earlier papers on disks, we prove here two main structure theorems for *nonsimply connected* embedded minimal surfaces of any given fixed genus.

The first of these asserts that any such surface *without* small necks can be obtained by gluing together two oppositely-oriented double spiral staircases.

The second gives a pair of pants decomposition of any such surface when there *are* small necks, cutting the surface along a collection of short curves. After the cutting, we are left with graphical pieces that are defined over a disk with either one or two sub-disks removed (a topological disk with two sub-disks removed is called a pair of pants).

Both of these structures occur as different extremes in the two-parameter family of minimal surfaces known as the Riemann examples.

Contents

0. Introduction	4
0.1. Brief outline of the paper and overview of the proofs	12
1. Multi-valued graphs	14
Part I. Results for disks from [CM04b]–[CM04e]	15
I.1. The lamination theorem and one-sided curvature estimate	15
Part II. The singular set S and limit lamination \mathcal{L}'	17
II.1. The singular set S	18
II.1.1. Convergence away from S	18
II.2. The local structure of \mathcal{L}' near a point of \mathcal{S}_{ulsc}	19
II.2.1. Collapsed leaves of \mathcal{L}'	22
II.3. The structure of the collapsed leaves of \mathcal{L}'	23

The authors were partially supported by NSF Grants.

-

^{© 2015} Department of Mathematics, Princeton University.

2	TOBIAS H. COLDING and WILLIAM P. MINICOZZI II	
	II.3.1. Proving property (1) of Proposition II.3.1:	
	Isolated removable singularities	24
	II.3.2. Each leaf is a limit of multi-valued graphs in the Σ_i 's	30
	II.3.3. Property (2) of Proposition II.3.1: Each collapsed leaf is stable	32
	II.3.4. Property (3) of Proposition II.3.1: Opposite orientations at	
	distinct points of $\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{ulsc}}$	39
Р	Part III. When the surfaces are ULSC:	
Т	he proof of Theorem 0.9	41
	III.0.5. Property (1) in Proposition III.0.2: Collapsed leaves are planar	42
	III.0.6. Property (2) in Proposition III.0.2: Ruling out just one	
	point of \mathcal{S}_{ulsc} in a leaf	42
	III.1. Properness and the limit foliation	47
	III.1.1. Establishing properness: The proof of (\star)	49
	III.2. Completing the proof of Theorem 0.9	57
	III.3. Sequences with fixed genus	58
	III.4. An application: A one-sided property for ULSC surfaces	58
Ρ	eart IV. When the surfaces are not ULSC:	
Т	The proof of Theorem 0.14	61
	IV.1. Proving $(C1)$ in Theorem 0.14:	
	A plane through each point of \mathcal{S}_{neck}	62
	IV.1.1. The proof of Proposition IV.1.1: A decomposition near each	
	point of $\mathcal{S}_{ m neck}$	64
	IV.1.2. Step (1): Decomposing Σ_j into ULSC pieces	74
	IV.1.3. Step (2): The ULSC pieces of Σ_j contain graphs	76
	IV.2. The ULSC regions of the lamination: $(C2)$ and (D)	
	in Theorem 0.14	79
	IV.2.1. The leaf Γ cannot cross the limit planes	80
	IV.2.2. The properties of a collapsed leaf Γ	82
	IV.2.3. $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of exactly two points	82
	IV.2.4. The proof of Proposition IV.2.2	89
	IV.2.5. The proof of $(C2)$ and (D) in Theorem 0.14	90
	IV.3. Putting it all together: The proof of Theorem 0.14	91
Ρ	Part V. The no mixing theorem, Theorem 0.4	91
Ρ	eart VI. Completing the proofs of Theorems 0.6 and 0.12	94
	VI.1. Blow up results for ULSC surfaces	94
	VI.1.1. Finding short separating curves	95
	VI.2. Complete leaves of \mathcal{L}'	97
	VI.3. Incomplete leaves of \mathcal{L}'	103

FIXED	GENUS

VI.3.1. If Γ is not flat, then $\Gamma_{\text{Clos}} \cap \mathcal{S}$ consists of at most two points	104
VI.3.2. The proof of Lemma VI.3.1	106
VI.4. The proofs of Theorems 0.6 and 0.12	114
Part VII. Modifications in the positive genus case	114
VII.1. The definitions and statements for positive genus	115
VII.1.1. The sets S_{nock} and S_{nlsc} for positive genus	115
VII.1.2. The statements of the theorems for positive genus	116
VII.2. The local structure near points in S_{plac} and S_{pack}	117
VII.2.1. Collapsed leaves	119
VII.3. Part III: When the surfaces are ULSC	120
VII.4. Parts IV and V: When the surfaces are not ULSC	122
VII.4.1. The local structure (β)	123
VII.5. Planes through S_{ulsc} and the proofs of (C2) and (D)	124
VII.6. The remaining cases of Proposition VII.4.2	125
VII.6.1. Part V: The no-mixing theorem	126
VII.6.2. Part VI: The leaves are all flat	126
Part VIII. Appendices	129
Appendix A. Surfaces with stable covers	129
A.1. Going from stability of a covering space to stability	
of a surface itself	129
A.2. A surface and stable cover with cyclic holonomy group	
where the previous lemma applies	132
A.3. A Bernstein theorem for <i>incomplete</i> surfaces	133
Appendix B. An extension of [CM08]	136
B.1. Chord-arc bounds for ULSC surfaces	136
B.2. Chord-arc and area bounds for surfaces with bounded curvature	137
Appendix C. Estimates for stable surfaces	137
C.1. The regularity of the distance function to the interior boundary	138
C.2. The proof of the "stable graph" proposition	141
Appendix D. Blowing up intrinsically on the scale of	
nontrivial topology	145
Appendix E. Minimal surfaces with a quadratic curvature bound	
in a half-space	147
References	149



The two main structure theorems for nonsimply connected surfaces:

Figure 1. Absence of necks: The surface can be obtained by gluing together two oppositelyoriented double spiral staircases.





The curves that we cut along.

Figure 2. Presence of necks: The surface can be decomposed into a collection of pair of pants by cutting along short curves.

0. Introduction

This paper is the fifth and final in a series where we describe the space of all properly embedded minimal surfaces of fixed genus in a fixed (but arbitrary) closed 3-manifold. We will see that the key is to understand the structure of an embedded minimal planar domain in a ball in \mathbb{R}^3 . Since the case of disks was considered in the first four papers, the focus here is on nonsimply connected planar domains; the two main structure theorems of this paper are illustrated in Figures 1 and 2.

We will first restrict to the case of planar domains, i.e., when the surfaces have genus zero. In particular, the main theorems will first be stated and proved for planar domains. We will see that the general case of fixed genus requires only minor changes. The necessary changes to the main theorems and the modifications needed for their proofs will be given in Part VII.

Sequences of planar domains that are not simply connected are, after passing to a subsequence, naturally divided into two separate cases depending on whether or not the topology is concentrating at points. To distinguish between these cases, we will say that a sequence of surfaces $\Sigma_i^2 \subset \mathbf{R}^3$ is uniformly locally simply connected (or ULSC) if for each compact subset K of \mathbf{R}^3 , there exists a constant $r_0 > 0$ (depending on K) so that for every $x \in K$, all $r \leq r_0$,¹ and every surface Σ_i ,

(0.1) each connected component of $B_r(x) \cap \Sigma_i$ is a disk.

¹If each component of the intersection of a minimal surface with a ball of radius r_0 is a disk, then so are the intersections with all sub-balls by the convex hull property (see, e.g., Lemma C.1 in [CM04e]). Therefore, it would be enough that (0.1) holds for $r = r_0$.

For instance, a sequence of rescaled catenoids where the necks shrink to zero is not ULSC, whereas a sequence of rescaled helicoids is.

Another way of locally distinguishing sequences where the topology does not concentrate from sequences where it does comes from analyzing the singular set. The singular set S is defined to be the set of points where the curvature is blowing up. That is, a point y in \mathbb{R}^3 is in S for a sequence Σ_i if

(0.2)
$$\sup_{B_r(y)\cap\Sigma_i} |A|^2 \to \infty \text{ as } i \to \infty \text{ for all } r > 0.$$

We will show that for embedded minimal surfaces, S consists of two types of points. The first type is roughly modelled on rescaled helicoids and the second on rescaled catenoids:

- A point y in \mathbb{R}^3 is in \mathcal{S}_{ulsc} if the curvature for the sequence Σ_i blows up at y and the sequence is ULSC in a neighborhood of y.
- A point y in R³ is in S_{neck} if the sequence is not ULSC in any neighborhood of y. In this case, a sequence of closed noncontractible curves γ_i ⊂ Σ_i converges to y.

The sets S_{neck} and S_{ulsc} are obviously disjoint and the curvature blows up at both, so $S_{\text{neck}} \cup S_{\text{ulsc}} \subset S$. Proposition I.0.19 of [CM04d] implies that, after passing to a further subsequence, each point $y \in S \setminus S_{\text{neck}}$ has a radius $r_y > 0$ so that each component of $B_{r_y}(y) \cap \Sigma_j$ is a disk for every j. In other words, this proposition implies that S is given as the disjoint union

(0.3)
$$\mathcal{S} = \mathcal{S}_{\text{neck}} \cup \mathcal{S}_{\text{ulsc}}.$$

Note that $S_{\text{neck}} = \emptyset$ is equivalent to that the sequence is ULSC as is the case for sequences of rescaled helicoids. On the other hand, $S_{\text{ulsc}} = \emptyset$ for sequences of rescaled catenoids. These definitions of S_{ulsc} and S_{neck} are specific to the genus zero case that we are focusing on now; the slightly different definitions in the higher genus case can be found around equation (VII.1.2).

We will show that every sequence Σ_i has a subsequence that is either ULSC or for which S_{ulsc} is empty. This is the next "no mixing" theorem. We will see later that these two different cases give two very different structures.

THEOREM 0.4. If $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ is a sequence of compact embedded minimal planar domains² with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$, then there is a subsequence with either $S_{\text{ulsc}} = \emptyset$ or $S_{\text{neck}} = \emptyset$.

In view of Theorem 0.4 and the earlier results for disks, it is natural to first analyze sequences that are ULSC, so where $S_{neck} = \emptyset$, and second analyze sequences where S_{ulsc} is empty. We will do this next.

²The theorem holds also for sequences with fixed genus; see Part VII.

As already mentioned, our main theorems deal with sequences $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ of compact embedded minimal planar domains with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$. We will assume here that these planar domains are not disks. (Recall that the case of disks was dealt with in [CM04b]–[CM04e].) In particular, we will assume that for each *i*, there exists some $y_i \in \mathbf{R}^3$ and $s_i > 0$ so that

(0.5) some component of $B_{s_i}(y_i) \cap \Sigma_i$ is not a disk.

Moreover, if the nonsimply connected balls $B_{s_i}(y_i)$ "run off to infinity" (i.e., if each connected component of $B_{R'_i}(0) \cap \Sigma_i$ is a disk for some $R'_i \to \infty$), then the results of [CM04b]–[CM04e] apply. Therefore, after passing to a subsequence, we can assume that the surfaces are uniformly not disks — namely, that there exists some R > 0 so that (0.5) holds with $s_i = R$ and $y_i = 0$ for all *i*.

In general, we will allow our sequence of surfaces to have bounded genus. Recall that for a surface Σ with boundary $\partial \Sigma$, the genus of Σ is the genus of the closed surface $\hat{\Sigma}$ obtained by adding a disk to each boundary circle. The genus of a union of disjoint surfaces is the sum of the genuses. Therefore, a surface with boundary has nonnegative genus; the genus is zero if and only if it is a planar domain. For example, the disk and the annulus both have genus zero; on the other hand, a closed surface of genus g with any number of disks removed has genus g.

Common for both the ULSC case and the case where S_{ulsc} is empty is that the limits are always laminations by flat parallel planes and the singular sets are always closed subsets contained in the union of the planes. This is the content of the next theorem.

THEOREM 0.6. Let $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ be a sequence of compact embedded minimal planar domains³ with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$. If

(0.7)
$$\sup_{B_1 \cap \Sigma_i} |A|^2 \to \infty,$$

then, after a rotation of \mathbf{R}^3 , there exists a subsequence Σ_j , a lamination $\mathcal{L} = \{x_3 = t\}_{\{t \in \mathcal{I}\}}$ of \mathbf{R}^3 by parallel planes (where $\mathcal{I} \subset \mathbf{R}$ is a closed set), and a closed nonempty set S in the union of the leaves of \mathcal{L} such that

- (A) For each $1 > \alpha > 0$, $\Sigma_j \setminus S$ converges in the C^{α} -topology to the lamination $\mathcal{L} \setminus S$.
- (B) $\sup_{B_r(x)\cap\Sigma_j} |A|^2 \to \infty$ as $j \to \infty$ for all r > 0 and $x \in S$. (The curvatures blow up along S.)

³The theorem holds also for sequences with fixed genus; see Part VII.

Before discussing the general ULSC case, it is useful to recall the case of disks. One consequence of [CM04b]–[CM04e] is that there are only two *local models* for ULSC sequences of embedded minimal surfaces. That is, locally in a ball in \mathbf{R}^3 , one of following holds:

- The curvatures are bounded and the surfaces are locally graphs over a plane.
- The curvatures blow up and the surfaces are locally *double spiral staircases*.

Both of these cases are illustrated by taking a sequence of rescalings of the helicoid; the first case occurs away from the axis, while the second case occurs on the axis. Namely, recall that the helicoid is the minimal surface Σ in \mathbf{R}^3 parametrized by

$$(0.8) (s\cos t, s\sin t, t) \text{ where } s, t \in \mathbf{R}.$$

If we take a sequence $\Sigma_i = a_i \Sigma$ of rescaled helicoids where $a_i \to 0$, then the curvature blows up along the vertical axis but is bounded away from this axis. Thus, we get that

- The intersection of the rescaled helicoids with a ball *away from* the vertical axis gives a collection of graphs over the plane $\{x_3 = 0\}$.
- The intersection of the rescaled helicoids with a ball *centered on* the vertical axis gives a double spiral staircase.

Loosely speaking, our next result shows that when the sequence is ULSC (but not simply connected), a subsequence converges to a foliation by parallel planes away from two lines S_1 and S_2 ; see Figure 3. The lines S_1 and S_2 are disjoint and orthogonal to the leaves of the foliation and the two lines are precisely the points where the curvature is blowing up. This is similar to the case of disks, except that we get two singular curves for nondisks as opposed to just one singular curve for disks. (The precise statement for disks is recalled in Part I.)



Singular lines S_1 and S_2 .

Figure 3. Theorem 0.9: Limits of sequences of nonsimply connected, yet ULSC, surfaces with curvature blowing up. The singular set consists of two lines S_1 and S_2 , and the limit is a foliation by flat parallel planes.

THEOREM 0.9. Let a sequence Σ_i , limit lamination \mathcal{L} , and singular set \mathcal{S} be as in Theorem 0.6.⁴ Suppose that each Σ_i satisfies (0.5) with $s_i = R > 1$ and $y_i = 0$. If every Σ_i is ULSC and

(0.10)
$$\sup_{B_1 \cap \Sigma_i} |A|^2 \to \infty,$$

then the limit lamination \mathcal{L} is the foliation $\mathcal{F} = \{x_3 = t\}_t$ and the singular set \mathcal{S} is the union of two disjoint lines \mathcal{S}_1 and \mathcal{S}_2 such that

- (C_{ulsc}) Away from $S_1 \cup S_2$, each Σ_j consists of exactly two multi-valued graphs spiraling together. Near S_1 and S_2 , the pair of multi-valued graphs form double spiral staircases with opposite orientations at S_1 and S_2 . Thus, circling only S_1 or only S_2 results in going either up or down, while a path circling both S_1 and S_2 closes up (see Figure 6).
- $(D_{\text{ulsc}}) \mathcal{S}_1$ and \mathcal{S}_2 are orthogonal to the leaves of the foliation.

Notice that Theorem 0.9 shows that if the fixed genus ULSC surfaces Σ_j have curvature blowing up, then they essentially have genus zero. More precisely, given an arbitrarily large ball $B_R \subset \mathbf{R}^3$, then $B_R \cap \Sigma_j$ has genus zero for j sufficiently large. To see this, combine the double spiral staircase structure near the two singular curves that holds for ULSC sequences (cf. Figure 6) with the smooth convergence elsewhere.

Despite the similarity of Theorem 0.9 to the case of disks, it is worth noting that the results for disks do not alone give this theorem. Namely, even though the ULSC sequence consists *locally* of disks, the compactness result for disks was in the *global* case where the radii go to infinity. One might wrongly think that Theorem 0.9 could be proven using the results for disks and a blow up argument. However, local examples constructed in [CM04a] show the difficulty with such an argument.⁵ We shall explain this further later together with what else is needed for the proof.

When the sequence is no longer ULSC, then there are other local models for the surfaces. The simplest example is a sequence of rescaled catenoids; the

⁴The theorem holds also for sequences with fixed genus with one minor change in the conclusion and one in the hypothesis. The change in the hypothesis is that we do not assume (0.5). The change in the conclusion is that there might be either one or two singular curves. Hypothesis (0.5) is used in the genus zero case to show that there cannot be just one singular curve. The reason that we will not assume (0.5) in the fixed genus case is that there can be either one or two singular curves in this case regardless; see Part VII.

⁵In [CM04a], we constructed a sequence of embedded minimal disks Σ_i in the unit ball B_1 with $\partial \Sigma_i \subset \partial B_1$ where the curvatures blow up only at 0. This sequence converges to a lamination of $B_1 \setminus \{0\}$ that cannot be extended smoothly to a lamination of B_1 ; that is to say, 0 is not a removable singularity. This should be contrasted with Theorem 0.9 where every singular point is a removable singularity for the limit foliation by parallel planes.

catenoid is the minimal surface in \mathbf{R}^3 parametrized by

(0.11)
$$(\cosh s \cos t, \cosh s \sin t, s)$$
 where $s, t \in \mathbf{R}$.

A sequence of rescaled catenoids converges with multiplicity two to the flat plane. The convergence is in the C^{∞} -topology except at 0 where $|A|^2 \to \infty$. This sequence of rescaled catenoids is not ULSC because the simple closed geodesic on the catenoid — i.e., the unit circle in the $\{x_3 = 0\}$ plane — is noncontractible and the rescalings shrink it down to the origin.

One can get other types of curvature blowup by considering the family of embedded minimal planar domains known as the Riemann examples.⁶ Modulo translations and rotations, this is a two-parameter family of periodic minimal surfaces, where the parameters can be thought of as the size of the necks and the angle from one fundamental domain to the next. By choosing the two parameters appropriately, one can produce sequences of Riemann examples that illustrate both of the two structure theorems (cf. Figures 1 and 2):

- (1) If we take a sequence of Riemann examples where the neck size is fixed and the angles go to $\frac{\pi}{2}$, then the surfaces with angle near $\frac{\pi}{2}$ can be obtained by gluing together two oppositely-oriented double spiral staircases. Each double spiral staircase looks like a helicoid. This sequence of Riemann examples converges to a foliation by parallel planes. The convergence is smooth away from the axes of the two helicoids. (These two axes are the singular set S where the curvature blows up.) The sequence is ULSC since the size of the necks is fixed and thus illustrates the first structure theorem, Theorem 0.9.
- (2) If we take a sequence of examples where the neck sizes go to zero, then we get a sequence that is *not* ULSC. However, the surfaces can be cut along short curves into collections of graphical pairs of pants. The short curves converge to points, and the graphical pieces converge to flat planes except at these points, illustrating the second structure theorem, Theorem 0.12 below.

With these examples in mind, we are now ready to state our second main structure theorem describing the case where S_{ulsc} is empty.

⁶See http://www.msri.org/publications/sgp/jim/geom/minimal/library/riemann/index. html for a description, as well as computer graphics, of these surfaces.

THEOREM 0.12. Let a sequence Σ_i , limit lamination \mathcal{L} , and singular set \mathcal{S} be as in Theorem 0.6.⁷ If $\mathcal{S}_{ulsc} = \emptyset$ and

(0.13)
$$\sup_{B_1 \cap \Sigma_i} |A|^2 \to \infty,$$

then $S = S_{neck}$ by (0.3) and

 (C_{neck}) Each point y in S comes with a sequence of graphs in Σ_j that converge to the plane $\{x_3 = x_3(y)\}$. The convergence is in the C^{∞} -topology away from the point y and possibly also one other point in $\{x_3 = x_3(y)\} \cap S$. If the convergence is away from one point, then these graphs are defined over annuli; if the convergence is away from two points, then the graphs are defined over disks with two subdisks removed.

Theorem 0.12, as well as Theorem 0.4, are proven by first analyzing sequences of minimal surfaces without any assumptions on the sets \mathcal{S}_{ulsc} and \mathcal{S}_{neck} . In this general case, we show that a subsequence converges to a lamination \mathcal{L}' divided into regions where Theorem 0.9 holds and regions where Theorem 0.12 holds. This convergence is in the smooth topology away from the singular set \mathcal{S} where the curvature blows up. Moreover, each point of \mathcal{S} comes with a plane and these planes are essentially contained in \mathcal{L}' ; see (P) below. The set of heights of the planes is a closed subset $\mathcal{I} \subset \mathbf{R}$ but may not be all of \mathbf{R} as it was in Theorem 0.9 and may not even be connected. The behavior of the sequence is different at the two types of singular points in \mathcal{S} the set S_{neck} of "catenoid points" and the set S_{ulsc} of ULSC singular points. We will see that S_{ulsc} consists of a union of Lipschitz curves transverse to the lamination \mathcal{L} . This structure of \mathcal{S}_{ulsc} implies that the set of heights in \mathcal{I} that intersect S_{ulsc} is a union of intervals; thus this part of the lamination is foliated. In contrast, we will not get any structure of the set of "catenoid points" S_{neck} ; see (D) below. Given a point y in \mathcal{S}_{neck} , we will get a sequence of graphs in Σ_i converging to a plane through y; see (C1) below. This convergence will be in the smooth topology away from either one or two singular points, one of which is y. Moreover, this limit plane through y will be a leaf of the lamination \mathcal{L} .

The precise statement of the compactness theorem for sequences that are neither necessarily ULSC nor with $S_{ulsc} = \emptyset$ is the following (see Figure 4):

⁷The theorem holds also for sequences with fixed genus with one small change in (C_{neck}) . Namely, the number of points in (C_{neck}) is bounded by two plus the bound for the genus; see Part VII.

THEOREM 0.14. Let $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ be a sequence of compact embedded minimal planar domains⁸ with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$. If

$$(0.15)\qquad\qquad\qquad \sup_{B_1\cap\Sigma_i}|A|^2\to\infty$$

then there is a subsequence Σ_j , a closed set S, and a lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$ so that

- (A) For each $1 > \alpha > 0$, $\Sigma_j \setminus S$ converges in the C^{α} -topology to the lamination \mathcal{L}' .
- (B) $\sup_{B_r(x)\cap\Sigma_j} |A|^2 \to \infty$ as $j \to \infty$ for all r > 0 and $x \in S$. (The curvatures blow up along S.)
- (C1) (C_{neck}) from Theorem 0.12 holds for each point y in S_{neck} . Further, the possible second singular point in the plane is also in S_{neck} .
- (C2) (C_{ulsc}) from Theorem 0.9 holds locally near S_{ulsc} . More precisely, each point y in S_{ulsc} comes with a sequence of multi-valued graphs in Σ_j that converge to the plane $\{x_3 = x_3(y)\}$. The convergence is in the C^{∞} -topology away from the point y and possibly also one other point in $\{x_3 = x_3(y)\} \cap S_{\text{ulsc}}$. These two possibilities correspond to the two types of multi-valued graphs defined in Section 1.
- (D) The set S_{ulsc} is a union of Lipschitz curves transverse to the lamination. The leaves intersecting S_{ulsc} are planes foliating an open subset of \mathbf{R}^3 that does not intersect S_{neck} . For the set S_{neck} , we make no claim about the structure.
- (P) Together (C1) and (C2) give a sequence of graphs or multi-valued graphs converging to a plane through each point of S. If P is one of these planes, then each leaf of L' is either disjoint from P or is contained in P.

Note that Theorem 0.14 is a technical tool that will be used to prove the main compactness theorem in the non-ULSC case, Theorem 0.12. In particular, Theorem 0.14 itself will be superseded by the stronger compactness theorems in the ULSC and non-ULSC cases, Theorems 0.9 and 0.12. This is because eventually we will know by the no mixing theorem that either $S_{neck} = \emptyset$ or $S_{ulsc} = \emptyset$, so that these cover all possible cases. Moreover, the assertions in Theorems 0.9 and 0.12 are stronger than those in Theorem 0.14.

After proving Theorem 0.14 in Part IV, we will be ready in Part V to prove the no mixing theorem, Theorem 0.4.

In Part VI, we will then complete the proof of Theorem 0.12. The main point left, which is not part of Theorem 0.14, is to prove that *every* leaf of the lamination \mathcal{L} in Theorem 0.12 is a plane. In contrast, Theorem 0.14 gives a

⁸The theorem holds also for sequences with fixed genus; see Part VII.



No structure for points in \mathcal{S}_{neck} .

Figure 4. Theorem 0.14: Limits of sequences of non-ULSC surfaces with curvature blowing up. The limit is a lamination of $\mathbf{R}^3 \setminus S$. The singular set S consists of two types of points — the ones in S_{neck} and the ones in S_{ulsc} . Note that the set S_{neck} is automatically closed, while the set S_{ulsc} is not. The set S_{ulsc} is a union of Lipschitz curves; the injectivity radius goes to zero at the endpoints of these curves, so these endpoints are in S_{neck} . Finally, the part of the lamination containing S_{ulsc} is foliated by planar leaves.

plane through each point of S_{neck} but does not claim that the leaves of \mathcal{L}' are planar.

Finally, since the no mixing theorem implies that Theorems 0.9 and 0.12 cover all cases, Theorem 0.6 will be a corollary of these two theorems.

We refer to the introduction of [CM04e] and the surveys [MP04], [P05], and [Ros03] for related results, including applications of the results of [CM04b]– [CM04e] as well as the results of this paper. There is much current research on minimal surfaces with infinite topology. Some of the results of the present paper were announced previously and have already been widely used to study infinite topology minimal surfaces; see, e.g., [MP04], [MPR04a], [MPR04b], [MPR], and [P05].

These results have more recently been used by Meeks, Perez, and Ros in their important paper on the uniqueness of the Riemann examples, [MPR15].

0.1. Brief outline of the paper and overview of the proofs. In Section 1, we will define the two notions of multi-valued graphs that will be needed to explain and prove the two main theorems.

Part I is devoted to recalling some of the earlier results for disks given in [CM04b]–[CM04e] and [CM08]. The first of these shows that embedded minimal disks are either graphs or are part of a double spiral staircase. The second result that we recall is the one-sided curvature estimate. Finally, we will recall the chord-arc bound for embedded minimal disks proven in [CM08].

In Part II, we will first define the singular set S and prove the convergence to the lamination \mathcal{L}' away from S. The rest of the part focuses on describing a neighborhood of each point in the ULSC singular set S_{ulsc} and the leaves

of \mathcal{L}' whose closure intersects \mathcal{S}_{ulsc} . A key point will be that the results of [CM04b]-[CM04e] for disks will give a sequence of multi-valued graphs in the Σ_j 's near each point $x \in \mathcal{S}_{ulsc}$. Moreover, these multi-valued graphs close up in the limit to give a leaf of \mathcal{L}' that extends smoothly across x. Such a leaf is said to be collapsed; in a neighborhood of x, the leaf can be thought of as a limit of double-valued graphs where the upper sheet collapses onto the lower. We will show that every collapsed leaf is stable,⁹ has at most two points of \mathcal{S}_{ulsc} in its closure, and these points are removable singularities. These results on collapsed leaves will be applied first in the USLC case in the next part and then later to get the structure of the ULSC regions of the limit in general, i.e., (C2) and (D) in Theorem 0.14.

In Part III, we prove Theorem 0.9, which gives the convergence of a ULSC sequence to a foliation by parallel planes away from two singular curves. Roughly speaking, there are two main steps to the proof:

- (1) Show that each collapsed leaf is in fact a plane punctured at two points of \mathcal{S} and, moreover, the sequence has the structure of a double spiral staircase near both of these points, with opposite orientations at the two points.
- (2) Show that leaves that are nearby a collapsed leaf of \mathcal{L}' are also planes punctured at two points of \mathcal{S} . (We call this "properness.")

In Part IV we consider general sequences of minimal surfaces that are neither necessarily ULSC nor with $S_{ulsc} = \emptyset$ and we prove the general compactness theorem, Theorem 0.14. Recall that this theorem asserts that the limit lamination \mathcal{L}' can be divided into two disjoint sub-laminations, one of which is the support of a region where (a subsequence of) the surfaces are ULSC and all of the results about ULSC sequences from Part III hold, such as the structure of the singular set and the multi-valued graphs structure. In the other region, curvature blowup comes exclusively from neck pinching and, thus, in this region there are no helicoid-like points. The key steps for proving the general structure theorem are the following:

- (1) Finding a *stable* plane through each point of S_{neck} . This plane will be a limit of a sequence of stable graphical annuli that lie in the complement of the surfaces.
- (2) Finding graphs in Σ_j that converge to a plane through each point of S_{neck} . To do this, we look in regions between consecutive necks and show that in any such region the surfaces are ULSC. The one-sided curvature estimate will then allow us to show that these regions are graphical.
- (3) Using (1) and (2) we then analyze the ULSC regions of a limit. That is, we show that if the closure of a leaf in \mathcal{L}' intersects \mathcal{S}_{ulsc} , then it has a

⁹More precisely, it is stable if it is orientable; otherwise, it has a stable double cover.

neighborhood that is ULSC. This will allow us to use the argument for the proof of Theorem 0.9 to get the same structure for such a neighborhood as we did in case where the entire surfaces were ULSC.

In Part V, we will use the structure obtained in Theorem 0.14 to show the no mixing theorem; Theorem 0.4. The key here is to show that if S_{ulsc} is nonempty, then S_{ulsc} cannot stop.

In Part VI, we will complete the proofs of Theorems 0.6 and 0.12. The only thing that remains to be proven is that *every* leaf of the lamination \mathcal{L}' is contained in a plane. We have already proven that the leaves of \mathcal{L}' are planes when the sequence is ULSC; thus, by the no mixing theorem, the only remaining case is when $\mathcal{S} = \mathcal{S}_{neck} \neq \emptyset$. We will divide the proof that the leaves of \mathcal{L}' are contained in planes into two cases, depending on whether or not the leaf is complete. In both cases, we will use a flux argument to rule out a nonflat leaf of \mathcal{L}' .

In Part VII, we describe the necessary changes to the main theorems and the modifications needed for their proofs when the sequence has positive genus.

Finally, we would like to thank the referee for an extremely careful reading and many useful comments that improved the paper.

1. Multi-valued graphs

To explain the theorems stated in the introduction and their proofs, we will need two notions of multi-valued graphs — namely, the one used in [CM04b]–[CM04e] and a generalization.

In [CM04b]–[CM04e], we defined multi-valued graphs as multi-sheeted covers of the punctured plane. To be precise, let D_r be the disk in the plane centered at the origin and of radius r, and let \mathcal{P} be the universal cover of the punctured plane $\mathbf{C} \setminus \{0\}$ with global polar coordinates (ρ, θ) so $\rho > 0$ and $\theta \in \mathbf{R}$. Given $0 \leq r \leq s$ and $\theta_1 \leq \theta_2$, define the "rectangle" $S_{r,s}^{\theta_1,\theta_2} \subset \mathcal{P}$ by

(1.1)
$$S_{r,s}^{\theta_1,\theta_2} = \{(\rho,\theta) \mid r \le \rho \le s, \, \theta_1 \le \theta \le \theta_2\}.$$

An N-valued graph of a function u on the annulus $D_s \setminus D_r$ is a single valued graph over (see Figure 5)

(1.2)
$$S_{r,s}^{-N\pi,N\pi} = \{(\rho,\theta) \, | \, r \le \rho \le s, \, |\theta| \le N \, \pi\}.$$

 $(\Sigma_{r,s}^{\theta_1,\theta_2}$ will denote the subgraph of Σ over the smaller rectangle $S_{r,s}^{\theta_1,\theta_2}$.) As in the earlier papers in the series, the multi-valued graphs that we will consider will never close up; in fact, they will all be embedded. Note that embedded corresponds to that the separation never vanishes. Here the separation w is the difference in height between consecutive sheets and is therefore given by

(1.3)
$$w(\rho,\theta) = u(\rho,\theta+2\pi) - u(\rho,\theta).$$

14



Figure 5. A multi-valued graph over the singly-punctured plane.

In the case where Σ is the helicoid (i.e., Σ can be parametrized by $(s \cos t, s \sin t, t)$ where $s, t \in \mathbf{R}$), then

(1.4)
$$\Sigma \setminus x_3 - \operatorname{axis} = \Sigma_1 \cup \Sigma_2,$$

where Σ_1 , Σ_2 are ∞ -valued graphs. Σ_1 is the graph of the function $u_1(\rho, \theta) = \theta$ and Σ_2 is the graph of the function $u_2(\rho, \theta) = \theta + \pi$. In either case the separation $w = 2\pi$.

Locally, the above multi-valued graphs give the complete picture for a ULSC sequence. However, the global picture can consist of several different multi-valued graphs glued together. To allow for this, we are forced to consider multi-valued graphs defined over the universal cover of $\mathbf{C} \setminus P$, where P is a discrete subset of the complex plane \mathbf{C} (see Figure 6). We will see that the bound on the genus implies that P consists of at most two points. The basic example of such a multi-valued graph comes from the family of minimal surfaces known as the Riemann examples.

Part I. Results for disks from [CM04b]–[CM04e]

The results for nonsimply connected minimal surfaces that are proven in this paper rely on the earlier results for disks given in [CM04b]–[CM04e]. For completeness and easy reference, we start by recalling those.

I.1. The lamination theorem and one-sided curvature estimate

The first theorem that we recall shows that embedded minimal disks are either graphs or are part of double spiral staircases; moreover, a sequence of such disks with curvature blowing up converges to a foliation by parallel planes away from a singular curve S. This theorem is modelled on rescalings Locally graphical except over two points; those points correspond to the two axes.



The spiral staircases around each of the axes connect to each other between the axes.

Figure 6. A multi-valued graph over the doubly-punctured plane. The spiral staircases near each puncture are oppositely-oriented.

of the helicoid, and the precise statement is as follows (we state the version for extrinsic balls; it was extended to intrinsic balls in [CM08]):

THEOREM I.1.1 (Theorem 0.1 in [CM04e]). Let $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ be a sequence of embedded minimal disks with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$. If

(I.1.2)
$$\sup_{B_1 \cap \Sigma_i} |A|^2 \to \infty$$

then there exists a subsequence, Σ_j , and a Lipschitz curve $\mathcal{S} : \mathbf{R} \to \mathbf{R}^3$ such that after a rotation of \mathbf{R}^3 ,

- (1) $x_3(\mathcal{S}(t)) = t$. (That is, \mathcal{S} is a graph over the x_3 -axis.)
- (2) Each Σ_j consists of exactly two multi-valued graphs away from S (which spiral together).
- (3) For each $1 > \alpha > 0$, $\Sigma_j \setminus S$ converges in the C^{α} -topology to the foliation, $\mathcal{F} = \{x_3 = t\}_t$, of \mathbf{R}^3 .
- (4) $\sup_{B_r(\mathcal{S}(t))\cap\Sigma_j} |A|^2 \to \infty$ for all $r > 0, t \in \mathbf{R}$. (The curvatures blow up along \mathcal{S} .)

The second theorem that we need to recall asserts that every embedded minimal disk lying above a plane, and coming close to the plane near the origin, is a graph. Precisely, this is the *intrinsic one-sided curvature estimate* that follows by combining [CM04e] and [CM08]:

THEOREM I.1.3. There exists $\varepsilon > 0$ so that if

$$(I.1.4) \qquad \qquad \Sigma \subset \{x_3 > 0\} \subset \mathbf{R}^3$$

is an embedded minimal disk with the intrinsic ball $\mathcal{B}_{2R}(x) \subset \Sigma \setminus \partial \Sigma$ and $|x| < \varepsilon R$, then

(I.1.5)
$$\sup_{\mathcal{B}_R(x)} |A_{\Sigma}|^2 \le R^{-2}.$$

Theorem I.1.3 is in part used to prove the regularity of the singular set where the curvature is blowing up.

Note that the assumption in Theorem I.1.1 that the surfaces are disks is crucial and cannot even be replaced by assuming that the sequence is ULSC. To see this, observe that one can choose a one-parameter family of Riemann examples that is ULSC but where the singular set S is given by a *pair* of vertical lines. Likewise, the assumption in Theorem I.1.3 that Σ is simply connected is crucial, as can be seen from the example of a rescaled catenoid; see (0.11). Under rescalings the catenoid converges (with multiplicity two) to the flat plane. Thus a neighborhood of the neck can be scaled arbitrarily close to a plane but the curvature along the neck becomes unbounded as it gets closer to the plane. Likewise, by considering the universal cover of the catenoid, one sees that embedded, and not just immersed, is needed in Theorem I.1.3.

Finally, we recall the chord-arc bound for embedded minimal disks proven in Theorem 0.5 of [CM08]:

THEOREM I.1.6. [CM08]. There exists a constant C > 0 so that if $\Sigma \subset \mathbf{R}^3$ is an embedded minimal disk, $\mathcal{B}_{2R} = \mathcal{B}_{2R}(0)$ is an intrinsic ball in $\Sigma \setminus \partial \Sigma$ of radius 2R, and $\sup_{\mathcal{B}_{r_0}} |A|^2 > r_0^{-2}$ where $R > r_0$, then for $x \in \mathcal{B}_R$, the intrinsic distance is bounded from above by the extrinsic distance as follows:

(I.1.7) $C \operatorname{dist}_{\Sigma}(x,0) < |x| + r_0.$

Part II. The singular set S and limit lamination \mathcal{L}'

The three main results of this part are the convergence to the lamination \mathcal{L}' away from a singular set \mathcal{S} , the description of a neighborhood of each ULSC singular point, and the description of the leaves of \mathcal{L}' whose closure intersects \mathcal{S}_{ulsc} . We will explain these in a bit more detail next.

We start by defining the singular set S; roughly speaking, S is the set of points where the curvature blows up (see Definition/Lemma II.1.1). The definition of S will immediately imply that S is a closed subset of \mathbf{R}^3 . We next show that in the open subset $\mathbf{R}^3 \setminus S$, a subsequence of the sequence of embedded minimal surfaces converges to a minimal lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$ (see Lemma II.1.2). The results of [CM04b]–[CM04e] give a precise description of a neighborhood of each point in S_{ulsc} . Namely, for j large, Σ_j must be a double-spiral staircase near each point in S_{ulsc} and the set S_{ulsc} must satisfy a local cone property that gives the regularity of the set. The description near a singular point and local cone property are given in Lemma II.2.3. We also recall in Lemma II.2.3 that, as $j \to \infty$, this sequence of double-spiral staircases near a singular point x closes up in the limit to give a leaf of \mathcal{L}' that extends smoothly across x. We will say that such a leaf is collapsed; in a neighborhood of x, the leaf can be thought of as a limit of double-valued graphs where the upper sheet collapses onto the lower.

Finally, we will show that every collapsed leaf is stable, has at most two points of S_{ulsc} in its closure, and these points are removable singularities. The key for proving stability is to use the separations of the limiting multi-valued graphs to construct a positive Jacobi field in the limit. The limit Jacobi field is not *a priori* well defined, but is instead well defined on a covering space of the collapsed leaf. However, we show in Appendix A that stability of a covering space implies stability of the surface itself as long as the covering space has sub-exponential area growth. We apply this to show that every collapsed leaf is stable. We will also use the fact that the surfaces Σ_j have bounded genus to show that each collapsed leaf has at most two points of S_{ulsc} in its closure.

These results on collapsed leaves will be applied first in the USLC case in the next part and then later to get the structure of the ULSC regions of the limit in general, i.e., (C2) and (D) in Theorem 0.14.

II.1. The singular set S

To define the singular set, recall from [CM04e] that for any sequence of surfaces (minimal or not) in \mathbb{R}^3 , after possibly going to a subsequence, then there is a well-defined notion of points in \mathbb{R}^3 where the second fundamental form of the sequence blows up. The set of such points will be referred to below as the singular set S and is given by an elementary and straightforward compactness argument.

DEFINITION/LEMMA II.1.1 (The singular set; Lemma I.1.4 in [CM04e]). Let $\Sigma_i \subset B_{R_i}$ with $\partial \Sigma_i \subset \partial B_{R_i}$ and $R_i \to \infty$ be a sequence of (smooth) compact surfaces. After passing to a subsequence, Σ_j , we may assume that for each $x \in \mathbf{R}^3$, either of the two following properties holds:

- $\sup_{B_r(x)\cap\Sigma_j} |A|^2 \to \infty$ for all r > 0. (The set of such points x will be denoted by S.)
- $\sup_{i} \sup_{B_r(x) \cap \Sigma_i} |A|^2 < \infty$ for some r > 0.

II.1.1. Convergence away from S. The first result that we will need is that in the open subset $\mathbb{R}^3 \setminus S$, a sequence of embedded minimal surfaces has

a subsequence that converges to a minimal lamination \mathcal{L}' of $\mathbb{R}^3 \setminus \mathcal{S}$. This is an easy consequence of that the curvature is bounded on compact sets in the complement of \mathcal{S} and is proven in the next lemma.

LEMMA II.1.2. Suppose that Σ_j and S are as in Lemma II.1.1. If in addition the Σ_j 's are minimal and embedded, then there exists a subsequence (still denoted by Σ_j) and a lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$ so that the following hold:

- $\Sigma_j \to \mathcal{L}'$ on compact subsets of $\mathbf{R}^3 \setminus \mathcal{S}$.
- The leaves of \mathcal{L}' are minimal.

Proof. For each compact subset K of $\mathbb{R}^3 \setminus S$, then Lemma II.1.1 gives an open covering of K by finitely many balls where the curvatures of the Σ_j 's are bounded (independent of j) in the concentric double balls. Both claims now follow from Proposition B.1 in [CM04e] and a diagonal argument. \Box

As in [CM04e], convergence to \mathcal{L}' in the above lemma means that if we think of the embedded surfaces Σ_j themselves as laminations, then the coordinate charts for these laminations converge in the C^{α} -norm for any $\alpha < 1$ and the leaves converge as sets. The convergence is actually C^{∞} -tangentially, meaning that if we write a leaf locally as a graph, then a sequence of local graphs in Σ_j converges smoothly to this leaf. This tangential regularity follows from the C^{α} -convergence and elliptic estimates. However, easy examples show that the convergence in the transversal direction may only be in the Lipschitz topology; cf. [Sol86].

Throughout the rest of this paper, we will assume that $\Sigma_j \subset B_{R_j}$ with $\partial \Sigma_j \subset \partial B_{R_j}$ and $R_j \to \infty$ is a sequence of (smooth) compact embedded minimal surfaces that converges off of a singular set S to a lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$ with minimal leaves. The lamination \mathcal{L}' is given by Lemma II.1.2. In order to obtain additional structure of S and \mathcal{L}' , we will need to also make topological assumptions about the surfaces Σ_j . We will always assume that the Σ_j 's have bounded genus. In Part III, we will assume that the surfaces Σ_j are ULSC, i.e, that $S_{\text{neck}} = \emptyset$; in Part IV, we will consider the other case where $S_{\text{neck}} \neq \emptyset$.

II.2. The local structure of \mathcal{L}' near a point of \mathcal{S}_{ulsc}

We will eventually show that all of the leaves of the lamination \mathcal{L}' are flat (see Theorem 0.6), but we will need to first establish some initial structure of \mathcal{L}' . The first step will be accomplished in this section where we describe the local structure of \mathcal{L}' near a point in \mathcal{S}_{ulsc} .

The next lemma is going to show that each ULSC singular point lies in the closure of a leaf of \mathcal{L}' that extends smoothly across the singular point and, furthermore, ULSC singular points are leaf-wise isolated and they satisfy a local cone property. To state this cone property, let $\mathbf{C}_{\delta}(z)$ be the (convex) double cone with vertex z, cone angle $(\pi/2 - \arctan \delta)$, and axis parallel to the x_3 -axis. That is (see Figure 7),

(II.2.1)
$$\mathbf{C}_{\delta}(z) = \{x \in \mathbf{R}^3 \mid (x_3 - z_3)^2 \ge \delta^2 \left((x_1 - z_1)^2 + (x_2 - z_2)^2 \right) \}.$$

The local cone property is now defined as follows. Given $\delta > 0$ and $r_0 > 0$, we will say that a subset $S_{\text{ulsc}} \subset \mathbf{R}^3$ has the *local cone property* if S_{ulsc} is nonempty and

As in [CM04e], we will see in Section III.2 that this local cone property directly gives Lipschitz regularity of the subset S_{ulsc} .



Figure 7. It follows from the one-sided curvature estimate that the ULSC singular set S_{ulsc} has the local cone property and, as we will see, this gives Lipschitz regularity.

We can now state the lemma that gives the regularity of the leaves through S_{ulsc} and the local cone property for S_{ulsc} . (For a surface Γ , \mathbf{n}_{Γ} is its Gauss map.)

LEMMA II.2.3. Given a point $x \in S_{\text{ulsc}}$, there exists $r_0 > 0$ so that $B_{r_0}(x) \cap \mathcal{L}'$ has a component Γ_x whose closure $\overline{\Gamma_x}$ is a smooth minimal graph containing xand with boundary in $\partial B_{r_0}(x)$ (so x is a removable singularity for Γ_x).

Furthermore, $\overline{\Gamma_x} \cap S = \{x\}$ and, after rotating \mathbf{R}^3 so that $\mathbf{n}_{\overline{\Gamma_x}}(x) = (0,0,1)$, the set S_{ulsc} satisfies the local cone property (II.2.2) for some $\delta > 0$ and the above r_0 . The rotation may vary with x, but the dependence is Lipschitz.

Proof. For simplicity, translate so that x = 0. Since $0 \notin S_{\text{neck}}$, there exists some $r_0 > 0$ so that the components of $B_{r_0}(0) \cap \Sigma_j$ are disks for every j; cf. (0.1).

The first two properties follow immediately from Theorem 5.8 in [CM04c]. (This theorem combines the existence of multi-valued graphs near a blow up point and the sublinear growth of the separation.) Namely, since $0 \in S$, we first get a sequence of points $y_j \in \Sigma_j$ with $|A|^2(y_j) \to \infty$ and $y_j \to 0$. Since the component of $B_{r_0}(0) \cap \Sigma_j$ containing y_j is a disk, Theorem 5.8 in [CM04c] then gives the following two properties:

- There is a rotation of \mathbf{R}^3 and a subsequence so that Σ_j contains a 2-valued minimal graph $\Sigma_{d,j} \subset \Sigma_j$ defined over an annulus $D_{r_0/C} \setminus D_{r_j}$ where $r_j \to 0$.
- As $j \to \infty$, the 2-valued graphs close up in the limit to converge with multiplicity two to a graph Γ_x over $D_{r_0/C} \setminus \{0\}$ with $x \in \overline{\Gamma_x}$.

Since any subsequence of a convergent sequence has the same limit, we conclude that Γ_x is contained in a leaf of \mathcal{L}' . Finally, x is a removable singularity for Γ_x by a standard removable singularity result for minimal graphs.

The cone property follows easily from Corollary I.1.9 in [CM04e], which gives a constant $\delta_0 > 0$ so that if $B_{2R} \cap \Sigma_j$ contains a 2-valued graph in $\{x_3^2 \leq \delta_0^2 (x_1^2 + x_2^2)\}$ over $D_R \setminus D_{r_j}$ and with gradient $\leq \delta_0$, then each component of

(II.2.4)
$$B_{R/2} \cap \Sigma_j \setminus (\mathbf{C}_{\delta_0}(0) \cup B_{2r_j})$$

is a multi-valued graph with gradient ≤ 1 . After possibly shrinking the radius above given by Theorem 5.8 in [CM04c], we can assume that Γ_x is a graph with small gradient, and hence Corollary I.1.9 in [CM04e] applies. It follows that

(II.2.5)
$$B_{r_0}(x) \cap \mathcal{S}_{\text{ulsc}} \subset \mathbf{C}_{\delta_0}(x).$$

Finally, the embeddedness of the Σ_j 's implies that two limit minimal graphs through nearby singular points must be disjoint. It is now easy to see that the map that takes a singular point y to the tangent plane of the limit minimal graph through y is Lipschitz, giving the last claim.

Lemma II.2.3 shows that each point $x \in S_{\text{ulsc}}$ is a removable singularity for a component Γ_x of $B_{r_0}(x) \cap \mathcal{L}'$ for some $r_0 > 0$. Furthermore, the local cone property implies that the intersection of $B_{r_0}(x) \cap \Sigma_j$ with the complement of (a tubular neighborhood of) a cone $\mathbf{C}_{\delta'}(x)$ (for some $\delta' > 0$) consists of two multi-valued graphs for j large. (The fact that there are exactly two is established in Proposition II.1.3 in [CM04e].) However, it is worth noting that these two properties alone *do not* imply that x is a removable singularity for the lamination \mathcal{L}' , but rather there are two possibilities:

- (P) The multi-valued graphs in the complement of the cone $\mathbf{C}_{\delta'}(x)$ close up in the limit.
- (N-P) These multi-valued graphs converge to a collection of graphs (such as Γ_x) and *at least one* multi-valued graph that spirals infinitely¹⁰ on one side of Γ_x .

¹⁰It may happen that a multi-valued graph spirals into Γ_x as in the example in [CM04a]. It is also possible that there is a gap between Γ_x and the spiraling multi-valued graph (e.g., if it spirals into a different graphical component as in examples constructed in [HW11] and [Kle12]).

In the first case (P) (we will call this "properness" below), the sequence converges to a foliation in a neighborhood of x. The second case (N-P) ("not proper") is illustrated in [CM04a] by a sequence of embedded minimal disks Σ_i in the unit ball B_1 with $\partial \Sigma_i \subset \partial B_1$ where the curvatures blow up only at 0 and

(II.2.6)
$$\Sigma_i \setminus \{x_3 = 0\}$$

converges to two embedded minimal disks

$$(\text{II.2.7})\qquad \qquad \Sigma^- \subset \{x_3 < 0\}.$$

 $(\text{II.2.8})\qquad \qquad \Sigma^+ \subset \{x_3 > 0\},$

each of which spirals into $\{x_3 = 0\}$ and thus is not proper. Thus, in the example from [CM04a], 0 is the first, last, and only point in S_{ulsc} and the limit lamination consists of three leaves: Σ^+ , Σ^- , and the punctured unit disk $B_1 \cap \{x_3 = 0\} \setminus \{0\}$. In this example of (N-P), the limit lamination cannot be extended smoothly to any neighborhood of 0.

To summarize, [CM04a] shows that (N-P) can occur for a sequence of disks $\Sigma_i \subset B_{R_i}$ with $\partial \Sigma_i \subset \partial B_{R_i}$; however, [CM04e] shows that (N-P) cannot occur for disks if the radii R_i go to infinity.

II.2.1. Collapsed leaves of \mathcal{L}' . One of the difficulties is that the leaves of the lamination \mathcal{L}' may not be complete; this occurs at points of \mathcal{S} . We will begin by analyzing a particular type of incomplete leaf that we will call collapsed.

To define this, note that Lemma II.2.3 shows that each point $x \in S_{\text{ulsc}}$ is a removable singularity for a component Γ_x of $B_{r_0}(x) \cap \mathcal{L}'$. We will say that the leaf Γ of \mathcal{L}' containing Γ_x is collapsed:

Definition II.2.9. A leaf Γ of \mathcal{L}' is collapsed if there exists some $x \in S_{\text{ulsc}}$ so that Γ contains the local leaf Γ_x given by Lemma II.2.3.

It follows from Lemma II.2.3 that every collapsed leaf is a limit leaf of \mathcal{L}' . For a sequence of rescaled helicoids converging to a foliation by parallel planes away from an axis, every leaf is collapsed. We will eventually show that every leaf of \mathcal{L}' whose closure contains a point of \mathcal{S}_{ulsc} is collapsed. However, it is worth pointing out that this is not obvious. For example, in case (N-P) of the previous section, we get leaves of \mathcal{L}' that spiral infinitely into the collapsed leaf but are not themselves collapsed. (We will eventually rule out this possibility using that the sequence of outer radii is going to infinity.)

We will describe the structure of the collapsed leaves in the rest of this part. It is useful to first define the closure Γ_{Clos} of a leaf Γ of \mathcal{L}' to be the union of the closures of all bounded (intrinsic) geodesic balls in Γ ; that is, we

22

fix a point $x_{\Gamma} \in \Gamma$ and set

(II.2.10)
$$\Gamma_{\text{Clos}} = \bigcup_{r} \overline{\mathcal{B}_{r}(x_{\Gamma})},$$

where $\mathcal{B}_r(x_{\Gamma})$ is the closure of $\mathcal{B}_r(x_{\Gamma})$ as a subset of \mathbf{R}^3 .

Clearly, a leaf Γ is complete if and only if $\Gamma_{\text{Clos}} = \Gamma$ and we always have that

(II.2.11)
$$\Gamma_{\text{Clos}} \setminus \Gamma \subset \mathcal{S}.$$

The incomplete leaves of \mathcal{L}' can be divided into several types, depending on how Γ_{Clos} intersects \mathcal{S} :

- Collapsed leaves, defined in Definition II.2.9, where $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ contains a removable singularity for Γ .
- Leaves Γ with $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}} \neq \emptyset$, but where Γ does not have a removable singularity. This would occur, for example, if Γ spirals infinitely into the collapsed leaf through $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ as in (N-P). (We will eventually show that this does not occur.)
- Leaves Γ where $\Gamma_{\text{Clos}} \setminus \Gamma \subset S_{\text{neck}}$; these will not be considered until Part IV.

II.3. The structure of the collapsed leaves of \mathcal{L}'

In the rest of this part, we will describe the structure of the collapsed leaves of \mathcal{L}' defined in Definition II.2.9. The most important properties of a collapsed leaf Γ are given in Proposition II.3.1 below, which describes a neighborhood of the points of \mathcal{S}_{ulsc} in Γ . The proposition shows that such a Γ is stable and that the closure of Γ intersects \mathcal{S}_{ulsc} in at most two points. These results apply without additional assumptions on the sequence Σ_j ; we will see in the next part that Γ has more structure when we assume, in addition, that the sequence is ULSC.

The next proposition establishes the key properties of a collapsed leaf in the general case.

PROPOSITION II.3.1. Each collapsed leaf Γ of \mathcal{L}' has the following properties:

(1) Given any $y \in \Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$, there exists $r_0 > 0$ so that the closure (in \mathbb{R}^3) of each component of $B_{r_0}(y) \cap \Gamma$ is a compact embedded disk with boundary in $\partial B_{r_0}(y)$.

Furthermore, $B_{r_0}(y) \cap \Gamma$ must contain the component Γ_y given by Lemma II.2.3 and Γ_y is the only component of $B_{r_0}(y) \cap \Gamma$ with y in its closure.

(2) If Γ is oriented, then it is stable. (Otherwise, its oriented double cover is stable.)

(3) Γ_{Clos} intersects S_{ulsc} in at most two points. If $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ contains two points, then the multi-valued graphs in the Σ_j 's spiral in opposite directions around the two corresponding axes (see Figure 8).

A double-valued graph in Σ_j .



Figure 8. The multi-valued graph converging to Γ in Proposition II.3.1.

Properties (2) and (3) in Proposition II.3.1 are self-explanatory. However, to appreciate property (1), it may be useful to observe one implication of (1) and to also see an example of what it rules out. First, (1) implies that $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ consists of a discrete set of points and each of these points is a removable singularity. Second, recall from (N-P) — "not proper" — that a priori there may be multi-valued graphs in $B_{r_0}(y) \cap \mathcal{L}'$ that spiral infinitely into Γ_y ; (1) above says that these "infinite spirals" are not contained in any collapsed leaf.

Throughout this section Γ will be a collapsed leaf of \mathcal{L}' . By definition, a leaf Γ is a connected open surface but may not be complete. (And, in fact, collapsed leaves are incomplete by definition.) We will let $K \subset \Gamma$ denote a connected open subset with compact closure in Γ . Finally, $T(K, \varepsilon)$ is the ε -tubular normal neighborhood of K; i.e.,

(II.3.2)
$$T(K,\varepsilon) = \{x + s \mathbf{n}_{\Gamma}(x) \mid x \in K, |s| < \varepsilon\}.$$

II.3.1. Proving property (1) of Proposition II.3.1: Isolated removable singularities. To prove (1) of Proposition II.3.1, we will show the following claim:

CLAIM. If $x \in S_{\text{ulsc}}$ is a singular point in the closure Γ_{Clos} of a collapsed leaf Γ , then there exists $r_0 > 0$ so that the component of $B_{r_0}(x) \cap \overline{\Gamma}$ containing x is the one from Lemma II.2.3.

Property (1) then follows from the following two properties of the component Γ_x from Lemma II.2.3:

- $\Gamma_x \cup \{x\}$ is a smooth minimal surface.
- $\overline{\Gamma_x} \cap \mathcal{S} = \{x\}.$

Thus, we will have at the same time shown that each ULSC singular point in the closure of Γ is a removable singularity and has a neighborhood *in* Γ where there are no other singular points, as desired.

Lemma II.3.3 below establishes the above claim about the singular points in the closure of a collapsed leaf. The lemma is best illustrated using the "not proper" example in (N-P). In (N-P), a sequence of embedded minimal disks converges in $B_1 \setminus \{0\}$ to a lamination with three leaves: the punctured disk $\Gamma_0 = D_1 \setminus \{0\}, \Gamma^+$ spiralling into Γ_0 infinitely from above, and Γ^- spiralling into Γ_0 infinitely from below; see Figure 9. Notice that all three leaves contain 0 in their closure. The leaf Γ_0 is collapsed at 0 (so 0 is a removable singularity for Γ_0), but Γ^+ and Γ^- cannot extend past the singularity 0. The conclusion of Lemma II.3.3 is that Γ^+ and Γ^- cannot be contained in any collapsed leaf of \mathcal{L}' .



The intersection of Γ^+ with a vertical line gives a discrete infinite set of points limiting down to a point in Γ_0 .

Figure 9. The "not proper" example (N-P): The sequence of disks Σ_j converges in $B_1 \setminus \{0\}$ to a lamination with three leaves: the punctured disk $\Gamma_0 = D_1 \setminus \{0\}$ (dotted), Γ^+ spiralling into Γ_0 infinitely from above, and Γ^- spiralling into Γ_0 infinitely from below (not pictured). The collapsed leaf Γ_0 is not discrete but Γ^+ and Γ^- are.

The above example from (N-P) also serves to illustrate the idea of the proof of Lemma II.3.3. Namely, a key distinction between the collapsed leaf Γ_0 versus Γ^+ and Γ^- is that Γ^+ and Γ^- are not limit leaves; rather, they are discrete in the following sense:

Given any point y in Γ^+ or Γ^- , there exists s > 0 so that $B_s(y) \cap \mathcal{L}'$ has only one connected component (i.e., the one containing y).

On the other hand, since Γ^+ and Γ^- spiral infinitely into Γ_0 , the leaf Γ_0 is not discrete in this sense. Likewise, the description of a neighborhood of a point in S_{ulsc} shows that a collapsed leaf is never discrete.

LEMMA II.3.3. Suppose that $x \in S_{\text{ulsc}}$ and Γ' is a component of $B_{r_0}(x) \cap \mathcal{L}'$ with x in its closure $\overline{\Gamma'}$. If Γ' is contained in a collapsed leaf of \mathcal{L}' , then Γ' must be the component Γ_x given by Lemma II.2.3.

Proof. Since the component Γ' of $B_{r_0}(x) \cap \mathcal{L}'$ contains the point $x \in \mathcal{S}_{ulsc}$ in its closure, embeddedness and the cone property imply that Γ' has one of the following two properties:

- (L1) Γ' is the component Γ_x given by Lemma II.2.3 and hence extends smoothly across x.
- (L2) Γ' is not the component Γ_x given by Lemma II.2.3.

In Lemma II.3.4 below we will prove that the leaves satisfying (L2) are discrete in the following sense:

Given any point y in a leaf of \mathcal{L}' satisfying (L2), there exists s > 0so that $B_s(y) \cap \mathcal{L}'$ has only one connected component (i.e., the one containing y).

Completing the proof assuming discreteness. Suppose now that Γ is collapsed, $y \in \Gamma$, and the ball $B_s(y)$ is disjoint from \mathcal{S} . Let $\Gamma_{y,s}$ be the component of $B_s(y) \cap \Gamma$ containing y. It follows from the Harnack inequality (since the curvature is locally bounded on Γ) that $\Gamma_{y,s}$ is the limit of distinct leaves of $B_s(y) \cap \mathcal{L}'$. In particular, Γ is not discrete and hence does not contain any leaves of $B_{r_0}(x) \cap \mathcal{L}'$ that satisfy (L2). This completes the proof of the lemma modulo Lemma II.3.4 below.

The next lemma shows that we always get discreteness for leaves of \mathcal{L}' that have a point of \mathcal{S}_{ulsc} in their closure but are not collapsed at this point. (Cf. the picture for Γ^+ and Γ^- in Figure 9.)

LEMMA II.3.4. Given any point y in a leaf of \mathcal{L}' satisfying (L2), there exists s > 0 so that $B_s(y) \cap \mathcal{L}'$ has only one connected component (i.e., the one containing y).

Proof. Suppose that a component Γ' of $B_{r_0}(x) \cap \mathcal{L}'$ contains the point $x \in S_{\text{ulsc}}$ in its closure but is *not* equal to Γ_x . It suffices to find one point in Γ' where the leaf is locally discrete. (Since the leaf is connected, the Harnack inequality then implies that every point is discrete.) We will next outline the argument to find this discrete point. The key will be to find a sequence

of curves $\gamma_j \subset \Sigma_j$ with uniformly bounded length, where one sequence of endpoints converges to a point in Γ' , the Σ_j 's are *uniformly* discrete at the second endpoint of γ_j , and the γ_j 's stay away from the singular set S. These properties are made precise in (G1)–(G4) below; see Figure 10. Since the γ_j 's stay away from S and have bounded length, a subsequence of the γ_j 's will converge to a curve γ in some leaf of \mathcal{L}' . However, one sequence of endpoints converges to a point in Γ' and so the whole curve γ is in Γ' . Finally, the second endpoint of γ will give the desired discrete point in Γ' .



The (punctured) graph Γ_x in $B_{r_0}(x) \cap \mathcal{L}'$.

Figure	10.	The	curves	γ_j	in	Σ_j .
--------	-----	-----	--------	------------	----	--------------

Before making this precise, we need a few simple preliminaries. First, since $\Gamma_x \cup \{x\}$ separates the ball $B_{r_0}(x)$ and $\Gamma' \subset B_{r_0}(x) \setminus (\Gamma_x \cup \{x\})$ is connected, we may assume that Γ' is contained in the component $B_{r_0}^+(x)$ of $B_{r_0}(x) \setminus (\Gamma_x \cup \{x\})$ that is above Γ_x . Second, after shrinking r_0 , we can assume that

(II.3.5)
$$B_{r_0}^+(x) \cap \mathcal{S} = \emptyset.$$

Namely, there would otherwise be a sequence of points in S approaching x from above with corresponding separating graphs converging to Γ_x from above, but this would contradict the spiraling.

As mentioned, the key point is to find a sequence of curves γ_j parametrized by arclength

(II.3.6)
$$\gamma_j : [0, \ell_j] \to B^+_{r_0}(x) \cap \Sigma_j$$

with the following properties (see Figure 10):

- (G1) The endpoints $\gamma_i(0)$ converge to a point $y \in \Gamma'$.
- (G2) The lengths ℓ_j are uniformly bounded; i.e., $\ell_j \leq \ell$ for every j.
- (G3) The minimal distance between γ_i and \mathcal{S} is at least $\mu > 0$.
- (G4) The Σ_j 's are "uniformly discrete" at the endpoint $\gamma_j(\ell_j)$; precisely, there exists $\delta > 0$ so that $B_{\delta}(\gamma_j(\ell_j)) \cap \Sigma_j$ is a (connected) graph over its tangent plane at $\gamma_i(\ell_j)$ with gradient bounded by one.

The discreteness follows immediately from (G1)–(G4). Namely, (G2) and (G3) imply that a subsequence of the curves γ_j converges to a curve γ contained in a leaf of \mathcal{L}' . Since the endpoints $\gamma_j(0)$ converge to the point y in the leaf Γ' , the entire curve γ must be contained in Γ' . Finally, (G4) implies that Γ' is discrete at the second endpoint of γ and, hence, discrete everywhere by the Harnack inequality.

Before establishing (G1)-(G4), we need to recall the following two additional facts:

(1) Existence of nearby points of large curvature. Given any constants C_1 and C_4 , there exists $\varepsilon > 0$ so that for any s > 0 and every j sufficiently large (depending also on s), there is a point

(II.3.7)
$$q_j \in B_s(x) \cap \Sigma_j \setminus B_{\varepsilon s}(x)$$

so that q_i is above $\Gamma_x \cup \{x\}$ and q_i satisfies

(II.3.8)
$$|A|^2(q_j) \ge C_4 C_1 |x - q_j|^{-2}.$$

(2) Curvature bound away from x. Given any $\mu > 0$, there exists a constant C_2 so that if y is any point in $B_{r_0}^+(x) \cap \Sigma_j \setminus B_{\mu}(x)$, then

(II.3.9)
$$|A|^2(y) \le C_2$$

Property (1) was proven in Corollary III.3.5 in [CM04d]. Property (2) follows easily since the singular set S does not intersect $B_{r_0}^+(x)$ by (II.3.5). (The proof of (2) can be made precise using Lemma II.1.1 and a covering argument.)

To complete the proof of discreteness, it suffices to establish (G1)–(G4). We will do this next. First, fix a point $y \in \Gamma'$ in a small ball $B_h(x)$ about x. (*h* will need to be sufficiently small relative to r_0 but otherwise does not matter.) Since $y \in \Gamma'$, we can choose a sequence of points $y_j \in \Sigma_j$ that converge to y. Now choose a constant s > 0 with s much smaller than |y - x|.

Observe that property (1) gives points $q_j \in B_s(x) \setminus B_{\varepsilon s}(x)$ in Σ_j satisfying (II.3.8). A simple blow up argument (e.g., Lemma 5.1 in [CM04c]) then gives points $p_j \in \Sigma_j$ near q_j and radii r_j so that

(II.3.10)
$$\sup_{B_{r_j}(p_j)\cap\Sigma_j} |A|^2 \le 4 |A|^2(p_j) = 4 C_1 r_j^{-2},$$

and

(II.3.11)
$$B_{r_j}(p_j) \subset B_{\frac{2|x-q_j|}{\sqrt{C_4}}}(q_j)$$

In particular, by taking C_4 large in property (1), we can assume that the ratio

(II.3.12)
$$\frac{r_j}{|p_j - x|}$$

28

is as small as we want and, hence, also that

(II.3.13)
$$B_{r_i}(p_j) \subset B_{2s}(x) \setminus B_{\varepsilon s/2}(x)$$

We called the pair (p_j, r_j) a blow up pair in [CM04e]. The point about such a pair is that Theorem 0.7 in [CM04e] gives multi-valued graphs

(II.3.14)
$$\Sigma_j^g \subset \Sigma_j$$

defined outside of a disk of radius r_j centered at p_j and whose initial separation is proportional to r_j . On the other hand, since $p_j \notin B_{\varepsilon s/2}(x)$, property (2) implies that there is a uniform upper bound for $|A|^2(p_j)$ — and, thus, a uniform lower bound for the initial scale r_j .

We will also need a positive lower bound for the minimum distance between Σ_j^g and x. The argument for this is very similar to an argument in Section III.2 of [CM04e]. We will sketch the argument next. The lower bound follows easily once we have a lower bound for the distance from $B_{r_j}(p_j)$ to Γ_x . Since $p_j \notin B_{\varepsilon s/2}(x)$, the one-sided curvature estimate gives a lower bound for the distance from p_j to Γ_x . (Otherwise p_j would lie in a narrow cone about Γ_x and the one-sided curvature estimate would contradict (II.3.10).) Using this and the fact that r_j is small relative to $|p_j - x|$ (see (II.3.12)) then gives the desired lower bound for the distance from $B_{r_j}(p_j)$ to Γ_x . We leave the details to the reader.

To summarize, we have established a positive lower bound for the distance from Σ_j^g to x and for the initial scale r_j . This lower bound on the initial scale also implies a lower bound for the separation between the sheets of $\Sigma_j^{g,11}$ Moreover, Proposition II.1.3 in [CM04e] says that Σ_j contains exactly two (oppositely-oriented) multi-valued graphs in this region; the uniform curvature upper bound given by (2) then also implies a uniform lower bound for the distance between these two multi-valued graphs.

As a consequence of these uniform bounds, a (sub) sequence of the two (oppositely-oriented) multi-valued graphs is guaranteed to converge (with multiplicity one) to two multi-valued graphs in $B_{r_0}^+(x) \cap \mathcal{L}'$, and these limit multivalued graphs will satisfy the same lower bounds. Fix a point z in one of the limit multi-valued graphs. We will now find the desired curves γ_j from y_j to points converging to z. ((G1) and (G4) will then automatically be satisfied.) We use two facts to find these curves. First, the chord-arc bound of Theorem I.1.6 allows us to connect y_j to the multi-valued graph Σ_j^g by a curve $\gamma_i^+ \subset \Sigma_j$ with length at most $C_3 h$ and, furthermore, we can assume that γ_i^+

¹¹The existence of some lower bound is easy and almost obvious; a fairly sharp lower bound is proven in Lemma III.1.6 in [CM04e].

is above Σ_j^{g} .¹² Now that γ_j^+ connects y_j to the multi-valued graph, we can use a curve γ_j^g in Σ_j^g to connect the endpoint of γ_j^+ to points $z_j \in \Sigma_j^g$ converging to z; the γ_j^g 's automatically have uniformly bounded length and also stay uniformly away from x. This completes the proof of (G1)–(G4) and, consequently, also completes the proof of discreteness.

II.3.2. Each leaf is a limit of multi-valued graphs in the Σ_j 's. Recall that, throughout this section, Γ is a leaf of \mathcal{L}' and $K \subset \Gamma$ is a connected open subset that has compact closure in Γ .

We will first show in Lemma II.3.15 that the Σ_j 's are locally graphical over Γ in a tubular neighborhood of K. Corollary II.3.18 uses the local description of Lemma II.3.15 to construct multi-valued graphs $\Sigma_j^g \subset \Sigma_j$ converging to K. Both Lemma II.3.15 and Corollary II.3.18 apply to any leaf Γ and do not require Γ to be collapsed.

The next lemma shows that Σ_j is locally graphical over Γ in a small tubular neighborhood of K.

LEMMA II.3.15. Given any $\delta > 0$, there exist $\varepsilon > 0$ and J so that if j > Jand $x \in T(K, \varepsilon) \cap \Sigma_j$, then $\mathcal{B}_{\varepsilon}(x) \subset \Sigma_j$ is a graph over (a subset of) Γ with gradient bounded by δ .

Proof. Since Γ is a leaf of \mathcal{L}' , it is disjoint from the singular set \mathcal{S} . Therefore, for each point $y \in \Gamma$, the convergence of the Σ_j 's to the lamination \mathcal{L}' away from \mathcal{S} gives a ball $B_{\varepsilon_y}(y)$ and a J_y so that if $x \in B_{\varepsilon_y}(y) \cap \Sigma_j$ for $j > J_y$, then $\mathcal{B}_{\varepsilon_y}(x) \subset \Sigma_j$ is a graph over (a subset of) Γ with gradient bounded by δ .

However, the closure \overline{K} of K in Γ is compact, so it can be covered by a finite subcollection of the half-balls, i.e.,

(II.3.16)
$$K \subset \bigcup_{i=1}^{m} B_{\frac{\varepsilon y_i}{2}}(y_i).$$

It is then easy to see that this implies the lemma with

(II.3.17)
$$\varepsilon = 1/2 \min \varepsilon_{y_i}.$$

The next corollary uses Lemma II.3.15 to get multi-valued graphs $\Sigma_j^g \subset \Sigma_j$ over K; see (A) below. Furthermore, (C) below shows that Σ_j^g contains a point p_j far from the boundary $\partial \Sigma_j^g$ of the multi-valued graph. More precisely, $\partial \Sigma_j^g$ divides naturally into two parts, depending on whether or not it projects to ∂K ; (C) shows that the point p_j is far from the part of $\partial \Sigma_j^g$ that is not over ∂K . We will later use (C) to get multi-valued graphs with many sheets converging to a collapsed leaf.

¹²More precisely, the curve does not go below the union of Σ_j^g and the extrinsic ball $B_{r_j}(p_j)$.

The corollary will be used to produce spiralling multi-valued graphs, but it allows the possibility that the multi-valued graphs are in fact single-valued and, thus, just ordinary graphs; this occurs, for instance, if K is a disk. In the trivial single-valued case, $\partial \Sigma_j^g$ is a graph over ∂K and (C) below is trivially valid. On the other hand, when Σ_j^g is multi-valued, then (C) gives that the number of sheets is large.

COROLLARY II.3.18. Fix a point $p_0 \in K$. Given any (small) constant $\delta > 0$ and a (large) constant N, there exist $\varepsilon > 0$ and J so that for each j > J, we get the following:

- (A) There is a connected open subset $\Sigma_j^g \subset T(K,\varepsilon) \cap \Sigma_j$ so that, for each $x \in \Sigma_j^g$, the intrinsic ball $\mathcal{B}_{\varepsilon}(x)$ is a graph over (a subset of) Γ with gradient bounded by δ .
- (B) The normal exponential map from $K \times (-\varepsilon, \varepsilon)$ gives a diffeomorphism to $T(K, \varepsilon)$. Let $\Pi : T(K, \varepsilon) \to K$ denote the projection to K and Π_j the restriction of Π to Σ_j^g .
- (C) There is a point $p_j \in \Sigma_j^g$ with $\Pi_j(p_j) = p_0$ satisfying

(II.3.19)
$$\operatorname{dist}_{\Sigma_{i}^{g}}(p_{j}, \partial \Sigma_{j}^{g} \setminus \Pi_{i}^{-1}(\partial K)) > N.$$

Proof. Lemma II.3.15 gives $\varepsilon > 0$ (depending only on δ) so that for every point x in $T(K, \varepsilon) \cap \Sigma_j$ the intrinsic ball $\mathcal{B}_{\varepsilon}(x)$ is a graph over (a subset of) Γ with gradient bounded by δ .

Since K has compact closure in the (open) surface Γ , we can shrink $\varepsilon > 0$ so that the normal exponential map from $K \times (-\varepsilon, \varepsilon)$ gives a diffeomorphism to $T(K, \varepsilon)$.

Furthermore, since the Σ_j 's converge to \mathcal{L}' in a neighborhood of the point $p_0 \in \Gamma$, there is a sequence of points $p_j \in \Sigma_j$ converging to p_0 . (In fact, there are many such sequences; just pick one.) Let Σ_j^g be the component of $T(K, \varepsilon) \cap \Sigma_j$ containing p_j .

It remains to prove that (C) holds for J sufficiently large. We will do this by contradiction, so suppose that no such J exists for some fixed N. In particular, we get infinitely many j's where there exist curves $\gamma_j \subset \Sigma_j$ with the following properties:

- γ_j starts at p_j and ends at a point in $\partial T(K, \varepsilon)$ that is distance ε from K.
- The length of γ_i is at most N.
- γ_j is contained in $T(K, \varepsilon)$.

After passing to a subsequence, the γ_j 's must converge to a curve $\gamma \subset T(K, \varepsilon)$ that is contained in some leaf of \mathcal{L}' . (We are using here that γ_j stays away from \mathcal{S} .) Since the p_j 's converge to p_0 , the curve γ starts at $p_0 \in K$ and, hence,

we have

(II.3.20)
$$\gamma \subset \overline{K}$$

However, this is impossible since the second endpoints of γ_j are all distance ε from K and, thus, could not have converged to a point in K. This contradiction completes the proof.

II.3.3. Property (2) of Proposition II.3.1: Each collapsed leaf is stable. The main result of this subsection is that each oriented collapsed leaf of \mathcal{L}' is stable. The proof of stability has the following three main steps:

- Corollary II.3.18 gives multi-valued graphs $\Sigma_j^g \subset \Sigma_j$ converging to K with large multiplicity. The Σ_j^g 's can be thought of as single-valued graphs over a (subset of a) covering space K_j over K.
- Corollary II.3.21 describes the covering spaces K_j by analyzing the "holonomy" action of $\pi_1(K)$ on the fibers. (The holonomy is defined below.)
- Lemma II.3.22 then shows that a subsequence of the K_j 's satisfies (G1) and (G2) in Appendix A, so we can apply Corollary A.20 to see that K is stable.

Since this applies for any such K, and Γ can be exhausted by such K's by Lemma II.3.25 below, we conclude that Γ itself is stable.

The next corollary describes what the multi-valued graphs Σ_j^g look like as we follow them around a simple closed curve γ in K. Obviously, the preimage $\Pi_j^{-1}(\gamma)$ consists of a disjoint union of connected simple curves in the topological annulus $\Pi^{-1}(\gamma)$; see Figure 11.



Components of Σ_j^g in this cylinder.

Figure 11. Each component of $\Pi_i^{-1}(\gamma)$ is locally a graph over γ .

32

Some of the components of $\Pi_j^{-1}(\gamma)$ are more important than others. To distinguish the components, we will say that one of these components is "short" if it has two endpoints contained in the same boundary circle of $\Pi^{-1}(\gamma)$; otherwise, we will say the component is "long" (so a long component either has no boundary, or it has endpoints in distinct boundary circles of $\Pi^{-1}(\gamma)$). The corollary describes these long components:

COROLLARY II.3.21. Suppose that $\Pi_j : \Sigma_j^g \to K$ is as in Corollary II.3.18. If $\gamma \subset K$ is a simple closed curve, then either (1A) or (1B) holds:

- (1A) Each long component of $\Pi_j^{-1}(\gamma)$ is closed and is a graph over γ ; see Figure 12.
- (1B) The long components of $\Pi_j^{-1}(\gamma)$ are disjoint simple curves spiralling together from one boundary circle of $\Pi^{-1}(\gamma)$ to the other; see Figure 13.

If, in addition, K contains a simple closed curve σ that circles $p \in \Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ but is contractible in $\Gamma \cup \{p\}$, then

(2) $\Pi_j^{-1}(\sigma)$ has a single long component;¹³ this long component spirals from one boundary circle of the topological annulus $\Pi^{-1}(\sigma)$ to the other.

Schematic picture of Corollary II.3.21 (the two boundary circles of $\Pi^{-1}(\gamma)$ are dotted):



Figure 12. Case (1A): The long components of $\Pi_j^{-1}(\gamma)$ are graphs.



Figure 13. Case (1B): The long components of $\Pi_j^{-1}(\gamma)$ are multivalued graphs spiralling together from one boundary circle of $\Pi^{-1}(\gamma)$ to the other.

¹³Note that even though $\Pi_j^{-1}(\sigma)$ is all of where $\Pi^{-1}(\sigma)$ intersects the multi-valued graph, $\Pi_j^{-1}(\sigma)$ is not all of $\Pi^{-1}(\sigma) \cap \Sigma_j$. At the least, there must be another oppositely-oriented component of $\Pi^{-1}(\sigma) \cap \Sigma_j$ that spirals together; cf. the example of rescaled helicoids.

Proof. Since we are working in the *compact* embedded surface Σ_j and not in the limit, each component of $\Pi_j^{-1}(\gamma)$ is a simple curve with compact closure. In particular, these curves cannot spiral infinitely. Moreover, since $\Pi_j^{-1}(\gamma)$ is contained in the multi-valued graph, each component of $\Pi_j^{-1}(\gamma)$ is also locally a graph over γ .

If any long component is closed (and hence a graph over γ), then it separates the two boundary components of the topological annulus $\Pi^{-1}(\gamma)$ and, by embeddedness, every long component must be a closed graph over γ ; this is Case (1A). Suppose, on the other hand, that one (and, hence, every) long component connects the two boundary components of $\Pi^{-1}(\gamma)$. In this case, the embeddedness of $\Pi_j^{-1}(\gamma)$ forces all of these curves to spiral together; this is Case (1B).

Suppose now that a simple closed curve $\sigma \subset K$ circles $p \in \Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ but is contractible in $\Gamma \cup \{p\}$ and, in particular, does not circle any other points in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$. It follows from Lemma II.2.3 (and its proof) that the long components of $\Pi_j^{-1}(\sigma)$ do not close up and, hence, we are in Case (1B). It remains to see that there is just one long component. This follows immediately from Proposition II.1.3 in [CM04e], which shows that $\Pi^{-1}(\sigma) \cap \Sigma_j$ consists of exactly two oppositely oriented double spiral staircases.¹⁴ Since Σ_j^g is a multivalued graph over the connected set K, and hence can achieve only one of these orientations, it can contain only one of these.

We will say that K is sufficiently large when it contains a simple closed curve σ that circles exactly one point p in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ but is contractible in $\Gamma \cap \{p\}$, i.e., when (2) applies in Corollary II.3.21. We will assume in the rest of this section that K is sufficiently large.

LEMMA II.3.22. If the Σ_j 's are planar domains, K is sufficiently large, and $\gamma \subset K$ is a simple closed curve, then there can be only one long curve in (1B) of Corollary II.3.21 for j sufficiently large. More generally, when the Σ_j 's have bounded genus, then we get a bound for the number of distinct curves in (1B).

Proof. We will give the proof for genus zero, i.e., when the Σ_j 's are planar domains; the easy modifications needed for the general case are left to the reader.

Let $\sigma \subset K$ be a simple closed curve circling $p \in \Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$, and so σ is contained in a small neighborhood of p. (This exists since K was assumed to be sufficiently large.) Let $\gamma \subset K$ be a second simple closed curve. After possibly perturbing σ slightly, we can assume that it is disjoint from γ . Fix

¹⁴Technically, this description applies only when σ is in a neighborhood of p. This is sufficient for us since our σ is homotopic to a curve in a neighborhood of p and Σ_j^g is locally graphical over K.

points $x \in \sigma$ and $y \in \gamma$, and let $\eta \subset K$ be a simple curve from x to y. Again, after perturbing things, we can assume that η intersects σ and γ only at its endpoints x and y; see Figure 14.



Figure 14. The proof of Lemma II.3.22: The curves σ , η , and γ in K.

Suppose now that $\Sigma_j^g \subset \Sigma_j$ contains two distinct long curves, γ_1 and γ_2 , in $\Pi_j^{-1}(\gamma)$ that spiral together; see Figure 15. We will show that this leads to a contradiction by constructing two simple closed curves, μ_1 and μ_2 , in Σ_j that have linking number one in Σ_j . This is impossible for a planar domain. (It implies that the genus is at least one.)

We will first construct the curve $\mu_2 \subset \Sigma_j$ out of four parts; see Figure 16. The first part of μ_2 is a 1-valued graph over γ that is contained in γ_2 and has both of its endpoints over y. These two endpoints are distinct since they are at different heights over y. The next two parts of μ_2 are graphs over η that connect these two endpoints to two distinct points over $x \in \sigma$. Finally, we close the curve up by connecting the two points over x by a multi-valued graph over σ . Here, we have used that $\Pi_j^{-1}(\sigma)$ has *exactly* one long component to show that these endpoints can be connected and to see that the curve connecting them is at least 2-valued. Furthermore, we have also implicitly used that j is large to ensure that we can find the graphs over η and to ensure that the endpoints of these over x lie in a long component of $\Pi_j^{-1}(\sigma)$. (We will use that j is large in the same way later in the paper, usually without mentioning that we are doing so.)

The curve $\mu_1 \subset \Sigma_j$ is constructed similarly, with two notable differences; see Figure 17. First, the 1-valued graph over γ is chosen to be in γ_1 this time, as opposed to γ_2 before. Consequently, the one-valued graphs over γ in μ_1 and μ_2 are disjoint and, furthermore, the graphs over η are at four distinct heights. Second, instead of closing μ_1 up with a multi-valued graph over σ , do it over a slight outward perturbation of σ (see Figure 17). This makes the two "closing up" curves for μ_1 and μ_2 disjoint. However, since there is just one long component over σ (and also over its slight outward perturbation), we see that the "closing up" curve for μ_1 must cross one of the graphs over η in μ_2 . Moreover, this intersection is transverse and the curves are otherwise disjoint. This implies that μ_1 and μ_2 have linking number one in Σ_j , which gives the desired contradiction. The proof of Lemma II.3.22:



Figure 15. Two curves γ_1 and γ_2 in Σ_j spiral together over γ , but only one curve spirals over σ .



Figure 16. The (dashed) simple closed curve μ_2 in Σ_j has four parts: a 1-valued graph over γ in γ_2 ; a multi-valued graph over σ ; two graphs over η .

The contradiction for the proof of Lemma II.3.22:



Figure 17. Repeating the construction with γ_1 in place of γ_2 gives a second simple closed curve μ_1 . Perturbing the 1-valued graph over σ slightly outside of σ , μ_1 and μ_2 intersect in exactly one point and do so transversely. Hence, μ_1 and μ_2 have linking number one, which is impossible in the planar domain Σ_j .
Each Σ_j^g is a multi-valued graph over K but can be thought of as a singlevalued graph over a domain K_j in some covering space of K. However, this covering space may depend on j. Therefore, in order to apply the results of Appendix A, we need to pass to a subsequence so that

- The K_j 's all lie in the same covering space \hat{K} (independent of j).
- The K_i 's exhaust \hat{K} .
- The holonomy group of the covering space \hat{K} is **Z**. (The definition of the holonomy group is recalled below.)

In order to achieve these three points, we need a few elementary facts about covering spaces. First, recall that a covering space $\hat{\Pi} : \hat{K} \to K$ with base point $x \in K$ is uniquely determined by the *holonomy homomorphism* Hol from $\pi_1(K)$ to the automorphisms of the fiber $\hat{\Pi}^{-1}(x)$. To define this homomorphism, suppose that

$$(\text{II.3.23}) \qquad \qquad \gamma: [0,1] \to K$$

is a curve with $\gamma(0) = \gamma(1) = x$ and \hat{x} is a point in $\hat{\Pi}^{-1}(x)$. The lifting property for covering spaces gives a unique lift¹⁵

(II.3.24)
$$\gamma_{\hat{x}} : [0,1] \to \hat{K}$$

of γ with $\gamma_{\hat{x}}(0) = \hat{x}$. We define $\operatorname{Hol}(\gamma)(\hat{x})$ to be the endpoint $\gamma_{\hat{x}}(1)$. Finally, define the holonomy group to be the image $\operatorname{Hol}(\pi_1(K))$.

We will use the following elementary lemma in the proof:

LEMMA II.3.25. Given a connected surface Γ without boundary, there exists a sequence of connected open sets K_j with compact closure that exhaust Γ . That is, we have $\Gamma = \bigcup_{j=1}^{\infty} K_j$ and $K_j \subset K_{j+1}$ for every j.

Proof. This follows immediately from the existence of a complete metric proven in [NO61]. \Box

We are now ready to prove that each oriented collapsed leaf is stable.

Proof of (2) in Proposition II.3.1. We will show that any connected open subset $K \subset \Gamma$ that has compact closure in Γ and is sufficiently large must be stable. Since Γ can be exhausted by such K's by Lemma II.3.25, we will conclude that Γ itself is stable.

Fix a point $x \in K$. By repeatedly applying Corollary II.3.18 with $\delta = 1/j$ and passing to a subsequence, we get a sequence of connected multi-valued graphs Σ_j^g over K, covering spaces $\Pi_j : \hat{K}_j \to K$, domains $K_j \subset \hat{K}_j$, and

¹⁵Recall that $\gamma_{\hat{x}}$ is said to be a *lift* of γ if $\gamma = \hat{\Pi} \circ \gamma_{\hat{x}}$. The lifting property for covering spaces says that we get a unique lift of γ for each choice of point \hat{x} with $\hat{\Pi}(\hat{x}) = x$.

functions $u_j: K_j \to \mathbf{R}$ with

(II.3.26)
$$|u_j| + |\nabla u_j| \le 1/j,$$

so that there is a bijection from K_j to Σ_j^g given by

(II.3.27)
$$x \to \Pi_j(x) + u_j(x) \mathbf{n}_{\Gamma}(\Pi_j(x)).$$

Furthermore, (II.3.19) gives a point $x_j \in \Sigma_j^g$ with $\Pi_j(x_j) = x$ satisfying

(II.3.28)
$$\operatorname{dist}_{\Sigma_{j}^{g}}(x_{j}, \partial \Sigma_{j}^{g} \setminus \Pi_{j}^{-1}(\partial K)) > j.$$

We must do two things in order to apply Corollary A.20 in Appendix A. Namely, we must pass to a subsequence so that the K_j 's all sit in the same covering space \hat{K} and we must show that the holonomy group of \hat{K} is **Z**. Once we have done these, (II.3.28) will imply that the K_j 's exhaust \hat{K} .

We will deal with the second one first; i.e., we will show that the holonomy group is always **Z**. This follows immediately from Corollary II.3.21. Namely, (2) in Corollary II.3.21 implies that the fiber over x in each K_j can be identified with **Z** and the holonomy from circling the point $p \in S_{\text{ulsc}}$ is just $n \to (n + 1)$ or $n \to (n - 1)$, depending on whether the multi-valued graph spirals up or down. Suppose now that γ is a simple closed curve through x representing a homotopy class $[\gamma]$ in $\pi_1(K)$. Furthermore, (1A) and (1B) in Corollary II.3.21 imply that either

- If (1A) holds, then $\operatorname{Hol}([\gamma])$ is the identity map; i.e., $n \to n$.
- If (1B) holds, then Hol([γ]) maps to n → (n ± k), where k is the number of disjoint curves spiralling together in (1B).

In particular, the image of the holonomy is always in \mathbf{Z} in either case.

Finally, we will use Lemma II.3.22 to prove that only a finite set of distinct covering spaces arise as one of the \hat{K}_j 's and, consequently, one of the \hat{K}_j 's occurs infinitely many times. We have already established that each holonomy group is **Z**, but the covering space is determined by the holonomy homomorphism (and not just the group). Each holonomy homomorphism

(II.3.29)
$$\operatorname{Hol}_j : \pi_1(K) \to \mathbf{Z}$$

is determined by the image of a fixed (finite) set of generators $\gamma_1, \ldots, \gamma_m$ of $\pi_1(K)$, so we need only to show a uniform bound for $\operatorname{Hol}_j(\gamma_n)$ for every j and n. However, (1B) in Corollary II.3.21 implies that $\operatorname{Hol}_j(\gamma_n)$ is just the number of disjoint (long) curves spiralling together in $\Pi_j^{-1}(\gamma_n)$ and Lemma II.3.22 bounds this uniformly, completing the proof.

Remark II.3.30. We have assumed throughout this subsection that the leaf Γ is oriented. When this is not the case, the same argument applies to show that the oriented double cover is stable.

II.3.4. Property (3) of Proposition II.3.1: Opposite orientations at distinct points of $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$.

Proof of (3) in Proposition II.3.1. We must show that if p and q are distinct points in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$, then the multi-valued graphs in Σ_j near p spiral in the opposite direction as the ones near q. Since there are only two possible directions, this implies that $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ contains at most two points. (If there were three such points, then two would have to be oriented the same way, which we will show is impossible.)

We will argue by contradiction, so suppose that the multi-valued graphs near p and q have the same orientation. In this case, we can choose a closed "figure eight" curve γ_j in Σ_j with the following properties (see Figure 18):

- γ_j is a graph over a fixed (immersed) figure eight curve γ in Γ that circles p and q in opposite directions. Let $r \in \gamma$ be the double point where γ is not embedded.
- The two points in γ_j above the double point r are in distinct sheets of Σ_j ; hence γ_j is embedded.



The second condition has a very useful consequence. Namely, the unit normal to Σ_j is always either upward or downward pointing along γ_j since Σ_j is graphical along γ_j ; therefore, elementary topology implies that

• The two points in γ_j above r are separated by an oppositely-oriented sheet of Σ_j .

We will now use these properties of the γ_j 's to find stable minimal surfaces Γ_j disjoint from the Σ_j 's that contain a graph near either p or q, contradicting that these points are in S_{ulsc} . Since Σ_j has genus zero, the curve γ_j separates in Σ_j ; let Σ_j^+ be one of the two components of $\Sigma_j \setminus \gamma_j$. Since $B_{R_j} \setminus \Sigma_j$ is mean convex in the sense of Meeks-Yau, the existence theory of [MY82b] gives a stable orientable embedded minimal planar domain

(II.3.31)
$$\Gamma_i^+ \subset B_{R_i} \setminus \Sigma_j \text{ with } \partial \Gamma_i^+ = \partial \Sigma_i^+.$$

Let Γ_j be the component of Γ_j^+ with $\gamma_j \subset \partial \Gamma_j$. Using estimates for orientable stable surfaces ([Sch83a]; cf. [CM02a]) and the fact that γ_j is a figure eight, it is now not hard to see that Γ_j must contain a graph near either p or q (see Figure 19). This can be seen as follows:

(1) After leaving the upper portion of γ_j over r, the stable surface Γ_j is separated from the lower portion of γ_j by an oppositely-oriented sheet of Σ_j and, hence, Γ_j has an *a priori* curvature bound there by [Sch83a]; cf. [CM02a].

To see that Γ_j does indeed leave the upper portion of γ_j over r, intersect Γ_j with a large transverse¹⁶ ball $B_R(z)$ to get a collection of closed curves and one segment σ_j where σ_j connects the upper and lower portions of γ_j (see Figure 20). Since these are separated near r by an oppositely-oriented sheet of Σ_j , the segment σ_j moves away from γ_j as desired.

- (2) Away from the singular points p and q, the surface Γ_j is locally pinched between sheets of Σ_j . Combining this pinching with the curvature bound from (1) implies that $\Gamma_j \to \Gamma$ away from p, q, and γ_j . (Here, "away" is with respect to distance along paths in $B_{R_j} \setminus \Sigma_j$.)
- (3) Combining the *a priori* bound of 1. with the flatness given by (2), unique continuation forces $\Gamma_j \to \Gamma_{\text{Clos}}$ even as it approaches *p* or *q*. (This unique continuation argument is spelled out in Lemma II.1.38 in [CM04d].) However, the one-sided curvature estimate, i.e., Theorem I.1.3, would then apply to the Σ_j 's near *p* or *q*, contradicting that $|A| \to \infty$ near *p* and *q*.

This contradiction shows that the multi-valued graphs near p and q are oppositely-oriented, completing the proof of (3).

This completes the proof of Proposition II.3.1.



The two parts of γ_j do not connect near r; hence, $\partial B_R(z) \cap \Gamma_j$ moves away from γ_j .

Figure 20. The stable surface moves away from its boundary near r.

¹⁶The application of transversality uses the regularity of $\tilde{\Gamma}_j$ up to the interior of $\tilde{\gamma}$. Local boundary regularity was established for two-dimensional minimal surfaces in [Hil69].

Remark II.3.32. The genus bound on the Σ_j 's can be used to directly see that $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ cannot contain three points. To see this, suppose that p, q, and r are three distinct points in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ and γ_{pq} is a geodesic in Γ from pto q. For j large, Theorem I.1.6 allows us to find simple closed curves $\gamma_{pq}^j \subset \Sigma_j$ with the following properties:

- γ_{pq}^{j} is contained in the ε -tubular neighborhood of γ_{pq} .
- $\gamma_{pq}^{j} \setminus (B_{\varepsilon}(p) \cup B_{\varepsilon}(q))$ consists of two graphs over γ_{pq} that are in distinct sheets of Σ_{j} .

Since Σ_j has genus zero, the curve γ_{pq}^j must separate Σ_j into two distinct components. However, it is easy to see that this is impossible by using the local connecting property near the third point r. Namely, we can take two points near p on opposite sides of γ_{pq}^j and connect each of them to $B_{\varepsilon}(r)$ by curves in Σ_j that do not intersect γ_{pq}^j . These two curves can then be connected to each other in $B_{\varepsilon}(r) \cap \Sigma_j$, giving the desired contradiction.

Part III. When the surfaces are ULSC: The proof of Theorem 0.9

In this part, we will prove Theorem 0.9, i.e., the main structure theorem for ULSC sequences where $S_{neck} = \emptyset$. The key will be to analyze the ULSC singular set S_{ulsc} and, in particular, the collapsed leaves of \mathcal{L}' . Although the emphasis will be on the ULSC case, many of the arguments will actually apply to a neighborhood of a collapsed leaf whose closure does not intersect S_{neck} . This will be used later when we analyze the general case.

In the previous section, we showed that a collapsed leaf Γ of \mathcal{L}' is a stable, incomplete minimal surface with isolated removable singularities at points in \mathcal{S}_{ulsc} . In general, Γ may have worse singularities at points of $\Gamma_{Clos} \cap \mathcal{S}_{neck}$, but we will assume that $\Gamma_{Clos} \cap \mathcal{S}_{neck} = \emptyset$ in this part.

In addition to what we have shown in the previous section, we need to establish two facts to complete the proof of Theorem 0.9. First, we must show that every collapsed leaf is a plane. (It then follows easily from embeddedness that all of these planes are parallel.) Since we have shown in Proposition II.3.1 that the collapsed leaves are stable with isolated removable singularities at each point of S_{ulsc} , this follows easily from the Bernstein theorem for complete stable surfaces in \mathbb{R}^3 . The second additional fact that must be established is the "properness" of the limit in the sense of [CM02b]. Roughly speaking, the local cone property already implies that the closed set S is *contained in* two Lipschitz curves each of which is transverse to the limit planes. The properness consists of showing that S actually fills out these curves completely; i.e., there cannot be a first or last point in S. See (\star) in Section III.1 for the precise statement. As in [CM02b], we will prove properness by showing that the vertical flux of a potential nonproper limit would have to be positive, which is impossible by Stokes' theorem. All of this will show the following (after a rotation of \mathbf{R}^3):

• The ULSC sequence of surfaces converges to the foliation by parallel planes

$$\mathcal{F} = \{x_3 = t\}_t$$

away from the singular set \mathcal{S} .

- (C_{ulsc}) from Theorem 0.9 holds.
- (D'_{ulsc}) : \mathcal{S} consists of two disjoint Lipschitz graphs $\mathcal{S}_1 : \mathbf{R} \to \mathbf{R}^3$ and $\mathcal{S}_2 : \mathbf{R} \to \mathbf{R}^3$ over the x_3 -axis.

From this, it follows immediately from the main theorem of [Mee04] that S_1 and S_2 are in fact straight lines *orthogonal* to the leaves of the foliation, giving (D_{ulsc}) from Theorem 0.9 and completing the proof of Theorem 0.9.

Recall that collapsed leaves are the leaves of \mathcal{L}' that "go through" a point of \mathcal{S}_{ulsc} , i.e., that contain the local leaf Γ_x given by Lemma II.2.3 for some $x \in$ \mathcal{S}_{ulsc} ; see Definition II.2.9. The next proposition establishes the key properties of a collapsed leaf in the ULSC case.

PROPOSITION III.0.2. Suppose that Γ is a collapsed leaf of \mathcal{L}' . If $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$, then

- (1) Γ_{Clos} is a plane.
- If, in addition, $S_{neck} = \emptyset$ (i.e., the sequence is ULSC), then
- (2) Γ_{Clos} intersects S_{ulsc} in exactly two points and the multi-valued graphs in the Σ_j 's spiral in opposite directions around the two corresponding axes (see Figure 8).

This proposition will be proven over the rest of this section.

III.0.5. Property (1) in Proposition III.0.2: Collapsed leaves are planar. To prove that Γ is flat, we first use property (3) in Proposition II.3.1 to see that Γ_{Clos} is the union of Γ together with at most two points in $\mathcal{S}_{\text{ulsc}}$ since we are assuming that $\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}} = \emptyset$. In particular, since each point in $\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{ulsc}}$ is a removable singularity by (1) in Proposition II.3.1, we conclude that Γ_{Clos} is a smooth complete surface without boundary.

Assuming first that Γ is oriented, (2) in Proposition II.3.1 implies that Γ is stable. We can then use a standard logarithmic cutoff argument at each point in $\Gamma_{\text{Clos}} \setminus \Gamma$ to conclude that Γ_{Clos} is itself stable. The Bernstein theorem for stable complete minimal surfaces, [FCS80], [dCP79], then implies that Γ_{Clos} is a plane, as desired. When Γ is not oriented, the preceding discussion applies to show that its oriented double cover is flat — and hence so is Γ . The obvious details are left to the reader.

III.0.6. Property (2) in Proposition III.0.2: Ruling out just one point of S_{ulsc} in a leaf. In contrast to property (1), we will need to use that the sequence

42

is ULSC in order to prove (2). We will later see that this assumption can be removed. However, the argument we will give to prove (2) in general will use the ULSC case that we are proving now. (This is why we are not proving the general case directly.)

We have shown in property (3) of Proposition II.3.1 that the closure of a collapsed leaf contains at most two ULSC singular points and that the Σ_j 's spiral in opposite directions around two such points. Hence, to prove (2) in Proposition III.0.2, we must show that Γ_{Clos} cannot intersect S in just one point.

Before proving (2), we need to recall a useful property of stable minimal surfaces. Namely, the following lemma shows that a stable surface that starts out on one side of a plane where the interior boundary is in a small ball is graphical away from its boundary (see Figure 21).

LEMMA III.0.3. There exists a small constant $0 < \delta < 1$ so that if $r_0 < \delta R_0$ and $\Gamma \subset B_{R_0}$ is a connected embedded stable minimal planar domain with nonempty inner boundary $\gamma = \partial \Gamma \setminus \partial B_{R_0}$ contained in the small ball $B_{\delta r_0}$, outer boundary $\partial \Gamma \setminus B_{\delta r_0}$ nonempty, and

(III.0.4)
$$B_{r_0} \cap \Gamma \cap \{x_3 = 0\} = \emptyset,$$

then Γ contains a graph over the annulus $D_{\delta R_0} \setminus D_{r_0} \subset \{x_3 = 0\}$. Moreover, this graph can be connected to the inner boundary γ by a curve in $B_{2r_0} \cap \Gamma$.



Stable Γ (dotted) becomes graphical.

Figure 21. Lemma III.0.3: The stable surface Γ starts off close to — but above — a disk and is forced to become graphical.

Proof. The proof has two steps. Namely, we first show that Γ contains an initial graph over a small annulus on the scale of δr_0 . The next step uses the initial graph to apply the "stable graph proposition" — Proposition C.2 in Appendix C — to get the desired graph over the large annulus $D_{\delta R_0} \setminus D_{r_0}$.

Producing the initial graph. Given a point $y \in \Gamma$ with $2\delta r_0 \leq |y| \leq R_0/4$, the estimates for stable surfaces of [Sch83a] gives C_s so that $\sup_{B_{|y|/3}(y)} |A| \leq C_s |y|^{-1}$. Since x_3 is a positive harmonic function on Γ , the Cheng-Yau gradient estimate then gives C' so that

(III.0.5)
$$\sup_{\mathcal{B}_{|y|/4}(y)} |\nabla x_3| \le C' \frac{y_3}{|y|},$$

where y_3 is the x_3 component of y. Moreover, the Harnack inequality gives that

(III.0.6)
$$\sup_{\mathcal{B}_{|y|/4}(y)} x_3 \le C_H y_3.$$

Thus, when $y_3 < |y|/C'$, then (III.0.5) implies that $\mathcal{B}_{|y|/4}(y)$ is a graph. Furthermore, (III.0.6) gives a constant $\beta > 0$ so that if $y_3 \leq \beta |y|$, then we can iterate the Harnack enough times to get a graph over $D_{2|y|} \setminus D_{|y|}$ all the while keeping x_3 less than |y|/C'; this is a graph and not a multi-valued graph since Γ is embedded and does not spiral infinitely.

In order to implement this argument, we have to find some "low point" y where $y_3 \leq \beta |y|$. This is given by a standard catenoid barrier argument (Lemma 3.3 in [CM02c]): Given any $\beta > 0$, there is a constant C_1 so that $\partial B_{C_1\delta r_0} \cap \Gamma$ contains such a low point and, moreover, it is in the same connected component of $B_{2C_1\delta r_0} \cap \Gamma$ that contains the inner boundary. Let Γ^g be the graph in Γ over the annulus $D_{2C_1\delta r_0} \setminus D_{C_1\delta r_0}$.

Applying Proposition C.2 to get the graph from r_0 to δR_0 . Let $\hat{\gamma}$ be the graph in Γ^g defined over the inner boundary $\partial D_{C_1 \delta r_0}$. The simple closed curve $\hat{\gamma}$ separates the planar domain Γ into two components; let Γ_0 be the "outer" one, so that γ is not in $\partial \Gamma_0$. Since the tubular neighborhood $\operatorname{An}_{C_1 \delta r_0}(\hat{\gamma})$ of radius $C_1 \delta r_0$ of $\hat{\gamma}$ in Γ_0 is contained in the graph Γ^g , we get uniform bounds for the area and total curvature of $\operatorname{An}_{C_1 \delta r_0}(\hat{\gamma})$. Now apply Proposition C.2 to $\hat{\gamma}$ to get a graph in Γ defined over the annulus $D_{R_0/\omega} \setminus D_{\omega C_1 \delta r_0}$. Finally, we fix $\delta > 0$ less than ω^{-1} and $(C_1 \omega)^{-1}$ so that $D_{R_0/\omega} \setminus D_{\omega C_1 \delta r_0}$ contains $D_{\delta R_0} \setminus D_{r_0}$.

We note next that the local cone property has two important consequences for the singular set S_{ulsc} and, in particular, for how the singular set changes as we move from one collapsed leaf to the next:

- (S1) S_{ulsc} cannot run off to infinity.
- (S2) Distinct points of S_{ulsc} in a collapsed leaf cannot combine in another collapsed leaf.

Proof of property (2) in Proposition III.0.2. To prove that $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ contains exactly two points, suppose for a moment that there was only one singular point. (There is always at least one by definition.) In particular, after a translation and rotation of \mathbf{R}^3 , we may suppose that

(III.0.7)
$$\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{ulsc}} = \{0\},$$

and the collapsed leaf through 0 is the punctured horizontal plane $\{x_3 = 0\}$ $\setminus \{0\}$. We will show that (III.0.7) implies that *every* leaf of \mathcal{L}' is a plane with one point removed, these planes foliate \mathbf{R}^3 , and as a consequence the intersection of any fixed ball with the surfaces Σ_j is simply connected for j sufficiently large. However, this is impossible since we have assumed an upper bound for the injectivity radii of the Σ_j 's in Theorem 0.9, so we conclude that (III.0.7) cannot hold.

Properness. We will show next that every open neighborhood of $\{x_3 = 0\}$ contains points of S_{ulsc} both above and below $\{x_3 = 0\}$; the proof will use only that $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$. We called this properness in [CM04e], and the argument is essentially the same, with one caveat: [CM04e] argues for embedded minimal disks, whereas presently we only know that the Σ_j 's are ULSC near 0. The disk hypothesis was used for two things in [CM04e]:

- (D1) The Σ_j 's are multi-valued graphs in the cone $\{|x_3| < \mu |x|\}$ for some $\mu > 0$.
- (D2) The portions of these two multi-valued graphs in a fixed ball combine to be part of a single embedded minimal disk in this small ball. (This disk property was used in [CM04e] to apply Stokes' theorem.)

The second fact (D2) holds in this case since $0 \in S_{ulsc}$. The first fact (D1) will follow immediately from the one-sided curvature estimate once we establish the following *scale invariant* ULSC property:

(D) There exists $\tau > 0$ so that, for $z \in \{x_3 = 0\}$ and j large, each component of $B_{\tau|z|}(z) \cap \Sigma_j$ that connects to the multi-valued graph in Σ_j is a disk.

We will next prove (D) by contradiction, using a variation of the "between the sheets" estimate of [CM04b]. To do this, assume that $\tau > 0$ is small and $z \in \{x_3 = 0\}$ is the first time that (D) fails (i.e., |z| is minimal); obviously, we must have $|z| > r_0$ since the $0 \in S_{\text{ulsc}}$. Fix a sequence of simple closed noncontractible curves $\gamma_j \subset B_{\tau|z|}(z) \cap \Sigma_j$. We will see that this leads to a contradiction:

- (1) See Figure 22. The existence results of Meeks-Yau, [MY82b], gives stable embedded connected minimal surfaces $\Gamma_j \subset B_{R_j} \setminus \Sigma_j$ with $\partial \Gamma_j \setminus \partial B_{R_j} = \gamma_j$ and $\partial \Gamma_j \cap \partial B_{R_j} \neq \emptyset$.¹⁷ Since these stable surfaces start out on one side of, but close to, the multi-valued graphs in Σ_j converging to $\{x_3 = 0\} \setminus \{0\}$, Lemma III.0.3 implies that the Γ_j 's quickly become graphical.
- (2) See Figure 23. The portion Σ_j^+ of Σ_j between $\partial B_{|z|/2} \cap \{x_3 = 0\}$ and the graph in Γ_j must be simply connected in extrinsic balls of radius $\tau |z|/2$ since it is trapped "between the sheets" of the multi-valued graph in Σ_j and

¹⁷This is a standard application of [MY82b]; see, e.g., Lemma 3.1 in [CM02c] for details.



Figure 22. The stable surface Γ_j is dotted.



Second stable surface $\tilde{\Gamma}_j$ would be forced to cut either Γ_j or Σ_j .

Figure 23. Σ_j is simply connected between the multivalued graph in Σ_j and the graph in Γ_j .

the graph in Γ_j (and these two can be connected). Namely, if Σ_j^+ contained a noncontractible curve in an extrinsic ball of radius $\tau |z|/2$ centered there, then we could apply [MY82b] to get a second stable surface $\tilde{\Gamma_j}$ disjoint from both Σ_j and Γ_j . The surface $\tilde{\Gamma_j}$ would be forced to become graphical (again by Lemma III.0.3) but would then have to cross the curve in $\Sigma_j \cup \Gamma_j$ that connects the two graphical regions. (Compare the proof of Theorem I.0.8 in [CM04b].)

(3) Since each ball of radius $\tau |z|/2$ centered on Σ_j^+ is simply connected by (2), the set Σ_j^+ is locally graphical by the one-sided curvature estimate and, since it contains a multi-valued graph but cannot pass through Γ_j , it spirals infinitely. This is impossible since each Σ_j is compact, giving the contradiction needed to establish (D).

Now that we have established (D), we can argue precisely as in [CM04e] using [CM02b] (see Lemma I.1.10 there) to prove that every open neighborhood of this plane contains points of S_{ulsc} both above and below this plane.

It follows easily from this properness — as in Lemma I.1.2 of [CM04e] — that an entire slab $\{-\varepsilon < x_3 < \varepsilon\}$ must be foliated by (the closures of) planar leaves of \mathcal{L}' . In order to extend this foliated structure to all of \mathbf{R}^3 , we will need to use the ULSC assumption next.

Using the ULSC hypothesis to repeat the argument. Since the set S_{ulsc} is automatically closed and transverse to these planes (by the one-sided curvature estimate), we see that a neighborhood of this plane is foliated by parallel planes and, in this neighborhood, S_{ulsc} is a single Lipschitz curve. However, since the constant $\tau > 0$ was uniform and did not depend on the particular singular point, we can now repeat the above argument to extend the set of foliated planes to the whole of \mathbf{R}^3 . Once we have the foliation by parallel planes, the ULSC condition and one-sided curvature estimate imply that S_{ulsc} is a discrete collection of transverse Lipschitz curves. The transversality implies

that each curve hits every leaf and hence there is only one curve. In sum, the sequence is converging to a foliation of \mathbf{R}^3 by parallel planes away from a single Lipschitz curve S_{ulsc} transverse to the planes. (This was exactly the result of [CM04b]–[CM04e] for sequences of disks.) This has two consequences:

- Near S_{ulsc} , the sequence looks like a double spiral staircase.
- Away from S_{ulsc} , the sequence is locally converging (with bounded curvature) to a foliation by parallel planes.

The second fact allows us to extend the double spiral staircase structure away from the singular curve S_{ulsc} , so that we get a sequence $R'_j \to \infty$ where the component of $B_{R'_j} \cap \Sigma_j$ intersecting $B_{R'_j/C}$ is a double spiral staircase. In particular, this component is also a disk. Since we have assumed that no such sequence of expanding disks in Σ_j exists, we rule out (III.0.7) as promised. \Box

III.1. Properness and the limit foliation

We have now shown that each collapsed leaf is a plane that is transverse to S_{ulsc} (with a definite lower bound on the angle of intersection). As in [CM04e], we must show that nearby leaves are also planes; we call this *properness* of the limit foliation. Since each singular point in S_{ulsc} has a plane through it, this properness will follow from showing that there cannot be a first or last such singular point. In fact, it is not hard to see that these properties are equivalent. Namely, a planar leaf nearby a collapsed leaf must also contain singular points since otherwise the one-sided curvature estimate would give a curvature bound at the singular point in the collapsed leaf. In [CM04e], a similar properness for disks (where each plane had only one puncture as opposed to the current situation of two) was proven using [CM02b].

Before giving the precise statement of properness, observe that we can rotate \mathbf{R}^3 so that the closure of each collapsed leaf of \mathcal{L}' is a horizontal plane, i.e., is given by $\{x_3 = t\}$ for some $t \in \mathbf{R}$. This is because the closure of each collapsed leaf is some plane by Proposition III.0.2 and these planes must all be parallel since the surfaces Σ_j are embedded. With this normalization, the precise statement of properness is

(*) If $t \in x_3(\mathcal{S}_{ulsc})$ and $\varepsilon > 0$, then $\mathcal{S} \cap \{t < x_3 < t + \varepsilon\} \neq \emptyset$ and $\mathcal{S} \cap \{t - \varepsilon < x_3 < t\} \neq \emptyset$.

We should point out that when the sequence is ULSC, (*) automatically implies that $\{x_3 = t\} \cap S_{\text{ulsc}}$ contains at least *two* points (cf. (S1) and (S2)). Namely, once $\{x_3 = t\} \cap S_{\text{ulsc}}$ contains one point p', then Proposition III.0.2 implies that there is a second point $q' \in \{x_3 = t\} \cap S_{\text{ulsc}}$ so that

$$(\text{III.1.1})\qquad\qquad \{x_3=t\}\setminus\{p',\,q'\}$$

is a collapsed leaf of \mathcal{L}' .



Low ends of Σ will be graphs.

Figure 24. The limit Σ when properness fails.

As in [CM04e] and [CM02b], the key to proving (\star) is a careful analysis of the vertical flux of the multi-valued graphs. Recall that if Σ is a minimal surface and $\sigma \subset \Sigma$ is a simple closed curve, then the vertical flux across σ is

(III.1.2)
$$\int_{\sigma} \frac{\partial x_3}{\partial n},$$

where $\frac{\partial x_3}{\partial n}$ is the derivative of x_3 in the unitary direction normal to σ but tangent to Σ . By Stokes' theorem, the integral (III.1.2) depends only on the homology class of σ since the coordinate function x_3 is harmonic on a minimal surface.

We will prove (\star) by contradiction as we now outline: If (\star) does not hold for t = 0, then we can assume, after possibly reflecting across $\{x_3 = 0\}$, that

$$(\text{III.1.3}) \qquad \qquad \mathcal{S} \cap \{0 < x_3 < \varepsilon\} = \emptyset.$$

We will use this to show that there is a unique leaf Σ of \mathcal{L}' that spirals into the plane $\{x_3 = 0\}$ from above and this leaf is a multiplicity one limit of the Σ_j 's. Moreover, near the singular points in $\{x_3 = 0\}$, Σ will be a double spiral staircase that spirals infinitely into the plane, with the two spirals oppositelyoriented (see (3) in Proposition II.3.1); see Figure 24. The ends of Σ coming from circling both double spiral staircases will be graphs lying above the plane $\{x_3 = 0\}$; we will see that this implies that each such end has nonnegative vertical flux. This structure of the ends also allows us to cutoff Σ below a carefully chosen horizontal plane $\{x_3 = \varepsilon\}$; it will be almost automatic that the boundary curve produced has positive vertical flux. Using the fact that the plane $\{x_3 = 0\}$ contains two points of \mathcal{S}_{ulsc} , we will find a sequence of separating curves in Σ whose vertical flux goes to zero. This gives a sequence of compact domains in Σ bounded at the top by a curve in the plane $\{x_3 = \varepsilon\}$ with positive flux, bounded at the bottom by curves with flux going to 0, and with boundary curves on the sides with nonnegative flux. Combining all of

this will give the desired contradiction since, by Stokes' theorem, the total flux of any compact domain must sum to zero.

There are, in general, two ways to show that the vertical flux in (III.1.2) is small; one can either show that the curve σ is short (since $|\nabla x_3| \leq 1$) or show that $|\nabla x_3|$ is small and the length of σ is bounded. Since the length of any closed noncontractible curve near $\{x_3 = 0\}$ is bounded away from zero, we must take the second approach here to bound the flux of the bottom boundary curves. (We will use the first approach in the next part near points in S_{neck} .) In our application, the harmonic function x_3 will be positive on the surface Σ and the estimate on $|\nabla x_3|$ will follow from the gradient estimate.

III.1.1. Establishing properness: The proof of (\star) . The next lemma shows that (\star) holds as long as we have properties (1) and (2) in Proposition III.0.2. In particular, since these properties always hold when the sequence is ULSC, we get (\star) in the ULSC case.

LEMMA III.1.4. If (1) and (2) in Proposition III.0.2 hold, then (\star) holds. That is, if the horizontal plane { $x_3 = t$ } is the closure of a collapsed leaf satisfying (2) in Proposition III.0.2, and $\varepsilon > 0$, then

(III.1.5)
$$\mathcal{S} \cap \{t < x_3 < t + \varepsilon\} \neq \emptyset \text{ and } \mathcal{S} \cap \{t - \varepsilon < x_3 < t\} \neq \emptyset.$$

Proof. For simplicity, we will assume that the Σ_j 's have genus zero. The general case follows with easy modifications.

We will argue by contradiction, so suppose that $S_{\text{ulsc}} \cap \{x_3 = 0\} = \{p, q\}, \varepsilon > 0$, and

(III.1.6)
$$\mathcal{S} \cap \{0 < x_3 < \varepsilon\} = \emptyset.$$

Fix a radius R > 0 so that the disk $D_R \subset \{x_3 = 0\}$ contains both p and q.

We will first record four consequences of (III.1.6) that will be proven below and then use these properties to rule out the possibility of such a nonproper limit (see Figure 25):

- (P1) There is exactly one leaf Σ of \mathcal{L}' in $\{x_3 > 0\}$ whose closure intersects the plane $\{x_3 = 0\}$. This leaf Σ is a multiplicity one limit of the Σ_j 's. Furthermore, after possibly reducing $\varepsilon > 0$, the leaf Σ is proper in compact subsets of $\{0 < x_3 < \varepsilon\}$.
- (P2) Each "low" end of Σ is an asymptotic graph with nonnegative vertical flux. Here "low" will be made precise below but roughly means starting off close to $\{x_3 = 0\}$ over the disk D_R containing p and q.
- (P3) Intersecting Σ with a carefully chosen horizontal plane where x_3 is constant will give a vertically separating curve $\gamma_+ \subset \Sigma$ with positive vertical flux. Here vertically separating means that if a curve in $\overline{\Sigma}$ is over D_R and

intersects both of

(III.1.7)
$$\{x_3 = 0\}$$
 and $\{x_3 = \varepsilon\}$,

then the curve also intersects γ_+ .

(P4) There is a sequence of vertically separating curves $\gamma_i \subset B_1 \cap \Sigma$ with $x_3(\gamma_i) \to 0$ and vertical flux going to zero.

(P3): Intersecting with a plane between two ends gives a curve γ_+ that is vertically separating and has positive flux.



will have ends that are graphs with nonnegative flux.

Figure 25. Properties (P2)–(P4) of Σ when properness fails.

The proof of (P1). Since p and q are locally the last points in S_{ulsc} , we are in case (L2) in the proof of Lemma II.3.3 near p and q. (Cf. the "not proper" example (N-P).) Consequently, Lemma II.3.4 implies that near pand q, but above $\{x_3 = 0\}$, the Σ_j 's converge with multiplicity one to double spiral staircases that spiral infinitely into $\{x_3 = 0\}$.¹⁸ More precisely, let $\Omega_p, \Omega_q \subset \{x_3 > 0\}$ be the convex cones with p and q as their vertices where the one-sided curvature estimate does not apply. Then near p, above $\{x_3 = 0\}$, but away from Ω_p , we know that \mathcal{L}' has exactly two leaves that spiral together infinitely into $\{x_3 = 0\}$. Furthermore, the same is true for q (though the spirals have the opposite orientation there). Finally, we also know that are multiplicity one limits and that they are isolated (i.e., not limit leaves).

The next step is to prove properness. Namely, we show that there exists some $\varepsilon_0 > 0$ so that Σ is proper in compact subsets of

$$(\text{III.1.8}) \qquad \qquad \{0 < x_3 < \varepsilon_0\}.$$

This properness will follow by combining the two following facts:

(Fact 1) There exists some $\varepsilon_1 > 0$ so that \mathcal{L}' does not contain any horizontal planes in the slab $\{0 < x_3 < \varepsilon_1\}$.

¹⁸This is proven within the proof of Lemma II.3.4.

(Fact 2) There exists $\varepsilon_2 > 0$ so that if $y \in \{0 < x_3 < \varepsilon_2\}$ is an "accumulation point" of Σ , then the leaf of \mathcal{L}' containing y is a horizontal plane. A point y is said to be "accumulation point"¹⁹ of Σ if there exists a sequence of points $y_j \in \Sigma$ so that

(III.1.9)
$$\lim_{j \to \infty} \operatorname{dist}_{\mathbf{R}^3}(y, y_j) = 0,$$

(III.1.10)
$$\lim_{j \to \infty} \operatorname{dist}_{\Sigma}(y_1, y_j) = \infty.$$

Obviously, the two facts together easily imply the properness of \mathcal{L}' in the slab between 0 and the minimum of ε_1 and ε_2 .

To prove (Fact 1), observe that the spiraling leaves go from some positive height all the way down to $\{x_3 = 0\}$ and the leaves cannot cross.

(Fact 2) follows easily from the proof of Lemma 1.3 in [MR05]. (In fact, one can even take $\varepsilon_2 = \varepsilon$ but we will not need this.) We will give the proof next for completeness. Suppose therefore that a leaf Σ accumulates at a point $y \in \{0 < x_3 < \varepsilon_2\}$ in a leaf $\Sigma \in \mathcal{L}'$. (y is in a leaf since the union of the leaves of \mathcal{L}' is a closed subset of $\mathbf{R}^3 \setminus \mathcal{S}$ and \mathcal{S} does not intersect the open slab $\{0 < x_3 < \varepsilon\}$.) We can use the Harnack inequality to extend the local sequence of graphs converging to $\hat{\Sigma}$ to graphs over a sequence of expanding subdomains of $\hat{\Sigma}$, eventually obtaining a positive Jacobi field on the universal cover of Σ (cf. Lemma 2.1 in [CM04c]). In particular, the existence of such a positive Jacobi field implies stability of the universal cover of $\hat{\Sigma}$. (See, e.g., Proposition 1.26 in [CM99]; it is not needed here, but in fact one could pass only to a double cover by [MPR10].) There are now two possibilities: either Σ is complete (which is good) or it has singularities at p and/or q. In the second case, Lemma A.35 implies that these singularities are removable and, thus, $\hat{\Sigma}$ is contained in $\{x_3 = 0\}$; since this is impossible, we conclude that Σ is complete and

(III.1.11)
$$\hat{\Sigma} \cap \{x_3 = 0\} = \emptyset.$$

Combining the local curvature estimate for stable surfaces, [Sch83a], [CM02a], with the gradient estimate and (III.1.11) gives $\varepsilon_2 > 0$ (depending only on ε) so that

(III.1.12)
$$\{x_3 \le \varepsilon_2\} \cap \hat{\Sigma} \text{ is locally graphical over } \{x_3 = 0\}.$$

(See, e.g., Lemma I.0.9 in [CM04b] for a detailed proof of (III.1.12).) However, the theory of covering spaces (see Lemma 1.4 in [MR05]) now implies that each component of $\{x_3 \leq \varepsilon_2\} \cap \hat{\Sigma}$ is globally a graph of a function u with $0 < u \leq \varepsilon_2$ over a domain $\Omega \subset \{x_3 = 0\}$ with boundary values $u \mid_{\partial\Omega} = \varepsilon_2$. However,

¹⁹These are often referred to as "limit points" of a leaf in the literature.

the proof of the strong halfspace theorem of Hoffman-Meeks, [HM90], implies that such a u must be identically equal to its boundary values and, by unique continuation, we see that $\hat{\Sigma}$ is a horizontal plane as claimed. This completes the proof of (Fact 2) and, hence, also of properness.

We now completely understand what is going on near p or q and away from Ω_p and Ω_q , and we have proven properness even in these bad regions (above the vertices). The next step is to see that there are no leaves of \mathcal{L}' that live entirely in one of these bad regions. Because of the strict convexity and properness, there cannot be any complete leaves. It remains to show that there cannot be a leaf that has p or q in its closure. However, since we have properness, this is ruled out immediately by a catenoid barrier argument (Lemma 3.3 in [CM02c]) that gives the existence of low points.

The next step is to understand a neighborhood of p and q. Since we have properness, the two separately spiraling multi-graphs must combine in the bad region to form a single double-spiral staircase; this is for topological reasons since we must have an oppositely-oriented component between the two consecutive sheets in a single multi-graph. Thus, there is only leaf leaf that spirals in $\{x_3 = 0\}$ near p and q; we will call this Σ .

Finally, it remains to show that if $\Sigma' \subset \{x_3 > 0\}$ is a leaf of \mathcal{L} that contains a point $q' \in \{x_3 = 0\}$ in its closure $\overline{\Sigma'}$, then $\Sigma' = \Sigma$. This has already been established when q' is near p or q since there is only leaf there (cf. the chord arc property of Theorem I.1.6). However, this easily gives the general case. Namely, first fix a compact disk

(III.1.13)
$$\overline{D_S} = \{x_1^2 + x_2^2 \le S^2\} \cap \{x_3 = 0\}$$

containing p, q, and q'. Since, by assumption (see (1) in Proposition III.0.2), we have that $\overline{D_S} \cap S_{\text{neck}} = \emptyset$, an easy compactness argument and the convex hull property give some $r_1 > 0$ so that

(III.1.14) each component of $B_{r_1}(y) \cap \Sigma_j$ is a disk for every $y \in \overline{D_S}$.

Hence, by the one-sided curvature estimate, the Σ_j 's are locally graphical in a cylindrical slab about $\{x_3 = 0\}$ and over D_S , as long as we stay away from p and q. If we now choose a point \bar{q} in Σ' with $|\bar{q} - q'|$ sufficiently small (depending on both |p - q'| and r_1), then we can repeatedly apply the Harnack inequality to connect \bar{q} by a curve in Σ' back to a small neighborhood of p. Therefore, since Σ is the only leaf in $\{x_3 > 0\}$ that intersects a sufficiently small neighborhood of p, we conclude that $\Sigma' = \Sigma$.

The proof of (P2). The proof of the asymptotic graph structure in (P2) will be similar to the proof of property (D1) in Section III.0.6. First, we fix a large constant $\Omega > 1$ and some disk $D_R \subset \{x_3 = 0\}$ containing both p and q. Since $\{x_3 = 0\} \setminus \{p, q\}$ is a leaf of \mathcal{L}' (and, in particular, disjoint from \mathcal{S}), an

easy covering argument gives an arbitrarily small constant $\varepsilon' > 0$ and constants $\mu_A > 0$ and C_A so that

(III.1.15)
$$\sup_{\{0 < x_3 < \mu_A\} \cap \{x_1^2 + x_2^2 < 4\Omega^2 R^2\} \setminus (B_{c'}(p) \cup B_{c'}(q))} |A|^2 \le C_A$$

The gradient estimate and (III.1.15) then give a constant $\mu > 0$ (and less than (μ_A) so that

See Figure 26. Each point $y \in \{0 < x_3 < \mu\} \cap \Sigma$ over ∂D_R is contained in a graph $\Sigma_y \subset \Sigma$ over the annulus $D_{\Omega R} \setminus D_R$ with $\Sigma_y \subset \{0 < x_3 < \varepsilon\}$. Furthermore, the graph Σ_{y} extends over (a large part of) D_{R} as a graph, connecting to the two double spiral staircase structures near p and q.





This allows us to make the notion of a "low" end precise. Namely, a low end is the component of $\Sigma \setminus \Sigma_y$ — with Σ_y as above — whose closure does not contain the boundary graph over ∂D_R .

It remains to show that each of the graphs Σ_y extends as a graph indefinitely, i.e., past $\partial D_{\Omega R}$. This is where we argue as the proof of (D) in Section III.0.6, proving the following *scale invariant* ULSC property:

(D') There exists $\tau > 0$ so that, for $z \in \{x_3 = 0\} \setminus D_R$ and j large, each component of $B_{\tau|z|}(z) \cap \Sigma_i$ that connects to the multi-valued graph in Σ_i is a disk.

Once we have shown (D'), then the one-sided curvature estimate and gradient estimate will allow us to extend Σ_y as a graph indefinitely. We will leave this easy extension argument to the reader.

We will next prove (D') by contradiction; suppose therefore that some Σ_y connects outside D_R to a component of $B_{\tau|z|}(z) \cap \Sigma_j$ which is not a disk. (Here and below we are identifying Σ_y with the portion of the Σ_j 's converging to it. This identification should not lead to confusion since we have already shown that the convergence is multiplicity one.) We observe that this "nondisk component" leads to a contradiction, roughly following the proof of (D) in Section III.0.6 as follows:

• See Figure 27. A curve γ_i that is noncontractible in $B_{\tau|z|}(z) \cap \Sigma_i$ is also noncontractible in Σ_j by the convex hull property. We can therefore apply [MY82b] to find a stable surface $\Gamma_i \subset B_{R_i} \setminus \Sigma_i$ with interior boundary γ_i as in (1) in Section III.0.6. As before, using the plane $\{x_3 = 0\}$ allows us to conclude by Lemma III.0.3 that these stable surfaces quickly become graphs as we move away from the interior boundary γ_i .

• See Figure 28. Fix a graph $\Sigma_{y'}$ above Σ_y , and let $\sigma_{y'}$ be the component of $\partial \Sigma_{y'}$ over ∂D_R . As in (2) in Section III.0.6, we could then apply [MY82b] to put in a second stable surface $\Gamma_j \subset B_{R_j} \setminus (\Gamma_j \cup \Sigma_j)$ with interior boundary equal to $\sigma_{y'}$. Furthermore, this surface also quickly becomes graphical by Lemma III.0.3. Finally, the graph in Γ_j must start off between the two graphs Σ_y and Γ_j because (1) its interior boundary $\sigma_{y'}$ was chosen to be above Σ_y , and (2) every point near D_R in Σ that connects back to the multi-valued graph in Σ must be below Γ_j .

This easily gives the desired contradiction: The construction of Γ_j guarantees that there is a curve in $\Sigma_j \cup \Gamma_j$ connecting Σ_y to the graph in Γ_j and, moreover, is a graph over the x_1 -axis except for in a small neighborhood of z. The graph in $\tilde{\Gamma_j}$ is consequently forced to intersect this curve, giving the desired contradiction.

The final part of (P2), i.e., the nonnegativity of the vertical flux of the low ends, follows immediately since the ends are asymptotic to planes (vertical flux zero) or *upper*-halves of catenoids (vertical flux positive). This is because the only other possibility would be an end asymptotic to the *lower* half of a catenoid, which is impossible since one of these would eventually go below the plane $\{x_3 = 0\}$. (Recall that any embedded minimal end with finite total curvature is asymptotic to either a plane or half of a catenoid by Proposition 1 in [Sch83b].)







Figure 28. The proof of (D'): $\tilde{\Gamma}_j$ must intersect $\Gamma_j \cup \Sigma_j$.

54



Positive distance between consecutive ends allows us to find transverse (dotted) plane between two.

Figure 29. The vertically separating curve γ_+ .

The proof of (P3). To prove (P3), recall first that there is a positive distance between consecutive ends by the maximum principle at infinity of [LR88]. This positive distance allows us to intersect Σ with a horizontal plane that intersects Σ transversely between the heights of two consecutive ends, giving a finite collection γ_+ of disjoint simple closed curves separating these ends; see Figure 29. This finiteness follows from the compactness of the level set that in turn used the properness of Σ . Since we will be considering the part of Σ below this plane, the outward normal derivative of x_3 is nonnegative at every point along γ_+ . However, this plane was chosen to be transverse to the surface, so this derivative must in fact be pointwise positive along γ_+ . That is, the flux integrand is pointwise positive along γ_+ so the vertical flux across γ_+ is clearly positive.

Finally, we will sketch briefly why the fact that Σ has genus zero implies that we can choose a single component of γ^+ that is vertically separating. We must show that only one component of $\Sigma \setminus \gamma^+$ connects γ^+ to the "the ceiling" $\{x_3 = \varepsilon\}$. The point is that if there were two such components of $\Sigma \setminus \gamma^+$, then we could solve a sequence of Plateau problems to get a stable surface Γ^+ between them with the following properties:

- $\partial \Gamma^+ \subset \gamma^+$ is a finite collection of disjoint simple closed curves in a plane.
- Γ^+ does not go below the plane containing γ^+ .
- Γ^+ is above one of the components of $\Sigma \setminus \gamma^+$ and below the other.

It is not hard to see that this is impossible. The connectedness of γ^+ is not actually necessary for the proof, so we will leave the details for the reader.

The proof of (P4). The last claim (P4) essentially follows from the description of Σ near the plane $\{x_3 = 0\}$ and the gradient estimate. To see this, let

$$(\text{III.1.16}) \qquad \qquad \gamma_{pq} \subset \{x_3 = 0\}$$

be the line segment from p to q. Using the description of Σ near $\{x_3 = 0\}$ and the chord arc bound of Theorem I.1.6, we can find simple closed curves $\gamma_i \subset \Sigma$ with the following properties (see Figure 30):

• γ_i is contained in the ε_i -tubular neighborhood of γ_{pq} where $\varepsilon_i \to 0$.



Away from p and q, the closed curve γ_i has two components; these lie in oppositely oriented graphs in Σ .

Figure 30. The separating curves γ_i .

- $\gamma_i \setminus (B_{\varepsilon_i}(p) \cup B_{\varepsilon_i}(q))$ consists of two graphs over γ_{pq} that are in *oppositely*oriented sheets of Σ .
- The length of $(B_{\varepsilon_i}(p) \cup B_{\varepsilon_i}(q)) \cap \gamma_i$ is at most $C \varepsilon_i$ for a fixed constant C.

It is easy to see that the γ_i 's are vertically separating since the sheets containing

(III.1.17)
$$\gamma_i \setminus (B_{\varepsilon_i}(p) \cup B_{\varepsilon_i}(q))$$

are oppositely-oriented (meaning that the unit normals point in nearly opposite directions). Namely, if we fix a spiralling curve in either of the two multi-valued graphs near either singular curve, then the third component to the unit normal to Σ does not change sign along this curve. Consequently, such a spiralling curve intersects exactly one of the γ_i 's and does so exactly once. It follows that the γ_i 's are vertically separating as claimed.

We can now use the gradient estimate to bound $|\nabla x_3|$ along γ_i away from p and q to see that the vertical flux on

(III.1.18)
$$\gamma_i \setminus (B_\delta(p) \cup B_\delta(q))$$

goes to zero for any fixed $\delta > 0$. To get this bound, note first that the height function x_3 is positive and harmonic on the multi-valued graphs in Σ that spiral infinitely into $\{x_3 = 0\}$ and these graphs have bounded curvature away from p and q. Therefore, the gradient estimate (for positive harmonic functions) implies that

(III.1.19)
$$\sup_{\gamma_i \setminus (B_{\delta}(p) \cup B_{\delta}(q))} |\nabla x_3| \to 0 \text{ uniformly as } i \to \infty.$$

Combining this gradient bound with the bound on the length of

(III.1.20)
$$(B_{\delta}(p) \cup B_{\delta}(q)) \cap \gamma_i$$

gives the last claim. This completes the proof of the properties (P1)-(P4) of Σ .

Using (P1)–(P4) to deduce a contradiction. We now see that these properties are contradictory. Namely, by Stokes' theorem, the total flux across γ_i , γ_+ , and the "ends" of Σ between γ_i and γ_+ must sum to zero. However, the flux across γ_+ is positive and every other flux is either nonnegative or approaches zero. This contradiction shows that (III.1.6) could not have held, proving the lemma.

III.2. Completing the proof of Theorem 0.9

We will now use the properties of the singular set S_{ulsc} and lamination \mathcal{L}' to show that \mathcal{L}' is a foliation by parallel planes with two Lipschitz curves removed, thereby completing the proof of Theorem 0.9. The two main steps are

- using properness (Lemma III.1.4) to see that the (collapsed) planar leaves of \mathcal{L}' intersect every height,
- using the local cone property to get regularity of S_{ulsc} .

Proof of Theorem 0.9. Lemma II.1.2 gives a subsequence Σ_j , singular set \mathcal{S} , and lamination \mathcal{L}' of $\mathbf{R}^3 \setminus \mathcal{S}$ with minimal leaves. The set $\mathcal{S} = \mathcal{S}_{ulsc}$ is nonempty by assumption.

For each point x in S_{ulsc} , properties (1) and (2) of Proposition III.0.2 give a (collapsed) leaf of \mathcal{L}' that is a plane with two points removed. (x is one of the two points.) It follows easily from the convergence to these planes and the embeddedness of Σ_j that all of the limit planes are parallel so, after a rotation of \mathbf{R}^3 , we can assume that these planes are horizontal, i.e., given as level sets $\{x_3 = t\}$. Furthermore, since $S_{\text{ulsc}} \subset \mathbf{R}^3$ is a nonempty closed set, the local cone property implies that $x_3(S_{\text{ulsc}}) \subset \mathbf{R}$ is also closed (and nonempty).

We will show first that the collapsed leaves (or, rather, their closures) foliate \mathbf{R}^3 — more precisely, that

(III.2.1)
$$x_3(\mathcal{S}_{\text{ulsc}}) = \mathbf{R}.$$

To prove (III.2.1), we assume that $\{x_3 = t_0\} \cap S_{\text{ulsc}} = \emptyset$ for some $t_0 \in \mathbf{R}$ and will see that this leads to a contradiction. Namely, since $x_3(S_{\text{ulsc}})$ is closed, there exists $t_s \in x_3(S_{\text{ulsc}})$ which is a closest point in $x_3(S_{\text{ulsc}})$ to t_0 . The desired contradiction now easily follows from Lemma III.1.4 since either $\{t_s < x_3 < t_0\} \cap S_{\text{ulsc}}$ or $\{t_0 < x_3 < t_s\} \cap S_{\text{ulsc}}$ is empty. We conclude therefore that $x_3(S_{\text{ulsc}}) = \mathbf{R}$.

Finally, the Lipschitz regularity of the curves now follows as in Lemma I.1.2 of [CM04e]; the same argument applies with obvious minor modifications to deal with the fact that each horizontal plane now contains two singular points as opposed to just one in [CM04e].

TOBIAS H. COLDING and WILLIAM P. MINICOZZI II

III.3. Sequences with fixed genus

Most of the arguments in the preceding sections have assumed that the surfaces Σ_j have genus zero as opposed to just some fixed finite genus. The arguments for the genus zero case are slightly simpler, however, the modifications needed for the general case are straightforward. The key point is that the infinite multiplicity of the multi-valued graphs converging to a collapsed leaf means that there is an arbitrarily large number of disjoint curves to choose from that "circle both axes" and thus have the desired properties for the preceding arguments. In the general case of finite (but nonzero) genus, we can therefore follow the preceding argument using the following lemma:

LEMMA III.3.1. If Σ is oriented with genus g and $\sigma_1, \ldots, \sigma_{g+1} \subset \Sigma$ are disjoint simple closed curves, then $\Sigma \setminus \bigcup_i \sigma_i$ is disconnected.

Proof. The first integral homology group of Σ is 2g-dimensional, and the intersection form is a bilinear form of full rank (cf. Lemma I.0.9 of [CM04d]). Therefore the maximal subspaces on which the intersection form vanishes have dimension g. Consequently, there is a nontrivial linear (integral) relation between the σ_i 's and the lemma follows easily.

III.4. An application: A one-sided property for ULSC surfaces

The compactness theorem for ULSC sequences, Theorem 0.9, can be used to prove estimates for embedded minimal surfaces that have a lower bound on their injectivity radius. We will prove several such estimates in this paper, including Lemma III.4.1 in this section. This lemma proves a one-sided property for nonsimply connected surfaces on the smallest scale of nontrivial topology, showing that an intrinsic ball in such a surface cannot lie on one side of a plane and have its center close to the plane on this scale. This result requires that we work on this scale since, after all, large balls in the catenoid can be rescaled to lie above a plane and yet come arbitrarily close to the plane.

The proof of the lemma divides naturally into two extreme cases, depending on whether the (inverse of the) curvature is comparable to the injectivity radius or is much larger. In the first case, the surface looks more like a catenoid while in the second it looks like a pair of oppositely-oriented helicoids joined together. In the first case, the lemma essentially follows from the logarithmic growth of the ends of the catenoid; the second case follows from the fact that these double-helicoids converge to a foliation of all of \mathbb{R}^3 by the compactness theorem for ULSC sequences, Theorem 0.9.

LEMMA III.4.1. Given any H > 0, there exists $C_1 > H$ so that if Σ is an embedded minimal planar domain with $0 \in \Sigma$, $\mathcal{B}_{4r_1}(0) \subset \Sigma$ is not a disk,

 $\mathcal{B}_{C_1r_1}(0) \cap \partial \Sigma = \emptyset$, and $\mathcal{B}_{r_1}(x)$ is a disk for each $x \in \mathcal{B}_{C_1r_1}(0)$, then

Proof. After rescaling, we can assume that $r_1 = 1$. We will argue by contradiction, so suppose that Σ_j is a sequence of embedded minimal planar domains containing 0 with

- (III.4.3) $\mathcal{B}_{i}(0) \subset (\Sigma_{i} \setminus \partial \Sigma_{i}) \cap \{x_{3} \leq H\},\$
- (III.4.4) $\mathcal{B}_4(0)$ is not a disk,
- (III.4.5) $\mathcal{B}_1(x)$ is a disk for each $x \in \mathcal{B}_i(0)$.

Observe first that Lemma B.1 in Appendix B gives a sequence $R_j \to \infty$ so that the component Σ_{0,R_j} of $B_{R_j} \cap \Sigma_j$ containing 0 is compact and has boundary in ∂B_{R_j} .²⁰ Replacing the sequence Σ_j by Σ_{0,R_j} gives a ULSC sequence — still denoted Σ_j — of embedded minimal planar domains in extrinsic balls whose radius goes to infinity. We will now divide the proof into two cases depending on whether or not the curvatures of the sequence blows up.

Case 1. Suppose first that $\sup |A|^2 \to \infty$ in some fixed ball of \mathbb{R}^3 (for some subsequence). We can then apply the compactness theorem for ULSC sequences, Theorem 0.9, to get a subsequence of the Σ_j 's that converges to a foliation by parallel planes away from two lines orthogonal to the leaves of the foliation. Since the foliation is of all of \mathbb{R}^3 , this contradicts the upper bound for x_3 in (III.4.3).

Case 2. Suppose now that $|A|^2$ is uniformly bounded on compact subsets of \mathbf{R}^3 for every Σ_j . In this case, a subsequence of the Σ_j 's converges smoothly to a minimal lamination \mathcal{L} of \mathbf{R}^3 by Proposition B.1 in [CM04e]. We will first see that (III.4.3)–(III.4.5) imply that there is a nonflat leaf Γ of \mathcal{L} satisfying

- (III.4.6) $\Gamma \subset \{x_3 \le H\},\$
- (III.4.7) $\mathcal{B}_1(x)$ is a disk for each $x \in \Gamma$.

To see this, first note that (III.4.4) implies that each $\mathcal{B}_4(0) \subset \Sigma_j$ cannot be written as a graph over any plane and hence contains a point y_j with $|A|^2(y_j) > \delta_0$ for some $\delta_0 > 0$. A subsequence of these points converges to a point y in some leaf — call it Γ — of \mathcal{L} with $|A|^2(y) \ge \delta_0$. Thus Γ is not flat. Equation (III.4.6) follows immediately from (III.4.3).

Observe next that the leaf Γ

• is a multiplicity one limit of the Σ_i 's,

²⁰This will be needed later to apply the compactness results for ULSC sequences.

• is locally isolated in \mathcal{L} in the sense that each point $y \in \Gamma$ has a neighborhood $B_r(y)$ so that $B_r(y) \cap \mathcal{L}$ consists only of the component of $B_r(y) \cap \Gamma$ containing y.

If either of these were not the case, then the universal cover of Γ would be stable and, hence, flat; cf. the proof of Corollary A.20 for more details. As a consequence, if K is any compact subset of Γ , then for j sufficiently large (depending on K), Σ_j contains a normal graph K_j over K; as $j \to \infty$, the K_j 's converge smoothly to K. Using this convergence and the convex hull property, it is easy to see that (III.4.5) implies (III.4.7). Moreover, this convergence and the fact that the Σ_j 's are planar domains implies that Γ is also a planar domain. To see this, suppose instead that Γ contains a pair of nonseparating curves γ and $\tilde{\gamma}$ that have linking number one (i.e., so they are transverse and intersect at exactly one point). Then for j large, we would get a similar pair of curves in Σ_j ; since this is impossible for planar domains, we conclude that Γ is also a planar domain.

We have shown that Γ is a planar domain satisfying (III.4.6) and (III.4.7). We can assume that $H = \sup_{\Gamma} x_3$. Recall that [MR05] gives

(III.4.8)
$$\sup_{\Gamma \cap \{H-1 < x_3 < H\}} |A|^2 = \infty.$$

Namely, by the first paragraph of the proof of Lemma 1.5 in [MR05], if instead we had

(III.4.9)
$$\sup_{\Gamma \cap \{H-1 < x_3 < H\}} |A|^2 < \infty,$$

then $\Gamma = \{x_3 = H\}$. However, Γ is not flat, so we conclude that (III.4.8) must hold.

We will now use (III.4.8) to define a new sequence of planar domains where we can argue to a contradiction as in Case 1. (Alternatively, we could apply (4.b) in Theorem 4 in [MR06].) Namely, (III.4.8) gives a sequence of points $p_n \in \Gamma \cap \{H - 1 < x_3 < H\}$ with

(III.4.10)
$$|A|^2(p_n) \to \infty.$$

By (III.4.7), we can apply Lemma B.1 to conclude that the component $\Gamma_{p_n,n}$ of $B_n(p_n) \cap \Gamma$ containing p_n is compact and has boundary in $\partial B_n(p_n)$. Translate $\Gamma_{p_n,n}$ by moving p_n to the origin to get a ULSC sequence

(III.4.11)
$$\Gamma_n = \Gamma_{p_n,n} - p_n$$

of compact embedded minimal planar domains with $\Gamma_n \subset B_n$, $\partial \Gamma_n \subset \partial B_n$, and $|A|^2(0) \to \infty$. As in Case 1, a subsequence converges to a foliation of all of \mathbb{R}^3 by parallel planes away from two lines orthogonal to the leaves of the foliation. However, by (III.4.6), the translated surfaces Γ_n are in the half-space $\{x_3 < 1\}$.

60

This contradiction completes the proof of the lemma in Case 2, completing the proof of the lemma. $\hfill \Box$

Remark III.4.12. Using this one-sided property, we can go back and prove a stronger version of Lemma B.1. This stronger result gives that ULSC surfaces are proper — as opposed to just knowing that each component in a ball is proper.

Part IV. When the surfaces are not ULSC: The proof of Theorem 0.14

We will now turn to the case where the sequence is not ULSC and there is consequently no longer a lower bound for the injectivity radius of the Σ_j 's in compact subsets of \mathbf{R}^3 . As we did in the ULSC case, we will initially argue for the genus zero case and then explain the easy modifications needed for the general case of fixed finite genus.

We have already defined the singular set S in Definition/Lemma II.1.1 to be the set of points where the curvature blows up. Furthermore, Lemma II.1.2 gives a subsequence Σ_j that converges to a minimal lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$. This gives (A) and (B) in Theorem 0.14.

In the ULSC case, every singular point was essentially the same; namely, in a neighborhood of each singular point, the surfaces were double-spiral staircases. However, we now have the possibility that the injectivity radius of the Σ_j 's is going to zero at the singular point. This occurs, for example, by taking a sequence of rescalings of a catenoid or one of the Riemann examples. Recall that the Riemann examples are singly-periodic embedded minimal planar domains that are topologically — and conformally — equivalent to an infinite cylinder with a one-dimensional lattice of punctures.

For the sequence of rescaled catenoids, the singular set S consists of just the origin and we get C^{∞} convergence to a single plane with multiplicity two away from the origin. Rescaling one of the Riemann examples gives a line of singular points and convergence to a foliation by parallel planes away from this line. By choosing different sequences of rescaled Riemann examples, we can get different singular sets — but we always get a foliation by parallel planes.

The local behavior of the surfaces near a singular point is quite different, depending on whether or not the injectivity radius is going to zero there. To account for this, we define the subset $S_{\text{neck}} \subset S$ to be the set of points where the injectivity radius goes to zero. By Proposition I.0.19 of [CM04d], after passing to a further subsequence, S is given as the disjoint union

(IV.0.1)
$$S = S_{neck} \cup S_{ulsc}.$$

Recall that S_{ulsc} was defined in the introduction to be the set of points where the curvature blows up but where the sequence is locally ULSC.

An overview of this part. In Section IV.1, we prove the main structure result for the non-ULSC part of the limit lamination \mathcal{L}' , i.e, (C1) in Theorem 0.14. Namely, we show that for each point y in \mathcal{S}_{neck} , we get a sequence of graphs in the Σ_j 's that converges to a plane through y. These graphs will be defined over a sequence of expanding domains and the convergence will be smooth away from y and possibly one other point in the limit plane. In the process of proving this, we will also establish (P) in Theorem 0.14.

In Section IV.2, we prove the main structure results for the ULSC part of the lamination, i.e., (C2) and (D) in Theorem 0.14. Namely, we show that this part of the lamination is actually a foliation by an open set of parallel planes in \mathbf{R}^3 , the ULSC singular set S_{ulsc} is a collection of Lipschitz curves transverse to these planes, and S_{neck} does not intersect these planes.

In Section IV.3, we combine all of this and complete the proof of Theorem 0.14.

Remark IV.0.2. Theorem 0.14 gives a flat leaf of \mathcal{L}' through every singular point in \mathcal{S} but does not show that all of the leaves of \mathcal{L}' are flat. This will be proven in Part VI.

IV.1. Proving (C1) in Theorem 0.14: A plane through each point of S_{neck}

For each point y in S_{neck} , we will prove in this section that there is a sequence of graphs in Σ_j converging to a plane through y. The graphs converge smoothly to the plane away from y and possibly one other point. (The other point is also in S_{neck} .)

The key tool in this section is Proposition IV.1.1, which allows us to decompose an embedded minimal planar domain $\Sigma \subset B_{r_0}$ with $\partial \Sigma \subset \partial B_{r_0}$ by "cutting it" inside a small ball \mathcal{B}_{r_1} whenever some component of $\mathcal{B}_{r_1} \subset \Sigma$ is not a disk. Moreover, the proposition uses a barrier construction to find a stable graph disjoint from Σ so that the pieces of Σ are on opposite sides of this graph; see Figure 31. The basic example to keep in mind is the catenoid: Cutting the catenoid along the unit circle in the $\{x_3 = 0\}$ plane gives two pieces; these pieces are on opposite sides of the stable graph $\{x_3 = 0\} \cap \{x_1^2 + x_2^2 > 1\}$.

PROPOSITION IV.1.1. There exists a constant C > 1 so that the following holds:

Let $\Sigma \subset B_{r_0}$ be an embedded minimal planar domain with $\partial \Sigma \subset \partial B_{r_0}$ and $0 \in \Sigma$. If $\mathcal{B}_{r_1} \subset \Sigma$ is not a topological disk for some $r_1 < r_0/C^2$, then there exists a stable embedded minimal surface $\Gamma \subset B_{r_0} \setminus \Sigma$ with $\partial \Gamma \subset \partial B_{r_0} \cup \mathcal{B}_{20r_1}$ and satisfying the following properties:

(A) A component Γ_0 of $B_{r_0/C} \cap \Gamma \setminus B_{Cr_1}$ is a graph with gradient bounded by one and so that $\partial \Gamma_0$ intersects both $\partial B_{r_0/C}$ and ∂B_{Cr_1} .



Figure 31. Proposition IV.1.1: A stable graph Γ_0 separates Σ into parts above and below the graph.

(B) There are distinct components H^+ and H^- of $B_{r_0/C} \setminus (\Gamma_0 \cup B_{Cr_1})$, a separating curve $\tilde{\sigma} \subset B_{Cr_1} \cap \Sigma$, and distinct components Σ^+ and Σ^- of $B_{r_0/C} \cap \Sigma \setminus \tilde{\sigma}$ so that $\Sigma^{\pm} \subset H^{\pm} \cup B_{Cr_1}$ and $\tilde{\sigma} = \partial \Sigma^+ \cap \partial \Sigma^-$.

Figure 31 suggests that the separating curve $\tilde{\sigma}$ is also the inner boundary of Γ , but we are not claiming this in general. Rather, we will first choose an inner boundary curve in the very small ball B_{r_1} , use this to produce a stable graph that separates in $B_{r_0/C} \setminus B_{Cr_1}$ with components of Σ on both sides, and then choose the curve $\tilde{\sigma}$ to separate these upper and lower pieces of Σ in $B_{r_0/C} \setminus B_{Cr_1}$.

As mentioned above, Proposition IV.1.1 will be the key tool for getting the limit plane through each point of S_{neck} that was promised in (C1) in Theorem 0.14. To see why, we will first use Proposition IV.1.1 to get a sequence of stable graphs that are disjoint from Σ_j and converge, away from y, to a plane through y. Since the outer radii R_j go to infinity, applying Proposition IV.1.1 to the sequence Σ_j will give a sequence of stable graphs that are disjoint from Σ_j and defined over larger and larger annuli centered at y. As $j \to \infty$, the inner radii of these annuli go to zero and the outer radii go to infinity. Consequently, the stable graphs will converge (subsequentially) to a minimal graph over a plane punctured at y and this graph will have y in its closure. By a standard removable singularity theorem, this limit graph extends smoothly across y to an entire minimal graph and, hence, is flat by the Bernstein theorem. The easy details will be left to the reader.

Now that we have this stable limit plane through $y \in S_{neck}$, the proof of (C1) in Theorem 0.14 will consist of two main steps. We will sketch these two steps next. (They are proven in Sections IV.1.2 and IV.1.3, respectively.)

- (1) Decomposing Σ_j into ULSC pieces. Let Σ_j^+ be the portion of Σ_j above the stable graph. (We argue similarly for the part below.) There are two possibilities:
 - Σ_j^+ is scale-invariant ULSC away from y. Namely, there exist C' > 0and a sequence $r_j \to 0$ so that if $x \notin B_{C'r_j}(y)$, then each component of $B_{|x-y|/C'}(x) \cap \Sigma_j^+$ is a disk. (This intersection is empty when |x| is larger than $C'R_j$.)

• Otherwise, we can apply Proposition IV.1.1 to cut along a second noncontractible curve; see Figure 32.

In the second case, we replace Σ_j^+ by the portion Σ_j^{+-} of Σ_j^+ that is *below* the second stable graph. After repeating this a finite number of times, we will eventually get down to a scale-invariant ULSC subset of Σ_j with two interior boundary components (one component for the cut near y and one component for the last cut that we make).

(2) The ULSC pieces contain graphs. In either case, Lemma 3.3 in [CM02c] will then give low points in Σ_j on either side of these stable graphs (see Figure 38). Here "low points" roughly means points close to the stable graph but away from its boundary. The one-sided curvature estimate²¹ from [CM04e] and the gradient estimate will imply that the low points in the resulting ULSC subsets of Σ_j are graphical. Piecing this together will easily give the desired global graphs.

In the first case in (2) above, the graphs in the Σ_j 's will be defined over annuli; in the second case, the graphs will be over pairs of pants, i.e., over disks with two subdisks removed. We will refer to the second case as a "pair of pants" decomposition; see Figure 32.



Graphical annuli (dotted) separate the "pairs of pants."

Figure 32. A pair of pants decomposition near a point where the injectivity radius goes to zero.

Steps (1)-(2) above are modelled on similar arguments for topological annuli in [CM02c]. Some new complications will arise here because of the more complicated topological types of the surfaces, especially in the second case in step (2).

IV.1.1. The proof of Proposition IV.1.1: A decomposition near each point of S_{neck} . The next lemma will first give stable surfaces disjoint from Σ and with "interior boundary" contained in a small ball. In order to prove Proposition IV.1.1, we will later show that these stable surfaces contain the desired graphs. More precisely, the next lemma assumes that a component of a minimal planar domain Σ in a small ball is not a disk so that it must contain a

64

 $^{^{21}}$ The one-sided curvature estimate is recalled in this paper in Theorem I.1.3.

simple closed curve $\tilde{\gamma}$ separating two components σ_1 and σ_2 of $\partial \Sigma$. The lemma then uses the separating curve $\tilde{\gamma}$ as "interior boundary" for a Plateau problem to get a stable minimal surface "between" σ_1 and σ_2 ; see (C) in Lemma IV.1.2.

LEMMA IV.1.2. Let $\Sigma \subset B_{r_0}$ be an embedded minimal planar domain with $\partial \Sigma \subset \partial B_{r_0}$. If $r_1 < r_0$ and a component Σ_{r_1} of $B_{r_1} \cap \Sigma$ is not a topological disk, then the following four properties hold (see Figure 33):

- (A) $\Sigma \setminus \Sigma_{r_1}$ has at least two connected components; each of these components has at least one component of $\partial \Sigma$ in its closure.
- (B) If σ_1 and σ_2 are components of $\partial \Sigma$ that are separated by Σ_{r_1} (i.e., σ_1 and σ_2 are in the closure of distinct components of $\Sigma \setminus \Sigma_{r_1}$), then we can choose a simple closed curve

(IV.1.3)
$$\tilde{\gamma} \subset \partial \Sigma_{r_1} \subset \partial B_{r_1}$$

that separates σ_1 and σ_2 in Σ .

- (C) There is a component Ω of $B_{r_0} \setminus \Sigma$ and an embedded stable minimal surface $\Gamma \subset \Omega$ with interior boundary $\partial \Gamma \setminus \partial B_{r_0}$ equal to $\tilde{\gamma}$, and so that Γ separates σ_1 and σ_2 in Ω .²²
- (D) Γ is area-minimizing amongst surfaces in Ω with boundary equal to $\partial\Gamma$.



Figure 33. Lemma IV.1.2: If the planar domain Σ contains a closed noncontractible curve in the small ball B_{r_1} , then Σ has distinct boundary components σ_1 and σ_2 . Moreover, there is a stable surface Γ that is disjoint from Σ and separates σ_1 and σ_2 in a component Ω of $B_{r_0} \setminus \Sigma$. The boundary of Γ has two parts, an outer boundary in ∂B_{r_0} and an inner boundary curve $\tilde{\gamma} \subset \partial B_{r_1} \cap \Sigma$.

²²One must be careful interpreting this "separation" since $\partial\Gamma$ may intersect σ_1 or σ_2 . In this case, we mean that Γ separates points in the interior of Σ that are arbitrarily close to σ_1 and σ_2 .

Proof. The first two claims, i.e., (A) and (B), follow from the fact that Σ is a planar domain and, by the maximum principle, any homologically nontrivial curve in Σ_{r_1} is also homologically nontrivial in Σ .

We will next solve the Plateau problem to get the desired stable surface in the complement of Σ . To do this, we need to choose the boundary of the stable surface and decide which of the two components of $B_{r_0} \setminus \Sigma$ that we will solve in.

To get the boundary, simply let Σ_0 be the component of $\Sigma \setminus \tilde{\gamma}$ that contains σ_1 in its boundary; we will minimize area among surfaces with boundary equal to

(IV.1.4)
$$\partial \Sigma_0$$
.

Note that $\partial \Sigma_0$ has "interior boundary" equal to $\tilde{\gamma}$ and σ_2 is not in $\partial \Sigma_0$.

We will use a simple linking argument to choose the domain Ω to solve in. First, fix a smooth curve $\eta \subset \Sigma$ from σ_1 to σ_2 that intersects $\tilde{\gamma}$ exactly once and does so transversely. (Such a curve exists since $\tilde{\gamma}$ separates σ_1 and σ_2 in Σ .) Since Σ is compact and embedded, we can "push η off of Σ " — on either side of Σ — to get curves η^+ and η^- that are disjoint from Σ and in distinct components of $B_{r_0} \setminus \Sigma$. It follows that the (mod 2) linking numbers of η^+ and η^- , respectively, with $\tilde{\gamma}$ differ by one.²³ In particular, one of these — say η^- — has linking number 1 (mod 2) with $\tilde{\gamma}$. Let Ω be the component of $B_{r_0} \setminus \Sigma$ that contains the *other* curve η^+ .

It follows that we have the following three properties:

- The domain Ω is mean convex in the sense of [MY82b].
- $\partial \Sigma_0$ is contained in $\partial \Omega$ and bounds the planar domain Σ_0 in $\partial \Omega$.
- $\tilde{\gamma} = \partial \Sigma_0 \setminus \partial B_{r_0}$ has linking number 1 (mod 2) with the curve η^- that is not in Ω .

Using the first two properties, a result of Hardt-Simon, [HS79], gives an embedded minimal surface $\tilde{\Gamma} \subset \Omega$ with $\partial \tilde{\Gamma} = \partial \Sigma_0$ and so that $\tilde{\Gamma}$ minimizes area amongst surfaces in Γ with the same boundary.²⁴ In particular, $\tilde{\Gamma}$ must be stable.

66

²³Recall that if $\eta \subset B_r$ is a curve with endpoints in ∂B_r and $\gamma \subset B_r$ is a closed curve, then their linking number is defined to be the number of times (mod 2) that η intersects a surface $\Gamma \subset B_r$ with $\partial \Gamma = \gamma$. As usual, we assume that Γ and η intersect transversely when counting intersections. The point is that this number does not depend on the particular choice of bounding surface Γ .

²⁴We could of course have applied a result of Meeks-Yau to get a stable planar domain. However, this planar domain would have minimized area only amongst planar domains. We will later use this minimizing property to bound the area of Γ by constructing comparison surfaces. It will be convenient not to have to restrict the topological type of the comparison surfaces.

The surface $\tilde{\Gamma}$ may have several components. We will use the third property above to show that the component Γ containing $\tilde{\gamma}$ in its boundary separates σ_1 and σ_2 in Ω and, hence, satisfies (C). First of all, $\tilde{\gamma}$ alone cannot be the entire boundary of $\partial \Gamma$; indeed, any surface $\Gamma_{\tilde{\gamma}} \subset B_{r_0}$ with $\partial \Gamma_{\tilde{\gamma}} = \tilde{\Gamma}$ would be forced to intersect the curve η^- that is not in Ω . Therefore, we must have that

(IV.1.5)
$$\partial B_{r_0} \cap \partial \Gamma \neq \emptyset$$

A similar use of the linking condition implies that Γ separates σ_1 and σ_2 in Ω , giving (C).

The minimizing property given in (D) will be used to give an upper bound for the area of Γ by constructing comparison surfaces with the same boundary. To carry this out, we will need two elementary lemmas. The first is a simple topological lemma showing that any collection of disjoint simple closed curves in a planar domain is homologous to a collection of *distinct* boundary curves. Moreover, together the initial curves and the boundary curves bound a subdomain of the planar domain. The second lemma uses this to construct comparison surfaces and hence, using (D), deduce an area bound for the surface Γ above.

LEMMA IV.1.6. Let P be a (possibly disconnected) compact planar domain with boundary ∂P . Given any collection $\sigma_1, \ldots, \sigma_n$ of disjoint simple closed curves in P, then there is a subdomain $P_0 \subset P$ with

(IV.1.7)
$$\partial P_0 = \left(\cup_{i=1}^n \sigma_i\right) \cup \left(\cup_{i=1}^m \eta_i\right),$$

where η_1, \ldots, η_m are distinct components of ∂P .

Remark IV.1.8. Before giving the proof of Lemma IV.1.6, it may be helpful to make two remarks. First, it is possible that m = 0 in IV.1.7, i.e., that $\partial P_0 = \bigcup_{i=1}^n \sigma_i$. Second, notice that all of the above curves — both the σ_i 's and η_i 's — are thought of as *unoriented* curves.

Proof of Lemma IV.1.6. Since we can consider each connected component of P separately, we may as well assume that $P \subset \mathbf{R}^2$ is connected. The set P_0 will be given as the level set $f^{-1}(1)$ of a map

(IV.1.9)
$$f: P \setminus (\bigcup_{i=1}^{n} \sigma_i) \to \{-1, +1\}.$$

To define f, first fix a point $p_0 \in P \setminus (\bigcup_{i=1}^n \sigma_i)$. For each point $p \in P \setminus (\bigcup_{i=1}^n \sigma_i)$, choose a curve $\gamma_p \subset P$ from p_0 to p that is transverse to the σ_i 's and let n(p) be the number of times that γ_p crosses $(\bigcup_{i=1}^n \sigma_i)$; see Figure 34. It follows from elementary topology that n(p) (mod two) does not depend on the choice of the curve γ_p and, hence, we can define the function f by

(IV.1.10)
$$f(p) = (-1)^{n(p)}$$



f changes sign each time we cross a (dotted) σ_i .

Figure 34. The proof of Lemma IV.1.6: The curves σ_i are dashed while ∂P is solid. After fixing a basepoint p_0 , we define a function f(p) to be +1 or -1, depending on whether you have to cross an even or odd number of the σ_i 's to get from p_0 to p. The domain P_0 is then defined to be the level set where f is +1. In the example pictured, P_0 has two components.

Define P_0 by

(IV.1.11)
$$P_0 = \{p \mid f(p) = +1\}.$$

It follows easily that f changes sign as we cross each σ_i and, therefore, each $\sigma_i \subset \partial P_0$ as desired.

LEMMA IV.1.12. Given a constant C_1 , there exists $C_2 > C_1$ so that if $\Sigma \subset B_{r_0}$, $\tilde{\gamma}$, and $\Gamma \subset \Omega \subset B_{r_0} \setminus \Sigma$ are as in Lemma IV.1.2 and for some r between r_0 and r_1 we have

(IV.1.13)
$$r^{-1} \operatorname{Length}(\partial B_r \cap \Sigma) \leq C_1,$$

then we get an area bound for $B_r \cap \Gamma$

(IV.1.14)
$$r^{-2}\operatorname{Area}(B_r \cap \Gamma) \leq C_2.$$

Proof. Note first that (IV.1.13) and Stokes' theorem (using that $\operatorname{div}(\nabla |x|^2) = 4$ on a minimal surface) give

(IV.1.15)
$$r^{-2} \operatorname{Area} (B_r \cap \Sigma) \le (2r)^{-1} \operatorname{Length} (\partial B_r \cap \Sigma) \le C_1/2.$$

The minimizing property (D) in Lemma IV.1.2 implies that $B_r \cap \Gamma$ is itself area minimizing among surfaces in its homology class in Ω . It will suffice therefore to construct a comparison surface in $\overline{\Omega}$ with bounded area. This follows from the following steps:

(1) The outer boundary of $B_r \cap \Gamma$ — i.e., $(\partial B_r) \cap \Gamma$ — sits inside the (possibly disconnected) planar domain $\Omega \cap \partial B_r$. Consequently, Lemma IV.1.6 gives a subset P_0 of the planar domain $\Omega \cap \partial B_r$ with

(IV.1.16)
$$(\partial B_r \cap \Gamma) \subset \partial P_0 \text{ and } \partial P_0 \setminus (\partial B_r \cap \Gamma) \subset (\partial B_r) \cap \Sigma.$$

Note that P_0 has bounded area since it is contained in ∂B_r .

- (2) Let Σ_0 denote the component of $B_r \cap \Sigma$ containing $\tilde{\gamma}$, and let Σ_0^+ be one of the components of $\Sigma_0 \setminus \tilde{\gamma}$. Note that Σ_0^+ has bounded area by (IV.1.15).
- (3) By the previous two steps, the boundary of $P_0 \cup \Sigma_0^+$ contains all of $\partial(B_r \cap \Gamma)$ (both $\tilde{\gamma}$ and the outer boundary). Ideally, we would have that $\partial(P_0 \cup \Sigma_0^+) = \partial(B_r \cap \Gamma)$ so that $P_0 \cup \Sigma_0^+$ would be a valid comparison surface. However, this does not have to be the case since $\partial(P_0 \cup \Sigma_0^+)$ may have some additional components. We can therefore assume that

(IV.1.17)
$$\partial(P_0 \cup \Sigma_0^+) \setminus \partial(B_r \cap \Gamma) \neq \emptyset.$$

(4) Observe that $\partial(P_0 \cup \Sigma_0^+) \setminus \partial(B_r \cap \Gamma)$ is itself the boundary of a surface in $\overline{\Omega}$, namely of the surface $P_0 \cup \Sigma_0^+ \cup (B_r \cap \Gamma)$. We can therefore solve the Plateau problem for a surface Γ in $B_r \cap \Omega$ with

(IV.1.18)
$$\partial \Gamma = \partial (P_0 \cup \Sigma_0^+) \setminus \partial (B_r \cap \Gamma) \subset (\partial B_r) \cap \Sigma$$

The length bound on $(\partial B_r) \cap \Sigma$ and the isoperimetric inequality for minimal surfaces then give an area bound for $\tilde{\Gamma}$.

(5) Finally, it follows that

(IV.1.19)
$$\partial(P_0 \cup \Sigma_0^+ \cup \Gamma) = \partial(B_r \cap \Gamma),$$

so that $P_0 \cup \Sigma_0^+ \cup \tilde{\Gamma}$ is the desired comparison surface.²⁵ This gives (IV.1.14) since each of the three pieces of the comparison surface has the desired area bound.

The next lemma gives an area estimate for the components of the stable surface constructed in Lemma IV.1.2 on the largest scale r_1 where Σ is ULSC. Recall that Σ contains a noncontractible curve $\tilde{\gamma}$ in ∂B_{r_1} , so Σ is not ULSC on scales larger than r_1 . On the other hand, assumption (IV.1.21) below gives that Σ is ULSC on this scale.

LEMMA IV.1.20. Given a constant $C_1 > 8$, there exists $C_2 > C_1$ so that if $0 \in \Sigma \subset B_{r_0}$, $\tilde{\gamma}$, and $\Gamma \subset \Omega \subset B_{r_0} \setminus \Sigma$ are as in Lemma IV.1.2 with $r_1 < r_0/C_2$ and, in addition,

(IV.1.21)
$$\mathcal{B}_{r_1/4}(x) \subset \Sigma \text{ is a disk for every } x \in \mathcal{B}_{C_2 r_1},$$

then each component Γ' of $B_{C_1 r_1} \cap \Gamma \setminus B_{8 r_1}$ that can be connected to $\tilde{\gamma}$ by a curve in $B_{17r_1} \cap \Gamma$ satisfies

(IV.1.22)
$$\operatorname{Area}\left(\Gamma'\right) \le C_2 r_1^2.$$

Proof. Observe first that the chord-arc bound for ULSC surfaces, Lemma B.1 in Appendix B, shows that the ULSC hypothesis (IV.1.21) also holds for

 $^{^{25}}$ This surface is not embedded and it may not even have the same topological type, but it is nonetheless a valid comparison surface.

x in the component of Σ containing 0 in an extrinsic ball $B_{C'_2 r_1}$ where C'_2 goes to infinity as C_2 does. In particular, after replacing Σ by this component, we may as well assume that

(IV.1.23)
$$\mathcal{B}_{r_1/4}(x)$$
 is a disk for every $x \in \Sigma$.

The proof of (IV.1.22) will be by contradiction, using a compactness argument. Suppose therefore that Σ_j , Γ_j is a sequence of counter-examples where (IV.1.22) fails with $C_2 = j \rightarrow \infty$. After translating and rescaling, we may assume that $r_1 = 4$.

We will consider two cases depending on whether $|A|^2 \to \infty$ on a compact set for the sequence Σ_j .

Case 1. Suppose first that $|A|^2 \to \infty$ in some fixed ball of \mathbb{R}^3 (for some subsequence of the Σ_j 's). We can then apply the compactness theorem for nonsimply connected ULSC sequences, Theorem 0.9, to get a subsequence of the Σ_j 's that converges to a foliation by parallel planes away from two lines orthogonal to the leaves of the foliation.

Observe that both of these orthogonal "singular" lines intersect the leaf through 0 inside the ball B_4 since Σ_j is assumed to be nonsimply connected inside that ball. It follows easily from the description of the convergence near the lines (as oppositely-oriented double spiral staircases) that any such component Γ'_j is sandwiched between the almost planar leaves of the foliation (for j sufficiently large). This sandwiching, together with interior curvature estimates for stable surfaces, implies that Γ'_j is itself a graph and hence has bounded area as desired.

Case 2. Suppose now that $|A|^2$ is uniformly bounded on compact subsets of \mathbf{R}^3 for every Σ_j . In this case, we will get uniform area bounds for the surfaces Σ_j in the ball B_{4C_1} . Once we have these area bounds for the Σ_j 's, then the comparison argument in Lemma IV.1.12 will give a uniform bound for area of the Γ_j 's in the same ball. However, we assumed that there was no such area bound for Γ_i 's. This contradiction will complete the proof of the lemma.

Therefore, to complete the proof of the lemma, it suffices to bound the area of Σ_j in B_{4C_1} . This area bound follows immediately from combining the following two facts:

• By Lemma B.3 in Appendix B, the uniform curvature bound implies uniform area bounds for each component of Σ_j in extrinsic balls. (The bound depends on the ball but not on j.) More precisely, if $\Sigma_{j,R}$ is a component of $B_R \cap \Sigma_j$, then Lemma B.3 implies that

(IV.1.24)
$$\operatorname{Area}\left(\Sigma_{j,R}\right) \le C_c R^2,$$

where the constant C_c depends only on the supremum of $|A|^2$ on $B_{C_0 R} \cap \Sigma_j$. The constant C_0 here is universal and does not depend on the upper bound for the curvature.

Even though each intrinsic ball of radius one in Σ_j is a disk, there is a component of B₄∩Σ_j that is not a disk. Therefore, it follows easily from a barrier argument and the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, that there exists R > 4C₁ so that only one component of B_R ∩ Σ_j intersects B_{4C1} for j large.

This follows exactly as does the analogous result for disks given in Corollary 0.4 in [CM04e]. Namely, if there were two such components, then we could put a stable surface between them. As R gets larger, this stable surface must be very flat and, thus, we have that each component of the Σ_j 's in a large ball $B_{R'}$ (for $R' \gg R$) is essentially contained on one side of a plane. Furthermore, this plane passes through B_{4C_1} . However, the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, with $r_1 = 1$ implies that the nonsimply connected component must grow out of this half-space in a ball of a fixed radius R'' that depends only on C_1 . This gives the desired contradiction for R sufficiently large, so we conclude that there could not have been two such components.

The last result that we will need to recall before proving Proposition IV.1.1 is the following elementary property of connected planar domains:

LEMMA IV.1.25. Let Σ be a connected planar domain and $\sigma_1, \ldots, \sigma_n$ the components of $\partial \Sigma$. Given k < n and a collection $\{\sigma_{i_1}, \ldots, \sigma_{i_k}\}$, there is a simple closed curve $\tilde{\sigma} \subset \Sigma$ that separates $\cup_{j \leq k} \sigma_{i_j}$ from $\partial \Sigma \setminus \bigcup_{j \leq k} \sigma_{i_j}$.

Proof of Proposition IV.1.1. We will first use a rescaling argument to locate the smallest scale of nontrivial topology, choose a noncontractible curve γ on this scale, and then solve the Plateau problem with γ as interior boundary. We will then obtain an area bound for the components of Γ on this scale. This area bound will allow us to apply the "stable graph proposition," Proposition C.2, to get the graph $\Gamma_0 \subset \Gamma$ and thus prove (A). Finally, in the last step of the proof, we will find the separating curve $\tilde{\sigma}$ and prove (B).

Blowing up on the smallest scale of nontrivial topology. Fix a large constant $C_1 > 5$ to be chosen. Applying the blow up lemma, Lemma D.1 in Appendix D, at 0 gives an intrinsic ball

(IV.1.26)
$$\mathcal{B}_{C_1s_1}(y_1) \subset \mathcal{B}_{5C_1r_1},$$

so that $\mathcal{B}_{4s_1}(y_1)$ is not a disk but $\mathcal{B}_{s_1}(y)$ is a disk for each $y \in \mathcal{B}_{C_1s_1}(y_1)$. We can now use this topologically nontrivial region of Σ to solve a Plateau problem. Namely, applying Lemma IV.1.2 to the component of $B_{4s_1}(y_1) \cap \Sigma$ containing $\mathcal{B}_{4s_1}(y_1)$ gives a simple closed noncontractible²⁶ curve $\gamma \subset B_{4s_1}(y_1) \cap \Sigma$, a mean convex domain $\Omega \subset B_{r_0} \setminus \Sigma$, and a stable embedded minimal surface

(IV.1.27)
$$\Gamma \subset \Omega$$
,

with interior boundary $\partial \Gamma \setminus \partial B_{r_0}$ equal to γ . Moreover, there are distinct components σ_1 and σ_2 of $\partial \Sigma \subset \partial B_{r_0}$ that are separated in Σ by γ and separated in Ω by Γ . Finally, Γ is area-minimizing amongst surfaces in Ω with boundary equal to $\partial \Gamma$.

An area estimate on the smallest scale of nontrivial topology. Suppose that Γ' is a component of $B_{24s_1}(y_1) \cap \Gamma \setminus B_{8s_1}(y_1)$ that can be connected to γ by a curve in $B_{17s_1}(y_1) \cap \Gamma$. If the constant C_1 from the previous step is sufficiently large (independent of Σ and Γ), then Lemma IV.1.20 gives a constant C_2 so that

(IV.1.28)
$$\operatorname{Area}\left(\Gamma'\right) \le C_2 \, s_1^2.$$

Finding the graph in Γ . Using the area bound (IV.1.28), we can apply Proposition C.2 to get that each component of $B_{r_0/\omega} \cap \Gamma \setminus B_{20\omega r_1}$ is a graph, where the constant $\omega > 1$ comes from Proposition C.2 and we have used that $4s_1 \leq 20r_1$ to estimate the inner radius.

A linking argument as in Lemma IV.1.2 then implies that one of these components Γ_0 has the property that $\Gamma_0 \cup B_{20\omega r_1}$ separates $B_{r_0/C}$ into components H^+ above and H^- below Γ_0 where $\sigma_1 \subset H^+$ and $\sigma_2 \subset H^{-27}$.

Finding the separating curve. To complete the proof, we need only find the separating curve $\tilde{\sigma} \subset \Sigma$ and prove (B). In doing this, we will increase the inner radius and decrease the outer radius.

The key to this step is to prove that there is a constant $C > 20\omega$ so that

(IV.1.29) only one component of $B_{Cr_1} \cap \Sigma$ intersects both H^+ and H^- .

Before proving (IV.1.29), it may be helpful to make a few remarks. First, it is not hard to see that (IV.1.29) is necessary to establish (B).²⁸ Second, it is easy to see that there must be at least one component of $B_{Cr_1} \cap \Sigma$ that intersects both H^+ and H^- . This is because Σ has boundary components $\sigma_1 \subset H^+$ and $\sigma_2 \subset H^-$ and the only way to connect these without crossing the annular graph

72

²⁶Since Σ has nonpositive curvature and γ is noncontractible in the intrinsic ball $\mathcal{B}_{4s_1}(y_1)$, it is also noncontractible in Σ .

²⁷Technically, this is not quite right since $\partial \Sigma$ is contained in the boundary of the larger ball B_{r_0} . Rather, the linking argument gives two components — call them $\tilde{\sigma}_1$ and $\tilde{\sigma}_2$ — of $\partial B_{r_0/C} \cap \Sigma$ that are separated by Γ_0 .

²⁸Namely, if there were two distinct components of $B_{Cr_1} \cap \Sigma$ that each connected H^+ and H^- , then it would be impossible to find a single connected curve in $B_{Cr_1} \cap \Sigma$ that separates H^+ and H^- .
Γ_0 is to go through the "hole" in the middle. Finally, the basic idea behind (IV.1.29) is that if there were two components passing through the "hole" in Γ_0 , then a barrier argument would also give a stable surface between the two components that also passes through the hole. However, such a stable surface would have to be very flat if C is much larger than 20ω , so it cannot pass through this hole. This is essentially the argument that we will give below, but it will take a little work to make it precise.

Once we establish (IV.1.29), the rest of the proof of the proposition will follow easily. Namely, if we let $\hat{\Sigma}$ be the component of $B_{Cr_1} \cap \Sigma$ intersecting both H^+ and H^- , then Lemma IV.1.25 gives a simple closed curve $\tilde{\sigma} \subset \hat{\Sigma}$ separating $H^+ \cap \partial \hat{\Sigma}$ from $H^- \cap \partial \hat{\Sigma}$. In particular, the components Σ^{\pm} of $B_{r_0/C} \cap \Sigma \setminus \tilde{\sigma}$ that intersect $\hat{\Sigma}$ must satisfy $\Sigma^{\pm} \subset H^{\pm} \cup B_{Cr_1}$.

Finally, to complete the proof of the proposition, it remains only to prove (IV.1.29). We will do this by contradiction, so suppose that $\hat{\Sigma}_1$ and $\hat{\Sigma}_2$ are distinct components of $B_{Cr_1} \cap \Sigma$ each of which intersects both H^+ and H^- . We will show that this is impossible for C large enough.

Since the only "hole" in the graph Γ_0 is in $B_{20\omega r_1}$, there must be components $\tilde{\Sigma}_i \subset \hat{\Sigma}_i$ of $B_{40\omega r_1} \cap \Sigma$ each of which intersects all three of H^+ , H^- , and $B_{20\omega r_1}$; see Figure 35. Label these components so $\gamma \cap \hat{\Sigma}_2 = \emptyset$. To get the contradiction, we will solve a Plateau problem to get a second stable graph that is between the $\tilde{\Sigma}_i$'s and also disjoint from the graph Γ_0 ; such a graph would be forced to sit on one side of Γ_0 and hence would not allow both of the $\tilde{\Sigma}_i$'s to get to both H^+ and H^- ; see Figure 36.

To set this up, note first that we can assume that Γ_0 is a graph with arbitrarily small gradient — say at most $\delta > 0$ — after possibly increasing C. This follows from estimates for minimal graphs; see Proposition 1.12 in [CM01]. After a rotation of \mathbf{R}^3 , we can assume that Γ_0 is a graph over the horizontal plane $\{x_3 = 0\}$.



Figure 35. The key step in finding the separating curve: Ruling out that two components of $B_{Cr_1} \cap \Sigma$ both intersect both H^+ and H^- .



Second stable graph is forced to cross Σ .

Figure 36. The contradiction: We cannot have a second stable graph that is between the two components of Σ and on one side of Γ_0 .

Fix a point y_1 in $B_{20\omega r_1} \cap \tilde{\Sigma}_1$, and choose a point y_2 in $\tilde{\Sigma}_2$ so that the segment γ_{y_1,y_2} from y_1 to y_2 is "almost vertical"; see Figure 35. More precisely, applying Lemma A.8 of [CM04b] (as in (I.0.20) of [CM04b]) gives $y_2 \in \tilde{\Sigma}_2$ with

(IV.1.30)
$$|\Pi(y_2 - y_1)| \le |y_2 - y_1| \cos \theta_0,$$

where Π is orthogonal projection to the horizontal plane and the constant θ_0 is defined in the appendix of [CM04b]. It now follows that there is a component $\tilde{\Omega}$ of $B_{Cr_1} \setminus (\Gamma \cup \Sigma)$ so (some subsegment of) γ_{y_1,y_2} is linked with $\partial \hat{\Sigma}_2$ in $\tilde{\Omega}$. Note that $\tilde{\Omega}$ is mean convex in the sense of [MY82b]. A result of [MY82a]–[MY82b] gives a stable embedded minimal surface

(IV.1.31)
$$\hat{\Gamma}^0 \subset \tilde{\Omega}$$

with $\partial \hat{\Gamma}^0 = \partial \hat{\Sigma}_2$. Now, we choose the constant C large (depending only on ω) so that the curvature estimates [Sch83a], [CM02a] for stable surfaces with the fact that $\hat{\Gamma}^0$ is disjoint from the graph Γ_0 implies that each component of $B_{100\omega r_1} \cap \hat{\Gamma}^0$ that intersects γ_{y_1,y_2} must be a very flat graph that does so exactly once (see (IV.1.30)). Since $\partial \hat{\Sigma}_2$ and γ_{y_1,y_2} are linked in $\tilde{\Omega}$, a component $\hat{\Gamma}$ of $B_{100\omega r_1} \cap \hat{\Gamma}^0$ intersects γ_{y_1,y_2} at least once. In particular, $\hat{\Gamma}^0$ separates $\tilde{\Sigma}_1$ and $\tilde{\Sigma}_2$, forcing one of these to lie on the same side of Γ_0 as does $\hat{\Gamma}^0$. This gives the desired contradiction; see Figure 36. Consequently, we conclude that $B_{Cr_1}(x_1) \cap \Sigma$ contains only one component $\hat{\Sigma}$ that intersects both H_1^+ and H_1^- ; i.e., (IV.1.29) holds.

IV.1.2. Step (1): Decomposing Σ_j into ULSC pieces. Suppose now that $0 \in S_{neck}$, so that Proposition IV.1.1 gives

(1) A sequence of stable graphs Γ_j that are disjoint from Σ_j and that converge to a punctured plane through 0; after rotating \mathbf{R}^3 , we can assume that the stable graphs converge to $\{x_3 = 0\} \setminus \{0\}$.

(2) A sequence of closed curves $\tilde{\sigma}_j \subset B_{r_j} \cap \Sigma_j$ with $r_j \to 0$ and so that $\tilde{\sigma}_j$ divides Σ_j into a component Σ_j^+ above $\{x_3 = 0\}$ and a component Σ_j^- below.²⁹

We will show next that each Σ_j^+ contains a large scale-invariant ULSC piece Σ_j^{ulsc} . Before stating this precisely, it may be helpful to recall two simple examples:

- If we consider a sequence of shrinking catenoids, then 0 is the only point in S_{neck} and each half of the catenoids is easily seen to be scale-invariant ULSC. (In fact, given any point $x \neq 0$ in one of the catenoids, the ball $B_{|\underline{x}|}(x)$ has two simply connected components.)
- Consider now a sequence of rescalings of one of the Riemann examples. In this case, S_{neck} is a line through the origin and Σ_j^+ is *not* scale-invariant ULSC. However, if we cut Σ_j^+ along a second short curve (the "neck" immediately above the first curve), then the resulting "pair of pants" is scale-invariant ULSC with respect to the distance to the closer of the two necks.

The precise statement of the decomposition into ULSC pieces is given in the next lemma. For simplicity, we will suppose that $0 \in S_{\text{neck}}, \Sigma_j^+ \subset \Sigma_j$ are as above, and the constant C is given by Proposition IV.1.1.

LEMMA IV.1.32. Let $\tilde{\Sigma}_j^+$ denote the connected component of $B_{R_j/(2C)} \cap \Sigma_j^+$ with $\tilde{\sigma}_j$ in its boundary, and fix a constant $\alpha > 1$. For each j sufficiently large, one of the following two cases holds:

- (1) $\tilde{\Sigma}_{j}^{+}$ is scale-invariant ULSC: Given any $x \notin B_{\alpha C r_{j}}$, then each component of $B_{|x|/(\alpha C)}(x) \cap \tilde{\Sigma}_{j}^{+}$ is a disk.
- (2) $\tilde{\Sigma}_{j}^{+}$ contains a noncontractible curve $\tilde{\sigma}_{j}^{+}$ in a ball $B_{s_{j}}(y_{j})$ with
- (IV.1.33) $|y_j| > \alpha C r_j \text{ and } s_j < |y_j|/(C\alpha),$

so that the component Σ_j^{ulsc} of $\tilde{\Sigma}_j^+ \setminus \tilde{\sigma}_j^+$ with $\tilde{\sigma}_j$ in its boundary is scaleinvariant ULSC: Given any $x \notin (B_{\alpha C r_j} \cup B_{\alpha C s_j}(y_j))$, then each component of

(IV.1.34)
$$B_{\frac{\min\{|x|, |x-y_j|\}}{C\alpha}}(x) \cap \Sigma_j^{\text{ulsc}}$$

is a disk.

Proof. The key for establishing this lemma is that the decomposition into a Σ_i^+ and a Σ_i^- can be repeated anywhere that the topology is concentrating.

²⁹More precisely, there are shrinking extrinsic balls B_{r_j} so that $\Sigma_j^+ \setminus B_{r_j}$ is above Γ_j and similarly for Σ_j^- .

Namely, suppose that (1) does not hold and, hence, there exists some z_1 in $B_{R_i/(2C)} \setminus B_{\alpha C r_i}$ so that some component of

(IV.1.35)
$$B_{|z_1|/(\alpha C)}(z_1) \cap \Sigma_i^+$$

is not a disk. We can repeat the argument of Proposition IV.1.1 to get a second stable graph Γ'_j , separating curve $\tilde{\sigma}'_j$, and components Σ_j^{++} and Σ_j^{+-} of $\Sigma_j^+ \setminus \tilde{\sigma}'_j$ that are above and below, respectively, the graph Γ'_j ; see Figure 37.³⁰ Observe that the "middle component" Σ_j^{+-} is between the two stable graphs and has only two components in its interior boundary. If Σ_j^{+-} satisfies (2), then we are done. Otherwise, there is a third noncontractible simple closed curve. We can repeat the argument to cut Σ_j^{+-} to get an even lower component Σ_j^{+--} . The key point is that this new surface Σ_j^{+--} also has only two components in its interior boundary. Since Σ_j is compact, this process must eventually terminate to give a lowest component $\Sigma_j^{+-\cdots-}$ that satisfies (2).



Eventually get down to a lowest component; this must be scale-invariant ULSC (or we could cut again).

Figure 37. Cutting repeatedly to get the pair of pants decomposition.

Remark IV.1.36. The reader may find the constant α in Lemma IV.1.32 somewhat mysterious. The point is that taking α large forces the two interior boundary components in case (2) to be relatively far apart. This will be used to guarantee that the ULSC piece is sufficiently large, i.e., goes all the way out to the outer boundary in $\partial B_{R_i/(2C)}$.

IV.1.3. Step (2): The ULSC pieces of Σ_j contain graphs. We will next find the graphs in Σ_j converging to the plane $\{x_3 = 0\}$ away from 0 and possibly one other point. The argument for this is slightly simpler in Case (1) where $\tilde{\Sigma}_j^+$ is itself scale-invariant ULSC and we do not need to cut along

³⁰Proposition IV.1.1 directly gives the second stable graph disjoint from Σ_j . However, the proposition does not explicitly give that the two stable surfaces can be chosen to be disjoint. This is easy to achieve since the components of $B_{R_j} \setminus (\Sigma_j \cup \Gamma_j)$ are also mean convex in the sense of Meeks-Yau.

a second curve, but this simpler case already illustrates the key ideas. The argument follows a similar one in [CM03].

Suppose now that Case (1) in Lemma IV.1.32 holds for (a subsequence of) the Σ_j 's. The existence of the graphs in the $\tilde{\Sigma}_j^+$'s converging to $\{x_3 = 0\} \setminus \{0\}$ follows immediately from combining three facts:

• As $j \to \infty$, the minimum distance between $\partial B_1 \cap \tilde{\Sigma}_j^+$ and $\{x_3 = 0\}$ goes to zero. This was actually proven in Lemma 3.3 of [CM02c], which gave the existence of low points in a *connected* minimal surface contained on one side of a plane and with interior boundary close to this plane.³¹ We will recall this lemma from [CM02c] next.

LEMMA IV.1.37 (Lemma 3.3 in [CM02c]; see Figure 38). If $0 < \varepsilon < 4r_0/5$ and $\Sigma \subset B_{r_0}$ is a connected immersed minimal surface with $B_{\varepsilon} \cap \Sigma \neq \emptyset$, $\Sigma \setminus B_{\varepsilon} \neq \emptyset$, and

(IV.1.38)
$$\partial \Sigma \subset B_{\varepsilon} \cup (\partial B_{r_0} \cap \{x_3 > -3r_0/5\}),$$

then

(IV.1.39)
$$\min_{\Sigma \cap \{x_1^2 + x_2^2 \ge (4r_0/5)^2\}} x_3 \le 4\varepsilon \cosh^{-1}(3r_0/\varepsilon) < 4\varepsilon \log(6r_0/\varepsilon).$$

A minimal surface Σ above a plane and with $\partial \Sigma$ intersecting a small ball.



 Σ must contain points near the plane but far out.

Figure 38. The existence of low points near a plane.

In fact, we will use a slight variation of this that follows from the same proof. Namely, under the same assumptions, we get that

(IV.1.40)
$$\min_{\partial B_{r_0} \cap \partial \Sigma} x_3 \le 4\varepsilon \cosh^{-1}(3r_0/\varepsilon) < 4\varepsilon \log(6r_0/\varepsilon).$$

• The one-sided curvature estimate and the scale-invariant ULSC property give a scale-invariant curvature estimate for the $\tilde{\Sigma}_i^+$'s in a narrow cone about

³¹The argument for this was by contradiction. Namely, if there were no low points, then we would get a contradiction from the strong maximum principle by first sliding a catenoid up under the surface and then sliding the catenoid horizontally away, eventually separating two boundary components of the surface. Here the strong maximum principle is used to keep the sliding catenoids and the surface disjoint. See, for instance, Corollary 1.18 in [CM99] for a precise statement of the strong maximum principle.

the plane $\{x_3 = 0\}$. Here we have used that the $\tilde{\Sigma}_j^+$'s stay on one side of the graphs Γ_j converging to $\{x_3 = 0\} \setminus \{0\}$. Similarly, this curvature estimate and the barrier limit plane imply that the $\tilde{\Sigma}_j^+$'s are locally graphical in a slightly narrower cone about $\{x_3 = 0\}$.

• The first step gives a sequence of points in the $\tilde{\Sigma}_j^+$'s converging to a point in $\partial B_1 \cap \{x_3 = 0\}$. The second step allows us to apply the Harnack inequality to build this out into expanding, locally graphical, subsets of the $\tilde{\Sigma}_j^+$'s that are converging to the plane.

These locally graphical regions piece together to give graphs over expanding annuli; the other possibility would be to form a multi-valued graph, but this is impossible since such a multi-valued graph would be forced to spiral infinitely (since it cannot cross itself and also cannot cross the stable graph Γ_j).

Finally, we will briefly describe the modifications needed for case (2) in Lemma IV.1.32 when the Σ_j^{ulsc} 's have two interior boundary components. The complication arises in the second step. Namely, we can no longer locally extend the graph over $\{x_3 = 0\} \setminus \{0\}$; this is because $\tilde{\Sigma}_j^+$ is not scale-invariant ULSC in the second ball $B_{s_j}(y_j)$. To deal with this, we will consider several different cases.

The two simplest possibilities are when the points y_j go to either zero or infinity. When $y_j \to 0$, then we can replace the radii r_j by another sequence $r'_j > \max\{r_j, |y_j|\}$ where $r'_j \to 0$; with the new choice of r'_j , the new $\tilde{\Sigma}_j^+$'s are ULSC and we can proceed as in Case (1). On the other hand, when $|y_j| \to \infty$, we can replace the outer radii R_j by $|y_j|$ and the new sequence of $\tilde{\Sigma}_j^+$'s will again be scale-invariant ULSC.

Suppose therefore that the points y_j converge to a finite point $y \neq 0$. We will consider two separate subcases here (we can reduce to these after taking subsequences):

- Suppose first that s_j goes to 0. In this case, the one-sided curvature estimate gives estimates for the Σ^{ulsc}_j's as long as we stay away from the points 0 and y. We can then argue as in (1) to get the desired graphs these graphs converge to {x₃ = 0} \ {0, y}.
- Suppose now that $\liminf s_j = s_{\infty} > 0$. In this case, the sequence is ULSC away from 0 but not scale-invariant ULSC. (That is, the injectivity radius stays away from zero, but it does not necessarily grow as we go away from 0.) To make this precise, we will need an additional property of the balls $B_{s_j}(y_j)$ that was not recorded in Lemma IV.1.32 but follows easily from its proof.³² Namely, we can assume that s_j is the "smallest scale of nontrivial topology."

 $^{^{32}}$ See "Blowing up on the smallest scale of nontrivial topology" in the proof of Proposition IV.1.1.

More precisely, we can assume that the component of $B_{s_j}(y_j) \cap \Sigma_j$ containing the second interior boundary curve has injectivity radius at least βs_j for some fixed constant $\beta > 0$. In particular, since $\liminf s_j > 0$, the onesided curvature estimate gives uniform estimates on these components of $B_{s_j}(y_j) \cap \Sigma_j$. We can now argue as in (1) to get the desired graphs; this time the graphs converge to $\{x_3 = 0\} \setminus \{0\}$. In particular, this contradicts the existence of the second separating curves in $B_{s_j}(y_j)$ and, thus, shows that the case $\liminf s_j = s_{\infty} > 0$ does not occur.

This completes the proof of (C1) in Theorem 0.14.

IV.2. The ULSC regions of the lamination: (C2) and (D) in Theorem 0.14

In this section, we will prove that the ULSC regions of the lamination have the same structure as in the globally ULSC case of Theorem 0.9. Namely, we will prove that

- The leaves intersecting the ULSC part of the singular set S_{ulsc} are parallel planes. Each plane intersects S_{ulsc} at two points.
- S_{ulsc} is a union of Lipschitz curves transverse to the leaves. The leaves intersecting S_{ulsc} foliate an open subset of \mathbf{R}^3 that does not intersect S_{neck} .

The key for the proof of these two properties will be to show that each collapsed leaf has a neighborhood that is ULSC; this will be done in Proposition IV.2.2. Recall that a leaf Γ of \mathcal{L}' is said to be collapsed if its closure Γ_{Clos} contains a point in $\mathcal{S}_{\text{ulsc}}$ and this point is a removable singularity for Γ ; see Definition II.2.9. We have already established a great deal of structure for collapsed leaves in Proposition II.3.1, and much of this will be used below.

Here, and elsewhere in this section, the closure Γ_{Clos} of a leaf Γ is defined to be the union of the closures of all bounded geodesic balls in Γ ; that is, we fix a point $x_{\Gamma} \in \Gamma$ and set

(IV.2.1)
$$\Gamma_{\text{Clos}} = \bigcup_{r} \overline{\mathcal{B}_r(x_{\Gamma})},$$

where $\overline{\mathcal{B}_r(x_{\Gamma})}$ is the closure of $\mathcal{B}_r(x_{\Gamma})$ as a subset of \mathbf{R}^3 . Eventually we will show that, for a collapsed leaf, Γ_{Clos} is a flat plane and hence, in particular, $\Gamma_{\text{Clos}} = \overline{\Gamma}$. However, a priori Γ may not be proper, and thus the two notions could a priori differ.

The main result of this section is Proposition IV.2.2 below showing that Γ_{Clos} does not intersect S_{neck} . Since S_{neck} is a closed subset of \mathbf{R}^3 , it follows that every compact subset of Γ_{Clos} has a neighborhood in \mathbf{R}^3 that does not intersect S_{neck} .

PROPOSITION IV.2.2. If Γ is a collapsed leaf of \mathcal{L}' , then $\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}} = \emptyset$.

The rough idea of the proof is to first show that $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of exactly two points (this is analogous to each leaf having at most two singular points in the ULSC case); see Corollary IV.2.6 below. Consequently, the union of Γ and the given point in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ will give a stable surface that is "complete away from a point" and, hence, flat by Lemma A.26. Finally, once we know that Γ is flat, it will be easy to check that $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$.

Before we can get into the proof just outlined, we will need to recall a little of the structure that has already been proven. We will do this in the next two subsections. The next subsection establishes a key property of the stable limit planes that we get through each point of S_{neck} . The second subsection below reviews the properties of a general collapsed leaf of \mathcal{L}' .

IV.2.1. The leaf Γ cannot cross the limit planes. The structure result (C_{neck}) from Theorem 0.12 gives graphs $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$ in Σ_j that converge to a plane P_z through each point $z \in S_{\text{neck}}$; see Figure 39. The graphs $\tilde{\Sigma}_j^+$ converge smoothly away from z and possibly one other point (call it z^+); this second point must also be in S_{neck} . Similarly, the $\tilde{\Sigma}_j^-$ converge away from z and possibly a point $z^- \in S_{\text{neck}}$. Furthermore, $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$ are separated in Σ_j by the curve σ_j .³³ One expects that the limit plane P_z should be the closure of a leaf of \mathcal{L}' , but this is not a priori clear; for instance, S might even be dense in P_z .



Figure 39. The structure near z: The two graphs $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$ are separated by curves σ_j shrinking to z. The multi-valued graph Σ_j^g comes near z.

Using this structure, the next lemma proves that the leaves of \mathcal{L}' do not cross any of these limit planes. This is almost obvious since the stable surfaces

80

³³Property (C_{neck}) from Theorem 0.12 holds at z by (C1) in Theorem 0.14; this was proven in Section IV.1. The last "separation" claim is not explicit in (C_{neck}) but follows immediately from Proposition IV.1.1; using the notation from the fact that proposition, we have that $\tilde{\Sigma}_j^+ \subset \Sigma_j^+$ and $\tilde{\Sigma}_j^- \subset \Sigma_j^-$.

converging to the limit plane are disjoint from the Σ_j 's that are converging to the leaves of \mathcal{L}' . The only possible difficulty comes from that the convergence is only away from the singular set \mathcal{S} , but this will be easy to handle. The lemma applies to an arbitrary leaf Γ of \mathcal{L}' ; i.e., we do not need Γ to be collapsed.

LEMMA IV.2.3. Suppose that Γ is an arbitrary leaf of \mathcal{L}' , collapsed or not. If z is any point in S_{neck} and P_z is the corresponding limit plane through z, then Γ does not cross P_z .

Proof. Fix an open connected set $K \subset \Gamma$ with compact closure in Γ , and recall that the Σ_j 's contain

- Graphs $\tilde{\Sigma}_{j}^{+}$ and $\tilde{\Sigma}_{j}^{-}$ that both converge to P_{z} away from a finite set \mathcal{P} of points; see Figure 39. Moreover, $\tilde{\Sigma}_{j}^{+}$ and $\tilde{\Sigma}_{j}^{-}$ are separated in Σ_{j} by curves σ_{j} shrinking to z.
- Connected subsets $\Sigma_j^g \subset \Sigma_j$ given by Lemma II.3.15 that are locally graphical over K and that converge with multiplicity to K. These locally graphical subsets might globally be graphs or multi-valued graphs over K.

We will show first that the Σ_j^g 's cannot intersect both $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$. First, using that z is not in Γ and $K \subset \Gamma$ has compact closure, we can fix a ball $B_s(z)$ so that

(IV.2.4)
$$B_s(z) \cap \Sigma_i^g = \emptyset$$

for all j sufficiently large. On the other hand, the curves σ_j separating Σ_j^+ and $\tilde{\Sigma}_j^-$ are shrinking to z. Therefore, the curves σ_j do not intersect Σ_j^g when j is large and, hence, the connected set Σ_j^g cannot intersect both $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$. Without loss of generality, we can assume that

(IV.2.5)
$$\Sigma_j^g \cap \tilde{\Sigma}_j^+ = \emptyset$$

We will next use (IV.2.5) to show that the two smooth open surfaces P_z and K do not have any points of transverse intersection. Namely, if P_z and K have points of transverse intersection, then, since the singular set \mathcal{P} for the convergence to P_z is finite, $P_z \setminus \mathcal{P}$ and K would also have points of transverse intersection. However, this contradicts (IV.2.5) since $\tilde{\Sigma}_j^+ \to P_z$ smoothly away from \mathcal{P} and $\Sigma_j^g \to K$.

Finally, recall that if a connected minimal surface intersects both sides of a plane, then the surface and plane must have a point of transverse intersection; this follows from the local structure of the nodal set of a harmonic function — see, e.g., Lemma 4.28 in [CM99]. Therefore, since K is connected and does not intersect P_z transversely at any point, we see that K must be on one side of P_z . Since this holds for every such K and these exhaust Γ by Lemma II.3.25, we see that Γ also lies on one side of P_z .

Throughout the rest of this section, Γ will be a collapsed leaf of \mathcal{L}' .

IV.2.2. The properties of a collapsed leaf Γ . Before getting into the proof, it may be useful to recall the properties of the leaf Γ . Eventually, we will use these properties to show that Γ_{Clos} is a plane.

- Γ is by definition an injective immersion of a *connected* surface without boundary, but *not* necessarily complete. Furthermore, the immersion is not necessarily proper.
- Since Γ is a leaf of \mathcal{L}' , it follows that Γ does not intersect \mathcal{S} and hence, since \mathcal{S} is closed (as a subset of \mathbf{R}^3), each point in Γ has a neighborhood where the curvatures of the Σ_j 's are uniformly bounded.
- The following local structure of Γ near a point of Γ_{Clos}∩S_{ulsc} was established in (1) in Proposition II.3.1:
- (Loc) Given any $y \in \Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$, there exists $r_0 > 0$ so that the closure (in \mathbb{R}^3) of each component of $B_{r_0}(y) \cap \Gamma$ is a compact embedded disk with boundary in $\partial B_{r_0}(y)$. Furthermore, $B_{r_0}(y) \cap \Gamma$ must contain the component Γ_y given by Lemma II.2.3, and Γ_y is the only component of $B_{r_0}(y) \cap \Gamma$ with y in its closure.
- Γ (or its oriented double cover) must be stable by (2) in Proposition II.3.1.
- Γ_{Clos} intersects S_{ulsc} in at most two points by (3) in Proposition II.3.1.

These properties will be essential for proving Proposition IV.2.2. The main difficulty will be that Γ is not complete. This occurs where Γ_{Clos} intersects S; see Figure 40 for such an example. By (Loc), the points in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ are isolated removable singularities of Γ and thus are easily dealt with. Consequently, the first step will be to control the number of points of $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$ when Γ is collapsed. This will be done in the next subsection.



Each point in $\partial \Gamma$ is in \mathcal{S} .

Figure 40. A priori Γ could be a punctured disk in a plane.

IV.2.3. $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of exactly two points. The next corollary is the first step needed for the proof of Proposition IV.2.2, which gives $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$. This corollary shows that if $\Gamma_{\text{Clos}} \cap S_{\text{neck}} \neq \emptyset$, then $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of exactly two points.

82

COROLLARY IV.2.6. Let Γ be a collapsed leaf of \mathcal{L}' . If $\Gamma_{\text{Clos}} \cap S_{\text{neck}} \neq \emptyset$, then $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of exactly two points with one each in S_{ulsc} and S_{neck} .

This will be an easy corollary of Lemma IV.2.7 below, which shows that the sheets of the multi-valued graphs over Γ connect in a small neighborhood of any singular point. Previously, we used the one-sided curvature estimate to establish a similar connecting property near a point of S_{ulsc} .

Before making this connecting property precise, we need to set up some notation. Recall that if K is a "sufficiently large" open connected subdomain of Γ with compact closure in Γ , then Corollary II.3.18 gives a sequence of multivalued graphs $\Sigma_j^g \subset \Sigma_j$ that converges to K with infinite multiplicity. Since Γ can be exhausted by a nested sequence K_j of such K's by Lemma II.3.25, we can assume that the following holds (after passing to a subsequence):

(Graph) Σ_j contains a *j*-valued graph Σ_j^g over K_j of a function whose values are bounded by 1/j and whose gradient is bounded by 1/j. Here $K_j \subset \Gamma$ is a nested sequence of connected open sets with compact closure in Γ with $\Gamma = \bigcup_j K_j$.

We actually know a good deal more about these multi-valued graphs, but this additional structure will not be needed until the proof of Corollary IV.2.6.

LEMMA IV.2.7. The sheets of the multi-valued graph Σ_j^g connect near $z \in \Gamma_{\text{Clos}} \cap S_{\text{neck}}$; see Figure 41. Precisely, given any r > 0, there exist $\delta > 0$ and J so that if $x \in B_{\delta}(z) \cap \Gamma$, j > J, and

(IV.2.8) z_j^+ and z_j^- are points in the multi-valued graph Σ_j^g over x, then z_i^+ and z_j^- are in the same connected component of $B_r(z) \cap \Sigma_j$.

Proof. We will argue by contradiction, so suppose there exists some r > 0so that for every $\delta > 0$, there exists $x \in B_{\delta}(z) \cap \Gamma$ and infinitely j's so that $B_r(z) \cap \Sigma_j$ has (at least) two distinct components that both contain points



The sheets all intersect the same component of $B_{\delta}(z) \cap \Sigma$ for j sufficiently large.

Figure 41. Lemma IV.2.7: The sheets of the multi-valued graph Σ_j^g must connect near $z \in \Gamma_{\text{Clos}} \cap S_{\text{neck}}$.

in Σ_j^g over x. After passing to a subsequence, we can assume that there is a sequence of points z_i^+ and z_i^- in Σ_j^g with

(IV.2.9)
$$z_i^+$$
 and z_j^- converging to z ,

so that z_j^+ and z_j^- are in distinct components of $B_r(z) \cap \Sigma_j$.

We will use the Σ_j 's as barriers for a Plateau problem to construct stable surfaces $\tilde{\Gamma}_j$ between these distinct components. Before constructing the stable surfaces $\tilde{\Gamma}_j$, recall the following useful consequence of the interior curvature estimates for stable surfaces of [Sch83a], [CM02a] (see, e.g., Lemma 2.2 in [CM99]):

(Stab) There exists a positive constant $\alpha < 1$ so that if Γ_s is a stable embedded minimal surface with $\partial \Gamma_s \subset \partial B_R$, then each component of $B_{\alpha R} \cap \Gamma_s$ is a graph over some plane with gradient bounded by one.

Set

(IV.2.10)
$$r' = \alpha r.$$

The properties of the stable graphs Γ_j . Below we will find (stable) graphs $\tilde{\Gamma}_j$ between these distinct components so that the following three properties hold:

(IV.2.11) $\tilde{\Gamma}_j \subset B_{r'}(z) \setminus \Sigma_j \text{ with } \partial \tilde{\Gamma}_j \subset \partial B_{r'}(z).$

(IV.2.12)
$$B_{\varepsilon_i}(z) \cap \tilde{\Gamma}_j \neq \emptyset$$
 where $\varepsilon_j \to 0$.

(IV.2.13) The multi-valued graph Σ_{i}^{g} in Σ_{j} intersects both sides of $\tilde{\Gamma}_{j}$.

Recall that a properly embedded (connected) surface in \mathbb{R}^3 will automatically have two sides. Properties (IV.2.12) and (IV.2.13) will follow from a standard linking argument.

Constructing the stable graphs $\tilde{\Gamma}_j$. We will construct $\tilde{\Gamma}_j$ in two steps, first finding stable surfaces $\tilde{\Gamma}_j^0$ in the larger ball $B_r(z)$ and then letting $\tilde{\Gamma}_j$ be an appropriate component of $B_{r'}(z) \cap \tilde{\Gamma}_j^0$ where r' is given by (IV.2.10).

To construct $\tilde{\Gamma}_{j}^{0}$, first let $\Sigma_{r,j}^{+}$ and $\Sigma_{r,j}^{-}$ be the distinct components of $B_{r}(z) \cap \Sigma_{j}$ that contain z_{j}^{+} and z_{j}^{-} , respectively. Let ℓ_{j} be a line segment connecting z_{j}^{+} and z_{j}^{-} ; see Figure 42. Fix a component ℓ_{j}^{0} of $\ell_{j} \setminus (\Sigma_{r,j}^{+} \cup \Sigma_{r,j}^{-})$ that also connects $\Sigma_{r,j}^{+}$ and $\Sigma_{r,j}^{-}$ but intersects $\Sigma_{r,j}^{+}$ and $\Sigma_{r,j}^{-}$ only at its endpoints $\partial \ell_{j}^{0}$. The surface $\Sigma_{r,j}^{+}$ sits in the boundary of two components of $B_{r}(z) \setminus \Sigma_{j}$ — one component on each side of $\Sigma_{r,j}^{+}$.³⁴ Let Ω_{j} be the component that ℓ_{j}^{0} points

³⁴In the simplest case where $\Sigma_{r,j}^+$ and $\Sigma_{r,j}^-$ are the only components of $B_r(z) \cap \Sigma_j$, we would choose Ω_j to be the component of $B_r(z) \setminus \Sigma_j$ between them, i.e., the component containing



Bold ℓ_j connects the two components $\Sigma_{r,j}^+$ and $\Sigma_{r,j}^-$ of $B_r(z) \cap \Sigma_j$

Figure 42. We argue by contradiction to show that the sheets connect near z. Assuming that they do not, we first construct stable surfaces $\tilde{\Gamma}_{i}$.

into as it leaves $\Sigma_{r,j}^+$; i.e., let Ω_j be the component of $B_r(z) \setminus \Sigma_j$ with (IV.2.14) $\Sigma_{r,j}^+ \subset \partial \Omega_j$ and $\Omega_j \cap (\ell_j^0 \setminus \partial \ell_j^0) \neq \emptyset$.

The point of choosing Ω_j in this way is that the curve ℓ_j^0 has linking number one with $\partial \Sigma_{r,j}^+$ in Ω_j . The domain Ω_j is mean convex and, hence, [MY82a]– [MY82b] give a stable embedded minimal surface $\tilde{\Gamma}_j^0 \subset \Omega_j$ with $\partial \tilde{\Gamma}_j^0 = \partial \Sigma_{r,j}^+$.

Since $\partial \tilde{\Gamma}_j^0$ has linking number one with ℓ_j^0 in Ω_j , the endpoints of ℓ_j^0 are separated in $B_r(z)$ by $\tilde{\Gamma}_j^0$. However, the endpoints of ℓ_j^0 connect to the endpoints of ℓ_j by curves in $\Sigma_{r,j}^+$ and $\Sigma_{r,j}^-$; these curves do not cross $\tilde{\Gamma}_j^0$ and, consequently, the endpoints of ℓ_j are also separated in $B_r(z)$ by $\tilde{\Gamma}_j^0$. We can therefore choose a component $\tilde{\Gamma}_j$ of $B_{r'}(z) \cap \tilde{\Gamma}_j$ that separates the endpoints of ℓ_j ; so

(IV.2.15)
$$\Gamma_j \cap \ell_j \neq \emptyset$$

Since z_j^+ and z_j^- go to z, it follows that each ℓ_j is contained in a ball $B_{\varepsilon_j}(z)$ where $\varepsilon_j \to 0$. In particular, (IV.2.15) gives (IV.2.12). Since the endpoints of ℓ_j are both in the multi-valued graph, we also get (IV.2.13).

Using the stable graphs Γ_j to show that $\Gamma \subset P_z$. Now that we have constructed the $\tilde{\Gamma}_j$'s, we are ready to return to the proof of the lemma. The first step will be to show that a subsequence of the $\tilde{\Gamma}_j$'s converges to a subset of P_z . First, by (Stab), the surface $\tilde{\Gamma}_j$ is a graph with gradient bounded by one. After passing to a subsequence, we can therefore assume that $\tilde{\Gamma}_j$ converges to a minimal graph $\tilde{\Gamma}$. Since $\varepsilon_j \to 0$, $\tilde{\Gamma}$ contains z. On the other hand, the graph $\tilde{\Sigma}_i^+ \subset \Sigma_j$ given by Theorem 0.12 (see Figure 39) does not intersect the graph

the interior of ℓ_j^0 . In general, there are other components of $B_r(z) \cap \Sigma_j$ intersecting ℓ_j^0 so we cannot do this. This slightly complicates the choice of Ω_j .

 $\tilde{\Gamma}_j$; since $\tilde{\Sigma}_j^+$ converges to P_z away from a finite set, we conclude that $\tilde{\Gamma}$ must be on one side of P_z . Since $\tilde{\Gamma}$ is on one side of P_z and intersects it at z, the strong maximum principle implies that $\tilde{\Gamma}$ is contained in P_z , as desired.

We will next show that

(IV.2.16)
$$\Gamma \subset P_z$$

First, by (IV.2.13), there is at least one sheet of the multi-valued graph Σ_j^g on each side of the graph $\tilde{\Gamma}_j$ for every j. Since both of these sheets converge to Γ by (Graph), we can fix a point \tilde{y} in $B_{r'}(z) \cap \Gamma$ that is both a limit of points $\tilde{y}_j^+ \in \Sigma_j^g$ above $\tilde{\Gamma}_j$ and a limit of points $\tilde{y}_j^- \in \Sigma_j^g$ below $\tilde{\Gamma}_j$. For each j, the line segment connecting \tilde{y}_j^+ to \tilde{y}_j^- must intersect $\tilde{\Gamma}_j$ at a point \tilde{y}_j . The sequence of points $\tilde{y}_j \in \tilde{\Gamma}_j$ must also converge to the common limit \tilde{y} of \tilde{y}_j^+ and \tilde{y}_j^- . In particular, we conclude that $\tilde{y} \in P_z \cap \Gamma$, so that $\Gamma \subset P_z$ by Lemma IV.2.3 and the strong maximum principle.

The contradiction: We cannot have $\Gamma \subset P_z$. To complete the proof of the lemma, we explain next how (IV.2.16) leads to a contradiction. Since we will need the same argument later, it will be useful to isolate it out as a claim:

CLAIM. (IV.2.16) cannot hold; i.e., we cannot have $\Gamma \subset P_z$.

Proof of Claim. We will argue by contradiction, so suppose that $\Gamma \subset P_z$. Since Γ_{Clos} is contained in the closure of Γ , it follows that $\Gamma_{\text{Clos}} \subset P_z$ as well. Since Γ is collapsed, Γ_{Clos} contains a point

(IV.2.17)
$$y_{\Gamma} \subset S_{\text{ulsc}} \cap \Gamma_{\text{Clos}} \subset P_z.$$

Both of the graphs $\tilde{\Sigma}_j^+$ and $\tilde{\Sigma}_j^-$ are converging to P_z away from punctures in $\mathcal{S}_{\text{neck}}$, so we get sequences of points $y_i^+ \in \tilde{\Sigma}_i^+$ and $y_i^- \in \tilde{\Sigma}_i^-$ with

(IV.2.18)
$$y_j^+ \to y_\Gamma \text{ and } y_j^- \to y_\Gamma.$$

We will next use the one-sided curvature estimate to prove that y_j^+ and y_j^- can be connected in Σ_j in any small neighborhood of y_{Γ} as $j \to \infty$. To see this, note first that y_{Γ} is in S_{ulsc} and hence each component of $B_r(y_{\Gamma}) \cap \Sigma_j$ is a disk for some r > 0; after possibly choosing r smaller, we can assume that $|y_{\Gamma} - z| > r$. If there were at least two of these disks in $B_r(y_{\Gamma}) \cap \Sigma_j$ intersecting the concentric sub-ball $B_{C''r}(y_{\Gamma})$ where C'' > 0 is a sufficiently small constant, then the one-sided curvature estimate would give a uniform curvature bound for each component of Σ_j in this sub-ball.³⁵ Since no such curvature bound

 $^{^{35}}$ We actually use a corollary of the one-sided curvature estimate recorded in Corollary 0.4 in [CM04e]. This corollary states that if there are two disjoint embedded minimal surfaces in a ball in \mathbb{R}^3 , both intersect a sufficiently small ball around the center, and one is a disk, then we get an interior curvature estimate for the disk-type component.

holds near a point of S by definition, we conclude that the points y_j^+ and $y_j^$ must be in the same connected component of $B_r(y_{\Gamma}) \cap \Sigma_j$ for all j sufficiently large.

This local connectedness near y_{Γ} will easily lead to a contradiction. This is because y_j^+ and y_j^- were proven to be separated in Σ_j by the curve σ_j and the σ_j 's are shrinking to the point $z \neq y_{\Gamma}$. This contradiction completes the proof of the claim and hence of the lemma.





Figure 43. The points \tilde{y}_j^+ and \tilde{y}_j^- converge to a point $\tilde{y} \in \Gamma$ from opposite sides of the graph $\tilde{\Gamma}_j$. Thus, the points $\tilde{y}_j \in \tilde{\Gamma}_j$ between them also converge to $\tilde{y} \in \Gamma$, giving Lemma IV.2.7.

Proof of Corollary IV.2.6. By assumption, $\Gamma_{\text{Clos}} \cap S$ contains at least one point in each of $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$ and $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$. We will argue by contradiction to prove that $\Gamma_{\text{Clos}} \cap S$ cannot contain a third point. For simplicity, we will assume that the Σ_j 's are planar domains; the general finite genus case follows with easy modifications.

The proof follows the argument given in Remark II.3.32 and the basic idea is simple:

Suppose that p, q, and r are distinct points in $\Gamma_{\text{Clos}} \cap S$. The local connecting property near each point of S allows us to construct closed noncontractible curves in the Σ_j 's that converge with multiplicity two to a curve in Γ connecting p and q. These curves must separate in the planar domain Σ_j . However, if we connect points on opposite sides of these curves to the third point r, the local connecting property near r gives a contradiction.

The only difficulty in carrying out this argument will be that the surface Γ is not complete.

Step 1: Choosing the singular points and curves in Γ . We will first choose the points p, q, and r in $\Gamma_{\text{Clos}} \cap S$. Let p be the given point in $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}}$, and then let q be a closest point in $\Gamma_{\text{Clos}} \cap S$ to p. (A priori there may be many possible choices.) Since we are arguing by contradiction, there is a third distinct point $r \in \Gamma_{\text{Clos}} \cap S$.

By our choice of q, there must be a minimizing geodesic $\gamma_{pq} : [0, L] \to \Gamma_{\text{Clos}}$ parametrized by arclength and with the following properties:

- $\gamma_{pq}(0) = p$ and $\gamma_{pq}(L) = q$.
- The interior of γ_{pq} is contained in Γ .

Since the closed geodesic γ_{pq} is compact, we must have

(IV.2.19)
$$\operatorname{dist}_{\mathbf{R}^3}(\gamma_{pq}, r) > 0$$

Since $p \in S_{\text{ulsc}}$ and Γ is collapsed, property (1) in Proposition II.3.1 gives a ball $B_{\delta}(p)$ and a component Γ_p of $B_{\delta}(p) \cap \Gamma$ so that $\Gamma_p \cup \{p\}$ is a smooth minimal graph. Since γ_{pq} is minimizing, it is not hard to see that $\partial \Gamma_p$ intersects γ_{pq} in a single point p'.

Fix a constant $\varepsilon > 0$ that is much smaller than the distance from r to γ_{pq} . By the definition of Γ_{Clos} , there must be a point r' in Γ that is distance ε from r. Since Γ is connected, we can choose a compact curve $\tilde{\gamma}_{p'r'}$ that is contained in Γ and connects p' to r'. The curve $\tilde{\gamma}_{p'r'}$ may intersect γ_{pq} many times, so we replace it with the component of $\tilde{\gamma}_{p'r'} \setminus \gamma_{pq}$ with r' in its boundary. This gives a curve in Γ from γ_{pq} to r' and whose interior does not intersect γ_{pq} . After adding a subsegment of γ_{pq} and perturbing the resulting curve slightly off of γ_{pq} , we get a compact curve $\gamma_{p'r'} \subset \Gamma$ from p' to r' and whose interior does not intersect γ_{pq} ; see Figure 44.

The point about the curve $\gamma_{p'r'}$ is that it will give a way to connect points near p' to r' in $\Gamma \setminus \gamma_{pq}$; see Figure 44. This will be especially useful since the curve $\partial \Gamma_p$ allows us to connect points near p' that are on the opposite sides of γ_{pq} .

Step 2: Choosing the curves in the Σ_j 's. We can now argue as in Remark II.3.32. The key point is that Theorem I.1.6 and Lemma IV.2.7 imply that the Σ_j 's are locally connected in a small neighborhood of any of the singular points in $\Gamma_{\text{Clos}} \cap \mathcal{S}$. These connecting properties allow us to find simple closed curves $\gamma_{pq}^j \subset \Sigma_j$ with the following properties:

- γ_{pq}^{j} is contained in the ε -tubular neighborhood of γ_{pq} .
- $\gamma_{pq}^{j} \setminus (B_{\varepsilon}(p) \cup B_{\varepsilon}(q))$ consists of two graphs over γ_{pq} that are in distinct sheets of Σ_{j} .

Step 3: The contradiction. Since Σ_j has genus zero, the curve γ_{pq}^j must separate Σ_j into two distinct components. However, it is easy to see that this is impossible by using the local connecting property near the third point r.





Figure 44. The curves γ_{pq} and $\gamma_{p'r'}$.

Namely, we can take two points near p on opposite sides of γ_{pq}^{j} and connect each of them to r' by curves in Σ_{j} that do not intersect γ_{pq}^{j} . One of these connecting curves will be a graph over $\gamma_{p'r'}$ while the other is a graph over $\partial \Gamma_{p} \cup \gamma_{p'r'}$. These two connecting curves can then be connected to each other in $B_{C\varepsilon}(r) \cap \Sigma_{j}$, giving the desired contradiction.

IV.2.4. The proof of Proposition IV.2.2. We can now use Corollary IV.2.6 and the properties of a collapsed leaf to prove Proposition IV.2.2. Recall that this proposition claims that $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$ whenever Γ is a collapsed leaf of \mathcal{L}' .

Proof of Proposition IV.2.2. We will argue by contradiction, so suppose that Γ is a collapsed leaf of \mathcal{L}' and $\Gamma_{\text{Clos}} \cap S_{\text{neck}} \neq \emptyset$. By Corollary IV.2.6, we know that $\Gamma_{\text{Clos}} \setminus \Gamma$ consists of one point y in S_{ulsc} and one point z in S_{neck} . Furthermore, property (1) in Proposition II.3.1 implies that the point y is a removable singularity for Γ so that $\Gamma \cup \{y\}$ is smooth and complete away from the point z.

We will show next that Γ is flat. The starting point for this is that Γ (or its oriented double cover) is stable by property (2) in Proposition II.3.1. A standard "logarithmic cutoff function" argument then implies that $\Gamma \cup \{y\}$ (or its oriented double cover) is also stable; we leave the simple argument to the reader. If $\Gamma \cup \{y\}$ had been complete, then the Bernstein theorem for stable surfaces would have implied that it was flat. However, even in this case where $\Gamma \cup \{y\}$ is complete away from the single point z, Lemma A.26 in Appendix A implies that Γ is flat.

We have now established that Γ is a plane with two points removed and one of these points (namely z) is in S_{neck} . Since Γ cannot cross the limit plane P_z through z by Lemma IV.2.3, it follows that

(IV.2.20)
$$\Gamma \subset P_z$$
.

However, we already saw in the claim at the end of the proof of Lemma IV.2.7 that (IV.2.20) is impossible. This contradiction completes the proof of the proposition. $\hfill \Box$

IV.2.5. The proof of (C2) and (D) in Theorem 0.14. We can now argue as in the ULSC case of Part III to prove (C2) and (D) in Theorem 0.14. As we saw in Part III, the argument requires that we establish the following four properties of an arbitrary collapsed leaf Γ of \mathcal{L}' :

- (0) $\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}} = \emptyset.$
- (1) Γ_{Clos} is a plane and $\Gamma_{\text{Clos}} \setminus \Gamma$ contains at most two points.
- (2) $\Gamma_{\text{Clos}} \setminus \Gamma$ contains exactly two points.
- (*) If $t \in x_3(\mathcal{S}_{ulsc})$ and $\varepsilon > 0$, then

(IV.2.21)
$$\mathcal{S}_{\text{ulsc}} \cap \{t < x_3 < t + \varepsilon\} \neq \emptyset \text{ and } \mathcal{S}_{\text{ulsc}} \cap \{t - \varepsilon < x_3 < t\} \neq \emptyset.$$

Once we show that (0), (1), (2), and (\star) hold, then (C2), and (D) in Theorem 0.14 follow exactly as in Part III; we will not repeat the argument here.

It suffices therefore to check that (0), (1), (2), and (\star) hold in this setting. The first two are quite easy: (0) is exactly the conclusion of Proposition IV.2.2 and (1) follows from (0) together with (1) in Proposition III.0.2.

We will prove (2) by contradiction, so suppose that $\Gamma_{\text{Clos}} \cap S_{\text{ulsc}} = \{0\}$. It follows immediately that the Σ_j 's contain multi-valued graphs over (subsets of) the punctured plane $\Gamma = \{x_3 = 0\} \setminus \{0\}$ that converge to Γ with infinite multiplicity. Moreover, (D) in the proof of property (2) in Proposition III.0.2 gives the following *scale invariant* ULSC property:

(D) There exists $\tau > 0$ so that, for $z \in \{x_3 = 0\}$ and j large, each component of $B_{\tau |z|}(z) \cap \Sigma_j$ that connects to the multi-valued graph in Σ_j is a disk.

Note that the proof of (D) did not use that the sequence was ULSC. It follows from (D) and the one-sided curvature estimate that the multi-valued graphs in Σ_j converging to Γ spiral through an entire cone about Γ . Note that the Σ_j 's are assumed to be uniformly nonsimply connected. Therefore, Proposition IV.1.1 gives stable graphs Γ_j that are disjoint from Σ_j . The Γ_j 's are graphs with bounded gradient that start out in a fixed ball and are defined over annuli with a fixed inner radius and with outer radii going to infinity. Standard results for exterior graphs then imply that the Γ_j 's grow sublinearly. Consequently, the Γ_j 's are eventually contained in the narrow cone that the multi-valued graphs in the Σ_j 's spiral through. However, this is impossible since the two are disjoint. This contradiction completes the sketch of the proof of (2); we leave the details to the reader.

Remark IV.2.22. The argument that we gave here is actually simpler than in the ULSC case; cf. (2) in Proposition III.0.2. However, we could not yet use

90

this argument for the ULSC case since Proposition IV.1.1 relies on the ULSC case.

Finally, (\star) follows from (1) and (2) together with Lemma III.1.4, which proved (\star) in the ULSC case. However, Lemma III.1.4 did not actually require the sequence to be ULSC, but rather required only that (1) and (2) above hold. This completes the sketch of (C2) and (D) in Theorem 0.14.

IV.3. Putting it all together: The proof of Theorem 0.14

We have now completed the proof of all six of the claims in Theorem 0.14 over the course of this part. (The six claims are (A), (B), (C1), (C2), (D), and (P) in Theorem 0.14.) For the reader's convenience, we will review next where each was proven.

Proof of Theorem 0.14. The singular set S is defined in Definition/ Lemma II.1.1, where we also prove property (B). Lemma II.1.2 gives a subsequence Σ_j that converges to a minimal lamination \mathcal{L}' of $\mathbf{R}^3 \setminus S$, thus giving (A). Property (C1), which describes the points in S_{neck} , is established in Section IV.1. Properties (C2) and (D), which describe the points in S_{ulsc} , are established in Section IV.2. Finally, Property (P), which shows that the leaves of \mathcal{L}' do not cross the limit planes given by (C1), is proven in Lemma IV.2.3. \Box

Part V. The no mixing theorem, Theorem 0.4

This part is devoted to the proof of the no mixing theorem, i.e., Theorem 0.4. Recall that this theorem asserts that the singular set S consists of either exclusively helicoid points or exclusively catenoid points, i.e., either $S_{\text{neck}} = \emptyset$ or $S_{\text{ulsc}} = \emptyset$. We have already shown in (D) of Theorem 0.14 that the leaves intersecting S_{ulsc} foliate an open subset of \mathbf{R}^3 that does not intersect S_{neck} . Using in part that S is closed, we will show that the closure of this foliated region will also intersect S; therefore, the boundary of the foliated region must intersect $S_{\text{neck}} = S \setminus S_{\text{ulsc}}$. We will prove the no mixing theorem by showing that also the closure of this foliated region does not intersect S_{neck} and, hence, the foliated region is either empty or all of \mathbf{R}^3 . The argument for this will be very similar to the one that we used earlier to show that the leaves intersecting S_{ulsc} do not intersect S_{neck} .

Proof of Theorem 0.4. Suppose that $S_{ulsc} \neq \emptyset$; we will show that $S_{neck} = \emptyset$. By (D) of Theorem 0.14, the set S_{ulsc} is a union of Lipschitz curves transverse to the leaves of the lamination and the leaves intersecting S_{ulsc} foliate an open subset of \mathbf{R}^3 that does not intersect S_{neck} ; we will call this foliated region the "ULSC region." Moreover, each connected component of the ULSC region contains exactly two curves in S_{ulsc} and each of these two curves intersects



Figure 45. A limit lamination where the ULSC region has nonempty boundary; we will rule this out.



Figure 46. Properties of a sequence Σ_j that converges to a lamination where the ULSC region has nonempty boundary.

each leaf exactly once. We will prove the theorem by showing that these ULSC leaves foliate all of \mathbf{R}^3 . It is easy to see that this is equivalent to showing that a curve in \mathcal{S}_{ulsc} cannot just stop.

Note first that each curve in S_{ulsc} is in fact a line segment *orthogonal* to the leaves of the lamination. This follows from the main theorem of [Mee04] since we have already proven here that the ULSC regions are foliated.

Suppose now that one of the line segments in S_{ulsc} does stop, i.e., has an endpoint z. Since the set S is closed and S_{ulsc} is open in S, the endpoint must be in S_{neck} . To complete the proof, we will use the properties of the sequence Σ_j to show that having z in S_{neck} leads to a contradiction. This contradiction follows from the following four steps:

(1) Since $z \in S_{\text{neck}}$, Proposition IV.1.1 gives a sequence of stable graphs $\Gamma_j \subset B_{R_j} \setminus \Sigma_j$ converging to a horizontal plane through z; this convergence is smooth away from the point z. Proposition IV.1.1 also gives separating curves $\gamma_j \subset \Sigma_j$ with $\gamma_j \to z$ where γ_j divides Σ_j into a component Σ_j^+ above Γ_j and a component Σ_j^- below.³⁶ After possibly reflecting about

³⁶When we say that " Σ_j^+ is above Γ_j ," we have to be a little bit careful since each Γ_j is defined only over an annulus. The precise statement is given in Proposition IV.1.1: There are shrinking balls $B_{r_j}(z)$ with $r_j \to 0$ so that $\Gamma_j \cup B_{r_j}(z)$ divides $B_{R_j/C}$ into components H_j^+ above Γ_j and H_j^- below Γ_j ; we then have that Σ_j^+ is contained in $H_j^+ \cup B_{r_j}(z)$.

the horizontal plane through z, we can assume that the segment in S_{ulsc} lies above the plane.

- (2) Since z is in the closure of S_{ulsc} , (C2) of Theorem 0.14 gives double spiral staircases in Σ_j^+ above the horizontal plane through z. More precisely, fix a ball $B_{2s}(z)$ that intersects only one of the two vertical line segments in S_{ulsc} approaching the plane through z from above. After possibly shrinking s, we can also assume that $B_{2s}(z) \cap \{x_3 > x_3(z)\}$ is contained in one connected component of the ULSC foliated region. Then, by (C2) of Theorem 0.14, in each compact subset of $B_{2s}(z) \cap \{x_3 > x_3(z)\}, \Sigma_j^+$ will consist of a double spiral staircase for all j sufficiently large. This will be used in (4) to pull the double spiral staircases into an appropriate region near the plane.
- (3) We will show next that Σ_j^+ must be ULSC away from z. Namely, we will show that
 - There exists some $\varepsilon > 0$ so that if $y \in (B_{2s}(z) \setminus B_s(z)) \cap \{x_3 = x_3(z)\},\$

then each component of $B_{\varepsilon s}(y) \cap \Sigma_j$ is a disk for j sufficiently large. We will prove this by contradiction, so suppose that there is a sequence of noncontractible curves $\tilde{\gamma}_j$ in $B_{\varepsilon s}(y) \cap \Sigma_j^+$; the constant ε will be given by Proposition IV.1.1. Applying Proposition IV.1.1 to the $\tilde{\gamma}_j$'s will lead to the desired contradiction. Namely, Proposition IV.1.1 gives a second stable graph $\tilde{\Gamma}_j$ that is disjoint from both Σ_j and Γ_j . Since $\tilde{\Gamma}_j$ starts off from $\tilde{\gamma}_j$, we see that $\tilde{\Gamma}_j$ is above Γ_j . However, it follows from (1) that the axis of the double spiral staircase in Σ_j^+ can be connected to Γ_j by a short curve σ_j in Σ_j ; here short means that the length of σ_j goes to zero as $j \to \infty$. In particular, the short curve σ_j does not pass through the "hole" in the annulus $\tilde{\Gamma}_j$. Therefore, the graph $\tilde{\Gamma}_j$ must intersect $\Sigma_j \cup \Gamma_j$, which is a contradiction.

- (4) Finally, (3) will allow us to apply the one-sided curvature estimate³⁷ to show that the Σ_j⁺'s continue to spiral as graphs below the plane {x₃ = x₃(z)}, contradicting (1). To do this, suppose that y is any given point in (B_{2s}(z) \ B_s(z)) ∩ {x₃ = x₃(z)}. Observe that (1) gives the sequence of stable Γ_j's disjoint from Σ_j⁺ converging in B_s(y) to the horizontal disk B_s(y) ∩ {x₃ = x₃(z)}. It follows from this and (3) that we can apply the one-sided curvature estimate to get that each component of B_{ε's}(y) ∩ Σ_j is a graph for j large; here ε' > 0 depends on ε and the constant from the one-sided curvature estimate. Since these components are graphical for every such point y and start out as part of a multi-valued graph, there are now two possibilities:
 - The multi-valued graph can be continued down to $x_3 = x_3(z) \varepsilon' s$.

³⁷The one-sided curvature estimate from [CM04e] is recalled in Theorem I.1.3.

• The multi-valued graphs spiral infinitely into some horizontal plane above $x_3 = x_3(z) - \varepsilon' s$.

The latter is impossible since each Σ_j is a compact surface. This completes the proof of the fourth step and, hence, gives the promised contradiction to (1).

Part VI. Completing the proofs of Theorems 0.6 and 0.12

The only thing that remains to be proven is that *every* leaf of the lamination \mathcal{L}'^{38} is contained in a plane. This is the remaining claim in Theorem 0.6 (the planar lamination convergence theorem). We have already proven that the leaves of \mathcal{L}' are planes when the sequence is ULSC; thus, by the no mixing theorem, the only remaining case is when $\mathcal{S} = \mathcal{S}_{neck} \neq \emptyset$.

Recall that each point in S_{neck} comes with a plane through it that is a limit of stable graphs in the complement of the sequence Σ_j . Since $S_{\text{neck}} \neq \emptyset$ by assumption, there is at least one such plane and, hence, every leaf of \mathcal{L}' is contained in a half-space (by Lemma IV.2.3). We will divide the proof that the leaves of \mathcal{L}' are contained in planes into two cases, depending on whether or not the leaf is complete. Recall that a leaf Γ is complete when $\Gamma_{\text{Clos}} = \Gamma$, where the closure Γ_{Clos} is defined by fixing a point $x_{\Gamma} \in \Gamma$ and setting

(VI.0.1)
$$\Gamma_{\text{Clos}} = \bigcup_{r} \overline{\mathcal{B}_{r}(x_{\Gamma})}$$

see (II.2.10). Lemma VI.2.1 proves that complete leaves of $\mathcal{L}'s$ are planes; the incomplete leaves will be shown to be planes in Lemma VI.3.1.

It may be useful to give an example of the kind of thing that we need to rule out and a rough idea of why it cannot happen. Suppose therefore that a leaf Γ of \mathcal{L}' contains infinitely many necks, one on top of the next, and that these necks "shrink" to a point $p \in S_{\text{neck}}$. It follows that we have a limit plane through p and that Γ is contained on one side of this plane. We will use a flux argument to rule out such an example. Roughly speaking, we will find a "top" curve with positive flux and then find a sequence of "bottom" curves shrinking to p whose flux goes to zero. We will then show that all of the ends of Γ between these curves are asymptotic to planes or upward sloping catenoids and hence make a nonnegative contribution to the total flux. This will give the desired contradiction since Stokes' theorem implies that the total flux is zero.

VI.1. Blow up results for ULSC surfaces

We will later need to analyze the structure of the sequence Σ_j in Theorem 0.6 (the planar lamination convergence theorem) near points where the

³⁸Recall that \mathcal{L}' is a lamination of $\mathbf{R}^3 \setminus \mathcal{S}$ given in Lemma II.1.2.

topology is concentrating. In doing so, it will often be useful to work on the smallest scale of nontrivial topology. (This is similar to blowing up on the scale of the curvature in the ULSC case; cf. the notion of blow up pairs in [CM04e].) This can be achieved using a simple rescaling argument given in Lemma D.1 in Appendix D.

The advantage of working on the smallest scale of nontrivial topology is that we can use the compactness theorem for ULSC sequences — Theorem 0.9 — to prove a great deal of structure for the surfaces on this scale. We will use two such structure results below:

- Lemma III.4.1 proves a one-sided property for nonsimply connected surfaces on the smallest scale of nontrivial topology. This shows that an intrinsic ball in such a surface cannot lie on one side of a plane and have its center close to the plane on this scale; in the extreme case as the radius of the intrinsic ball goes to infinity, the surface would be forced to grow out of any half-space.
- Lemma VI.1.1 finds short curves on the smallest scale of nontrivial topology separating the ends. These short curves will be used in the flux argument for the main results in this part.

VI.1.1. Finding short separating curves. The next lemma finds short separating curves and stable graphs near points with small injectivity radius. These curves will separate the surface into two parts: a part Σ^+ above the graph and a part Σ^- below the graph; cf. Figure 31. Earlier, in Proposition IV.1.1, we found separating curves contained in small extrinsic balls; in fact (1) and (2) below are proven in Proposition IV.1.1. The new point here is the bound on the length of the curves in (3) below. To prove this length bound, we will work on the smallest scale of nontrivial topology.

LEMMA VI.1.1. Let $\Sigma \subset B_{r_0}$ be an embedded minimal planar domain with $\partial \Sigma \subset \partial B_{r_0}$. Suppose also that $\mathcal{B}_{r_1} \subset \Sigma$ is not a disk and $\Gamma \subset B_{r_0} \setminus \Sigma$ is the stable surface given by Proposition IV.1.1.³⁹ The following three properties hold (the first two are just Proposition IV.1.1):

- (1) Given $\tau > 0$, there exists $C \ge 1$ so that a component Γ_0 of $B_{r_0/C} \cap \Gamma \setminus B_{Cr_1}$ is a graph of a function v with $|\nabla v| \le \tau$ and $\partial \Gamma_0$ intersects both $\partial B_{r_0/C}$ and ∂B_{Cr_1} .
- (2) There are distinct components H^+ and H^- of $B_{r_0/C} \setminus (\Gamma_0 \cup B_{Cr_1})$, a separating curve $\tilde{\sigma} \subset B_{Cr_1} \cap \Sigma$, and components Σ^{\pm} of $B_{r_0/C} \cap \Sigma \setminus \tilde{\sigma}$ so that $\Sigma^{\pm} \subset H^{\pm} \cup B_{Cr_1}$ and $\tilde{\sigma} \subset \partial \Sigma^{\pm}$.

 $^{^{39}\}Gamma$ is given by solving a Plateau problem using a noncontractible curve in $\mathcal{B}_{r_1} \subset \Sigma$ as "interior" boundary.

(3) There exists C_1 so that if $r_0 \ge C_1 r_1$ and $\mathcal{B}_{r_1/4}(x)$ is a disk for each $x \in \mathcal{B}_{C_1 r_1}$, then $\tilde{\sigma}$ is homologous to a collection of curves whose total length is at most $C_1 r_1$.

Remark VI.1.2. We will call this the "short curve lemma" and call the separating curve $\tilde{\sigma}$ the "short curve." Of course, $\tilde{\sigma}$ itself may not be short; rather it is homologous to a collection of curves whose total length is at most $C_1 r_1$.

Proof. Claims (1) and (2) are proven in Proposition IV.1.1. We will prove the last claim by contradiction, so suppose that (3) fails with $C_1 = j$ for a sequence Σ'_j . After rescaling, we can assume that $r_1 = 4$.

Observe first that Lemma B.1 in Appendix B gives a sequence $R_j \to \infty$ so that the component Σ_j of $B_{R_j} \cap \Sigma'_j$ containing 0 is contained in the intrinsic ball \mathcal{B}_{4j} and, hence, each intrinsic ball of radius one in Σ_j is a disk. The Σ_j 's therefore give a ULSC sequence of embedded minimal planar domains in extrinsic balls whose radii go to infinity. As in the proof of Lemma III.4.1, we will now divide into two cases depending on whether or not the curvatures of the sequence blows up.

Case 1. Suppose first that there exists some R so that

(VI.1.3)
$$\limsup_{j \to \infty} \sup_{B_R \cap \Sigma_j} |A|^2 = \infty.$$

We can then apply the compactness theorem for ULSC sequences, Theorem 0.9, to get a subsequence of the Σ_j 's that converges to a foliation by parallel planes away from two lines orthogonal to the leaves of the foliation. It follows easily from the description of the convergence near the lines (as double spiral staircases) that we get a uniform length bound for the separating curve. This length bound is proven in "*The proof of* (P4)" within the proof of Lemma III.1.4 and will not be repeated here.

Case 2. Suppose now that $|A|^2$ is uniformly bounded on each compact subset of \mathbf{R}^3 for the sequence Σ_j . The length bound in this case follows by combining three facts:

• By Lemma B.3 in Appendix B, the uniform curvature bounds implies uniform area bounds for each component of Σ_j in extrinsic balls. (The bound depends on the ball but not on j.) More precisely, if $\Sigma_{j,R}$ denotes the component of $B_R \cap \Sigma_j$ containing 0, then Lemma B.3 implies that

(VI.1.4)
$$\operatorname{Area}\left(\Sigma_{j,R}\right) \le C_c R^2,$$

where the constant C_c depends only on

(VI.1.5)
$$\sup_{B_{C_0R}\cap\Sigma_j}|A|^2$$

The constant C_0 here is universal and does not depend on the upper bounds for the curvature.

• The area bound (VI.1.4) and the co-area formula give uniform length bounds for the boundary $\partial \Sigma_{j,R}$ for most values of R. Precisely, at least one half of the R's between $R_0/2$ and R_0 must satisfy

(VI.1.6) Length
$$(\partial \Sigma_{j,R}) \leq 2 C_c R_0 \leq 4 C_c R$$
.

• In the proof of Proposition IV.1.1, the components of $\partial \Sigma_{j,R}$ were divided into two groups, depending on whether they connected to Σ_j^+ or Σ_j^- ; the separating curve $\tilde{\sigma}_j$ was then chosen to separate these two groups.⁴⁰ However, if we do not ask for a *single* connected separating curve, then we can instead use either of the two groups to separate. Finally, (VI.1.6) gives a uniform bound for the total length.

Therefore, in either case, we get uniform length bounds, hence proving the lemma. $\hfill \Box$

VI.2. Complete leaves of \mathcal{L}'

We will show next that any complete leaf Γ of the lamination \mathcal{L}' is a plane. (\mathcal{L}' is the lamination of $\mathbf{R}^3 \setminus \mathcal{S}$ given in Lemma II.1.2.) Such a leaf Γ is a complete embedded minimal surface in \mathbf{R}^3 , but is not a priori known to be proper.

LEMMA VI.2.1. Suppose that Γ is a complete leaf of the lamination \mathcal{L}' ; i.e, suppose that $\Gamma_{\text{Clos}} = \Gamma$. Then Γ must be a plane.

Proof. We will assume that Γ is not a plane and show that this leads to a contradiction. The ULSC case was already completed in Theorem 0.9, so we can assume that

(VI.2.2)
$$S = S_{neck} \neq \emptyset$$

by the no mixing theorem (Theorem 0.4). Recall that, by Proposition IV.1.1, each point in S_{neck} comes with a plane through it that is a limit of stable graphs in the complement of the sequence Σ_j .⁴¹ Furthermore, by Lemma IV.2.3, the leaves of \mathcal{L}' do not cross any of these planes. Since there is at least one such

⁴⁰This is described in more detail in the proof of Proposition IV.1.1.

⁴¹More precisely, repeatedly applying Proposition IV.1.1 gave a sequence of stable graphs defined over larger and larger annuli, and this sequence converges to a limit graph over a punctured plane. The limit graph is bounded at the puncture and extends smoothly across the puncture to an entire graph; consequently, the limit graph is a plane by the Bernstein theorem.

plane, Γ is contained in a half-space. After a translation and a rotation, we may assume that $\Gamma \subset \{x_3 \ge 0\}$ and

(VI.2.3)
$$\inf_{\Gamma} x_3 = 0.$$

CLAIM. If Γ satisfies (VI.2.3) and is not a plane, then there is a sequence of points $p_n \in \Gamma$ satisfying

$$(\text{VI.2.4}) \qquad \qquad i(p_n) \to 0$$

(VI.2.5)
$$x_3(p_n) \to 0$$

Here $i(p_n)$ is the injectivity radius of Γ at p_n . Since Γ is a complete smooth surface, (VI.2.4) immediately implies that

(VI.2.6)
$$\operatorname{dist}_{\Gamma}(p_1, p_n) \to \infty.$$

Proof of Claim. Since Γ satisfies (VI.2.3) but is not a plane, [MR05] (see the first paragraph of the proof of Lemma 1.5 there) implies that for any $\varepsilon > 0$, we have

(VI.2.7)
$$\sup_{\Gamma \cap \{0 < x_3 < \varepsilon\}} |A|^2 = \infty.$$

Therefore, we get a sequence of points p_n in Γ satisfying (VI.2.5) and with

(VI.2.8)
$$|A|^2(p_n) \to \infty$$

Equation (VI.2.4) then follows immediately from this and the one-sided curvature estimate. *QED of Claim*

Return to the proof of the lemma. We will use the claim above to deduce a flux contradiction (similar to the proof of (\star) in the ULSC case given in Section III.1.1) as follows:

- (a) The leaf Γ must be a multiplicity one limit of the Σ_j 's. To see this, observe that if this were not the case, then the universal cover of Γ would be stable and, hence, flat; cf. the proof of Corollary A.20 for more details.
- (b) Blowing up at p_1 to get a separating curve. Fix a large constant $C_1 > 1$. (It will be chosen depending on both Lemma VI.1.1 — the "short curve lemma" — and Lemma III.4.1 — the "one-sided lemma" for nonsimply connected surfaces.) Applying the blow up lemma, Lemma D.1, at p_1 gives an intrinsic ball

(VI.2.9)
$$\mathcal{B}_{C_1s_1}(y_1) \subset \mathcal{B}_{5C_1i(p_1)}(p_1),$$

so that $\mathcal{B}_{4s_1}(y_1)$ is not a disk but $\mathcal{B}_{s_1}(y)$ is a disk for each $y \in \mathcal{B}_{C_1s_1}(y_1)$.

Taking C_1 sufficiently large, it follows that $\mathcal{B}_{C_1s_1}(y_1)$ satisfies the hypotheses of both Lemma VI.1.1 and Lemma III.4.1 (where the constant H

in Lemma III.4.1 is set equal to a large constant $C_2 > 1$). Hence, the short curve lemma, Lemma VI.1.1, gives an initial short separating curve

(VI.2.10)
$$\gamma_1 \subset B_{4s_1}(y_1) \cap \Gamma$$

and a stable graph

(VI.2.11)
$$\Gamma_0 \subset \mathbf{R}^3 \setminus \Gamma.$$

Since the surface Γ is not known to be proper in all of \mathbb{R}^3 , the graph Γ_0 would at first appear to be defined only over a bounded annulus. However, the multiplicity one convergence of (a) implies that the short curve $\gamma_1 \subset \Gamma$ is actually a smooth limit of curves $\gamma_{1,j}$ contained in the proper surfaces Σ_j . We can therefore apply the barrier construction to these curves in Σ_j and take the limit of the resulting stable graphs to get the desired Γ_0 as a graph defined outside the ball $B_{Cs_1}(y_1)$. It follows that the graph Γ_0 is asymptotic to either a plane or an upward sloping half-catenoid. (The other possibility would be a downward sloping half-catenoid, which is clearly impossible since Γ is above $\{x_3 = 0\}$.)

Moreover, since $\Gamma \subset \{x_3 \ge 0\}$, the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, guarantees that

(VI.2.12)
$$C_2 s_1 < x_3(y_1),$$

where $C_2 > 1$ is a large fixed constant. (We can make C_2 as large as we want by increasing C_1 .) By the same argument, the extrinsic ball $B_{C_2 s_1}(y_1)$ does not intersect any of the horizontal planes associated to the singular set S.

Finally, since $x_3(p_n)$ and $i(p_n)$ both go to zero, we can pass to a subsequence of the p_n 's so that

(VI.2.13)
$$\sup_{\Gamma} x_3 > \sup_{B_{C_{s_1}}(y_1)} x_3,$$

and then for $n \ge 1$,

(VI.2.14)
$$\inf_{B_{C s_n}(y_n)} x_3 > \sup_{B_{C s_{n+1}}(y_{n+1})} x_3.$$

(c) Γ is the only leaf of \mathcal{L}' that intersects $\overline{B_{Cs_1}(y_1)}$. A barrier argument and the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, give $\tilde{C} > C$ so that if $C_1 \geq \tilde{C}$, then only one proper component of $B_{\tilde{C}s_1}(y_1) \cap \mathcal{L}'$ intersects $\overline{B_{Cs_1}(y_1)}$.⁴² Here \tilde{C} depends only on C.

 $^{^{42}}$ This follows exactly as does the analogous result for disks given in Corollary 0.4 in [CM04e]. Namely, if there were two such components, then we could put a stable surface between them. Interior estimates for stable surfaces then imply that each of the original components lies on one side of a plane that comes close to the center of the ball. However,

Stable graph Γ_0 is dashed.



Figure 47. (b): Lemma VI.1.1 gives a short curve $\gamma_1 \subset \Gamma$ and a stable graph Γ_0 .



Figure 48. (c'): $\mathbf{R}^3 \setminus (\Gamma_0 \cup B_{Cs_1}(y_1))$ has components H^+ above and H^- below $\Gamma_0 \cup B_{Cs_1}(y_1)$. Γ is the only leaf of \mathcal{L}' intersecting both H^+ and H^- .

To complete the argument for (c), we need to verify that each component of any leaf of \mathcal{L}' in $B_{\tilde{C}s_1}(y_1)$ is proper. Fortunately, this will follow directly from Lemma B.3, which gives the compactness of each component of an embedded minimal surface in a ball B_R if there is *some* curvature bound in the fixed larger ball $B_{C_d R}$.⁴³ Namely, (VI.2.12) implies that $B_{2C_d \tilde{C}s_1}(y_1)$ is disjoint from \mathcal{S} so long as C_2 is sufficiently large and, thus, every leaf of \mathcal{L}' has bounded curvature in $B_{C_d \tilde{C}s_1}(y_1)$, although the bound may be very large. However, Lemma B.3 requires only some curvature bound to get that each component of $B_{\tilde{C}s_1}(y_1)$ has bounded area and compact closure.

- (c') Γ is the only leaf of \mathcal{L}' that intersects both sides of $B_{Cs_1}(y_1) \cup \Gamma_0$. Since the graph Γ_0 is a limit of surfaces that are disjoint from the Σ_j 's, it follows that none of the leaves of \mathcal{L}' can cross Γ_0 . However, Γ_0 is a graph over an annulus, so the leaves of \mathcal{L}' may "go through the hole" to get from one side of Γ_0 to the other; this is exactly what Γ does. However, by (c), Γ is the only leaf that intersects $B_{Cs_1}(y_1)$, so we conclude that Γ is the only leaf of \mathcal{L}' that intersects both sides of $B_{Cs_1}(y_1) \cup \Gamma_0$.
- (d) Repeating (b) at each p_n . Using (VI.2.4), we can argue as in (b) at each point p_n to get shrinking curves

(VI.2.15)
$$\gamma_n \subset B_{4s_n}(y_n)$$
 where $\mathcal{B}_{C_1s_n}(y_n) \subset \mathcal{B}_{5C_1i(p_n)}(p_n)$

as well as stable graphs that are defined outside $B_{Cs_n}(y_n)$ and are disjoint from Γ . Since $i(p_n) \to 0$, Lemma VI.1.1 gives that the flux across γ_n also goes to zero. Furthermore, (VI.2.6) guarantees that the shrinking curves are separated; the points p_n may be close in \mathbb{R}^3 but they are far apart in Γ .

this would contradict the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, so we conclude that there could not have been two such components.

⁴³Clearly, it is crucial here that C_d does not depend on the bound for the curvature.

Let Γ_n denote the connected component of $\Gamma \setminus (\gamma_1 \cup \gamma_n)$ containing both γ_1 and γ_n in its boundary. Note that we used that γ_1 and γ_n are separating to guarantee that such a component exists.

Finally, let E_n denote the "sandwiched" region in \mathbb{R}^3 that is between the stable graphs associated to p_1 and p_n together with the balls $B_{C_1 s_1}(y_1)$ and $B_{C_1 s_n}(y_n)$. We will need the following two properties of E_n :

(VI.2.16)
$$\Gamma_n \subset E_n \text{ and } \overline{E_n} \cap \mathcal{S} = \emptyset.$$

The first property follows immediately from (2) in Lemma VI.1.1. To see the second, note that a point of S in $\overline{E_n}$ would come with a horizontal plane through it that is disjoint from Γ ; this is impossible since the connected leaf Γ intersects both above and below E_n .

(e) Γ_n is properly embedded. We will prove this by contradiction, so suppose that some Γ_n is not proper. In this case, we would be able to choose a sequence $y_i \in \Gamma_n$ with

(VI.2.17) dist_{$$\Gamma_n$$} $(y_1, y_j) \to \infty$ and $|y_j - y| \to 0$ for some $y \in \Gamma_n \subset E_n$.

Since the union of the leaves of \mathcal{L}' is closed in $\mathbb{R}^3 \setminus \mathcal{S}$ and $\overline{E_n} \cap \mathcal{S} = \emptyset$, the point y must be contained in some leaf $\tilde{\Gamma}$ of \mathcal{L}' . As we have used several times, this implies that the universal cover of $\tilde{\Gamma}$ must be stable; cf. the proof of Corollary A.20 for more details. Since Γ is not stable (see the proof of (a)), it follows that $\tilde{\Gamma} \neq \Gamma$.

We claim that

(VI.2.18)
$$\Gamma$$
 is complete.

Proof of (VI.2.18). We know from (c) that Γ is the only leaf of \mathcal{L}' that intersects $\overline{B_{Cs_1}(y_1)} \cup \overline{B_{Cs_n}(y_n)}$, so

(VI.2.19)
$$y \notin B_{Cs_1}(y_1) \cup B_{Cs_n}(y_n)$$

It is also easy to see that the stable graphs that form the top and bottom of the boundary of E_n cannot be contained in leaves of \mathcal{L}'^{44} , so we must have that

(VI.2.20)
$$y \in E_n \setminus (\overline{B_{Cs_1}(y_1)} \cup \overline{B_{Cs_n}(y_n)}).$$

However, (c') then implies that the entire leaf $\tilde{\Gamma}$ must be trapped inside of E_n . Since $\overline{E_n} \cap S = \emptyset$, it follows that $\tilde{\Gamma}$ must be complete.

QED of (VI.2.18).

⁴⁴These stable graphs were obtained using limits of solutions to Plateau problems using the Σ_j 's as barriers.

Now that we have established (VI.2.18), the Bernstein theorem for stable surfaces implies that $\tilde{\Gamma}$ is a plane. Since $\tilde{\Gamma}$ does not cross $\{x_3 = 0\}$, it must be a horizontal plane. However, this is impossible since $\tilde{\Gamma} \cap \Gamma = \emptyset$ and Γ intersects both above and below E_n . Therefore, we conclude that Γ_n must be proper.

(f) The ends of Γ_n are graphs. We claim next that for each fixed n, there is a constant r_n so that

(VI.2.21)
$$\Gamma_n \cap \{x_1^2 + x_2^2 \ge r_n^2\}$$

consists of a finite collection of graphs over $\{x_3 = 0, x_1^2 + x_2^2 \ge r_n^2\}$.

We will show first that $\Gamma_n \cap \{x_1^2 + x_2^2 \ge r_n^2\}$ is locally graphical. The starting point is to observe that Γ_n is contained in the sandwich E_n and the height of this sandwich grows at most logarithmically. Therefore, by the one-sided curvature estimate, it suffices to prove that Γ_n is scale-invariant ULSC with respect to the distance to 0; see, for instance, (D) in Section III.0.6. This follows from the "between the sheets" argument that we have used several times already, so we will just sketch the proof this time. Namely, since Γ_n is connected, we can fix a curve $\sigma_n \subset \Gamma_n$ that connects γ_1 to γ_n ; we will choose r_n so that

(VI.2.22)
$$\sigma_n \subset \{x_1^2 + x_2^2 < r_n^2/4\}.$$

If $\Gamma_n \cap \{x_1^2 + x_2^2 \ge r_n^2\}$ were to contain a scale-invariant small neck, then a barrier argument would give a stable surface Γ_{barrier} in the complement of Γ_n that is also sandwiched in E_n . This sandwiching and the curvature estimates for stable surfaces imply that the stable surface Γ_{barrier} is graphical away from its boundary. Since the curve σ_n is away from the boundary of the stable surface and connects the top and bottom of the sandwich, the stable surface $\tilde{\Gamma}$ is forced to intersect the curve σ_n , giving the desired contradiction.

After increasing r_n , we conclude that $\Gamma_n \cap \{x_1^2 + x_2^2 \ge r_n^2\}$ is locally graphical and hence a union of graphs over $\{x_3 = 0, x_1^2 + x_2^2 \ge r_n^2\}$. (The other possibility is that it could contain a multi-valued graph; as we have argued before, this is impossible since such a multi-valued graph would have to spiral through the separating plane.) The properness of Γ_n proven in (e) implies that there can only be finitely many such graphs.

Note that, by the isoperimetric inequality, this gives area bounds for Γ_n in compact subsets of \mathbf{R}^3 .

(g) Slicing Γ_n with a plane to get the top curve. Each graphical end of each Γ_n is above $\{x_3 = 0\}$ and, consequently, is asymptotic to either a plane or to an upward sloping half-catenoid. Since there are only finitely many such planes for each n, we can choose a height h between $\sup_{\gamma_2} x_3$ and $\inf_{\gamma_1} x_3$ that misses all of the heights of the planar ends for every Γ_n and so that the

102

plane $\{x_3 = h\}$ intersects Γ transversely. It follows that $\{x_3 = h\}$ intersects each Γ_n transversely in a finite collection of simple closed curves. Note that this plane separates γ_1 from γ_n (and, in particular, does not intersect $\partial \Gamma_n$).

Let Γ'_n denote the component of {x₃ < h} ∩ Γ_n with γ_n in its boundary.
(h) The flux contradiction. The boundary of each Γ'_n consists of a "bottom curve" γ_n together with a collection of closed "top curves" in the plane {x₃ = h}. The collection of top curves is "increasing" in the following sense:

(VI.2.23)
$$\{x_3 = h\} \cap \partial \Gamma_n \subset \{x_3 = h\} \cap \partial \Gamma_{n+1}$$

Generally, one might expect equality in (VI.2.23); however, if Γ_{n+1} contained a catenoidal end that was not in Γ_n , then we would have a strict containment.

The integrand for the vertical flux is pointwise positive along the increasing boundary in $\{x_3 = h\}$ and, hence, the vertical flux of Γ'_n across $\{x_3 = h\}$ is positive and nondecreasing as a function of n. On the other hand, the flux across the bottom curve γ_n goes to zero as $n \to \infty$ by (d). We can therefore fix some large n so that (the absolute value of) the flux across γ_n is less than the flux across $\{x_3 = h\}$. Since Γ'_n has only finitely many ends and each of these ends has nonnegative flux at infinity, the total flux of Γ'_n is positive. This gives the desired contradiction since, by Stokes' theorem, the total flux of Γ'_n must be zero.

Remark VI.2.24. The above argument did not really need that the leaf Γ was complete in order to conclude that it must be flat. Rather, we showed that Γ must be flat as long as there exists a sequence of points $p_n \in \Gamma$ satisfying

(VI.2.25)
$$x_3(p_n) \to 0, \ i(p_n) \to 0, \ \text{and} \ \frac{i(p_n)}{\operatorname{dist}_{\Gamma_{\operatorname{Clos}}}(p_n, \mathcal{S})} \le C_0$$

where C_0 is a fixed constant that does not depend on Γ . This will be useful when we consider incomplete leaves in the next section.

VI.3. Incomplete leaves of \mathcal{L}'

It remains to show that each incomplete leaf Γ of \mathcal{L}' also must be flat. We do this in the next lemma.

LEMMA VI.3.1. Suppose that Γ is an incomplete leaf of the lamination \mathcal{L}' ; i.e, suppose that $\Gamma_{\text{Clos}} \neq \Gamma$. Then Γ is contained in a plane.

As proven earlier in Theorem 0.9, every leaf of \mathcal{L}' is flat when the sequence is ULSC. Therefore, by the no mixing theorem; i.e., Theorem 0.4, we can assume that $\mathcal{S}_{ulsc} = \emptyset$ and $\mathcal{S} = \mathcal{S}_{neck}$. Before getting into the proof, it is useful to consider an example of what a possible incomplete nonflat leaf Γ of \mathcal{L}' would have to look like. By assumption, $\Gamma_{\text{Clos}} \setminus \Gamma \neq \emptyset$ and, hence, $\Gamma_{\text{Clos}} \cap S_{\text{neck}} \neq \emptyset$. Since each point of $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$ comes with a plane through it and none of the leaves of \mathcal{L}' can cross these planes, such a Γ would be contained in either

- an open slab between two singular planes, or
- an open half-space bounded by a singular plane.

Note that, by the strong maximum principle, Γ cannot intersect a singular plane and, hence, we can take the above slab and half-space to be open. We will see in the next subsection that $\Gamma_{\text{Clos}} \cap S$ consists of only one point in the boundary plane(s).

The basic idea behind the proof of Lemma VI.3.1 is again that a potential counterexample would lead to a flux contradiction. Much of the argument is very similar to the complete case:

- Γ will be scale-invariant ULSC away from the singular points.
- Γ will be proper in an open slab or open half-space.
- The ends of Γ will be asymptotic to planes or upward-sloping catenoids.
- We will slice between two planar ends to get a "top curve" with strictly positive flux.
- We will find a sequence of "bottom curves" where the flux goes to zero.

The main difficulty lies in finding the sequence of "bottom curves" where the flux goes to zero. One expects that the injectivity radius of Γ goes to zero as we approach the singular points. However, the rate at which it does so may be quite slow, so we cannot find large regions in Γ "on the smallest scale of nontrivial topology" as the injectivity radius goes to zero. The key for overcoming this will be to get some additional control over Γ near a singular point; in particular, we will prove scale-invariant curvature and area bounds for Γ near each singular point. Once we have this, we can use the co-area formula to find a sequence of "bottom curves" whose length goes to zero.

VI.3.1. If Γ is not flat, then $\Gamma_{\text{Clos}} \cap S$ consists of at most two points. As mentioned, we have already shown that the complete leaves of \mathcal{L}' must be flat, so the remaining case is when

(VI.3.2)
$$\Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}} \neq \emptyset$$

Each point of $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$ comes with a plane through it and none of the leaves of \mathcal{L}' can cross this plane. Hence, by the strong maximum principle, this plane does not intersect any of the nonflat leaves of \mathcal{L}' . The starting point for Lemma VI.3.1 is to show that this plane contains exactly one point of $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$; see Lemma VI.3.3 below. It follows immediately from this that $\Gamma_{\text{Clos}} \cap S$ consists of at most two points for any nonflat Γ .

104

LEMMA VI.3.3. Suppose that $\Gamma \subset \{x_3 > 0\}$ is a nonflat leaf of \mathcal{L}' with $0 \in \Gamma_{\text{Clos}} \cap S_{\text{neck}}$ and $\{x_3 = 0\}$ is the associated stable limit plane through 0. Then we must have

(VI.3.4)
$$\Gamma_{\text{Clos}} \cap \{x_3 = 0\} = \{0\}$$

In fact, if $\Gamma' \subset \{x_3 > 0\}$ is any nonflat leaf of \mathcal{L}' with $\Gamma'_{\text{Clos}} \cap \{x_3 = 0\} \neq \emptyset$, then

(VI.3.5)
$$\Gamma'_{\text{Clos}} \cap \{x_3 = 0\} = \{0\}.$$

Proof. We will first argue by contradiction to prove (VI.3.4). Suppose therefore that there exists $p \neq 0$ with

(VI.3.6)
$$p \in \Gamma_{\text{Clos}} \cap \{x_3 = 0\}.$$

We begin by constructing a curve γ in Γ that connects Γ to 0 — or a singular point near 0 — and stays away from p. Precisely, γ will have the following properties:

(VI.3.7)
$$\gamma: [0,1) \to B_{|p|/3} \cap \Gamma,$$

(VI.3.8)
$$\operatorname{Length}(\gamma) \le |p|/3,$$

(VI.3.9)
$$\lim_{t \to 1} \gamma(t) \in \Gamma_{\text{Clos}} \cap \{x_3 = 0\}.$$

To construct γ , first use the definition of Γ_{Clos} to choose a point $y \in \Gamma$ so that the closure of $\mathcal{B}_{|p|/6}(y) \subset \Gamma$ contains 0. Then choose a sequence of length minimizing curves in Γ that start at y and whose second endpoints converge to 0. The Arzela-Ascoli theorem gives a subsequence of these curves that converges to a curve $\tilde{\gamma}$ that starts at y, ends at 0, and is contained in $\overline{\mathcal{B}}_{|p|/6}(y)$. Finally, let γ be the component of $\Gamma \cap \tilde{\gamma}$ that starts at y.

Note that the curve $\tilde{\gamma}$ might hit another point of \mathcal{S} before it gets to 0. However, this point must be close to 0 and, hence, far from p; this is all that the argument will use. For simplicity, we will assume that 0 was the first point of \mathcal{S} hit by $\tilde{\gamma}$ so that $\lim_{t\to 1} \gamma(t) = 0$.

Since γ is contained in Γ , we get a sequence of curves $\gamma_j : [0, t_j] \to \Sigma_j$ with $t_j \to 1$ and so that the γ_j 's converge to γ . In particular, $\gamma_j(t_j) \to 0$.

CLAIM. The injectivity radius of Σ_j at $\gamma_j(t_j)$ must go to zero.

Proof of Claim. Proposition IV.1.1 gives a stable graph disjoint from Σ_j for each j, and this sequence is converging to $\{x_3 = 0\} \setminus \{0\}$ as j goes to infinity. Moreover, exactly one component of Σ_j in a small ball near 0 intersects both sides of the stable graph. The injectivity radius of this component (obviously) goes to zero as j goes to infinity. It follows that every other component sits on one side of this stable graph; see (B) in Proposition IV.1.1. In particular, if the component of $B_{\varepsilon} \cap \Sigma_j$ containing $\gamma_j(t_j)$ were a disk for some fixed $\varepsilon > 0$

and all sufficiently large j, then the one-sided curvature estimate would imply that this component was graphical in a neighborhood of 0. Moreover, by the strong maximum principle, this sequence of graphs would have to converge to a subset of $\{x_3 = 0\}$. However, these graphs contain subsets of γ_j that are converging to (a component of)

(VI.3.10)
$$B_{\varepsilon} \cap \gamma \subset \Gamma.$$

It follows that $\{x_3 = 0\}$ would have to contain a (smooth) point of the leaf Γ , violating the strong maximum principle. QED for Claim

We can repeat the construction of γ near p to get curves $\gamma'_j : [0, t'_j] \to \Sigma_j$ converging to a curve $\gamma' : [0, 1) \to \Gamma$ so that the endpoints $\gamma'_j(t'_j)$ converge to a singular point near p. For simplicity, we will assume that this second singular point is actually equal to p. Arguing as in the claim, we see that the injectivity radius of Σ_j at $\gamma'_i(t'_j)$ also goes to zero.

We can now apply Proposition IV.1.1 to shrinking balls centered at $\gamma_j(t_j)$ and $\gamma'_j(t'_j)$ to get *disjoint* stable graphs Γ_j and Γ'_j that are disjoint from Σ_j , and so

(VI.3.11) $\Gamma_j \to \{x_3 = 0\} \setminus \{0\} \text{ and } \Gamma'_j \to \{x_3 = 0\} \setminus \{p\}.$

Since Γ_j and Γ'_j are disjoint, one must be above the other. After passing to a subsequence (and possibly switching Γ and Γ'), we can assume that Γ_j is always above Γ'_j . It follows easily from the barrier construction used for the proof of Proposition IV.1.1 that the curve γ'_j must also be below the graph Γ_j .⁴⁵ However, this forces γ'_j to converge to a curve in $\{x_3 = 0\}$, contradicting the strong maximum principle as in the proof of the claim above. This completes the proof of (VI.3.4).

Finally, when $p \in \Gamma'_{\text{Clos}} \cap \{x_3 = 0\}$, the same argument applies with obvious changes. Hence, we also get (VI.3.5).

VI.3.2. The proof of Lemma VI.3.1. As mentioned earlier, we can assume that we are in the case where $S_{ulsc} = \emptyset$, and we will use a flux argument to rule out the possibility of a nonflat leaf of \mathcal{L}' .

Proof of Lemma VI.3.1. We will prove the lemma by contradiction, so suppose that $\Gamma \subset \{x_3 > 0\}$ is a nonflat leaf of \mathcal{L}' with $0 \in \Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}}$ and $\{x_3 = 0\}$ is the associated stable limit plane through 0. By Lemma VI.3.3, there are two possibilities:

- $\Gamma_{\text{Clos}} \cap \mathcal{S} = \{0\}.$
- $\Gamma_{\text{Clos}} \cap S = \{0, p\}$ for some point p with $x_3(p) > 0$.

106

⁴⁵Namely, the stable graph is actually a subset of a stable surface that is disjoint from Σ_j and has interior boundary lying in Σ_j ; this interior boundary connects within Σ_j to the curve γ'_j . This barrier construction is given in Lemma IV.1.2.

 Γ is scale-invariant ULSC near 0. More precisely, there exist $\delta \in (0, 1)$ and $r_0 > 0$ so that

(VI.3.12)
$$\mathcal{B}_{\delta|x|}(x)$$
 is a disk for every $x \in B_{r_0} \cap \{x_3 > 0\} \cap \mathcal{L}'$.

Recall that the argument used to prove that complete leaves of \mathcal{L}' must be flat actually gave a stronger statement; see Remark VI.2.24. This stronger statement implies that (VI.3.12) holds.

 Γ has quadratic curvature blowup near 0. We will next use a compactness argument to prove that there exist constants C_d and $r_1 > 0$ so that

(VI.3.13)
$$|A|^2(x) \le C_d |x|^{-2}$$
 for every $x \in B_{r_1} \cap \{x_3 > 0\} \cap \mathcal{L}'.$

The constant C_d above might depend on \mathcal{L}' , but it will be fixed throughout this proof.

Proof of (VI.3.13). We will argue by contradiction, so suppose that there is a sequence of points $q_n \in \Gamma$ with $q_n \to 0$ and

(VI.3.14)
$$|q_n|^2 |A|^2(q_n) > n.$$

The idea of the proof is that dilating \mathcal{L}' by the factor $|q_n|^{-1}$ about the point q_n gives a sequence of laminations

(VI.3.15)
$$\mathcal{L}_n = |q_n|^{-1} \left(\mathcal{L}' - q_n \right)$$

with $|A|^2(0) > n$ and so that ∂B_1 intersects $S_{\text{neck}}(\mathcal{L}_n)$; here $S_{\text{neck}}(\mathcal{L}_n)$ is the singular set for the rescaled lamination \mathcal{L}_n . Moreover, (VI.3.12) gives a uniform lower bound for the injectivity radius of the leaves of $\{x_3 > 0\} \cap \mathcal{L}_n$ in $B_{1/2}$; see below for more details. Consequently, as n goes to infinity, a subsequence of the \mathcal{L}_n 's would converge to a lamination \mathcal{L}_∞ with

(VI.3.16)
$$0 \in \mathcal{S}_{ulsc}(\mathcal{L}_{\infty}) \text{ and } \partial B_1 \cap \mathcal{S}_{neck}(\mathcal{L}_{\infty}) \neq \emptyset.$$

However, this would contradict the no mixing theorem, so we conclude that the sequence q_n could not have existed.

We need two things to make this outline rigorous. First, we do not have a compactness theorem for sequences of laminations, but rather only for sequences of embedded minimal surfaces. This is easily dealt with since the limit \mathcal{L}_{∞} can be realized as a limit of a diagonal sequence of rescalings of the Σ_j 's; we will omit this standard argument. Second, in (VI.3.12) above we showed only that the leaves of \mathcal{L}' were scale-invariant ULSC near 0; what we need instead is that the sequence Σ_j is itself scale-invariant ULSC near 0. More precisely, we must show that there exists some $\delta_0 \in (0, 1)$ so that for each fixed n, we have

(VI.3.17) every component of $B_{\delta_0|q_n|}(q_n) \cap \Sigma_j$ is a disk for j large.

Since the Σ_j 's are converging to \mathcal{L}' away from \mathcal{S} , — and Γ does not intersect \mathcal{S} — then the component of $B_{\delta|q_n|}(q_n) \cap \Sigma_j$ that is converging to Γ is a disk with large curvature. However, the intrinsic version of the one-sided curvature estimate implies that this is the only component of this ball intersecting a smaller concentric sub-ball about q_n . This gives the remaining ingredient needed to make the proof rigorous.

Extending flatness. We claim that there exist constants $C_{\text{flat}} > 0$ and $r_2 > 0$ so that if $r < r_2, x \in \partial B_r \cap \Gamma$, and

(VI.3.18) $\mathcal{B}_{r/4}(x)$ is a graph with gradient less than C_{flat} over $\{x_3 = 0\}$,

then x is contained in a graph $\Gamma_x \subset \Gamma$ defined over (at least) the annulus

(VI.3.19)
$$\{x_3 = 0, r^2/4 < x_1^2 + x_2^2 < r_2^2\}.$$

In other words, once Γ becomes very flat, then it extends to a very flat graph defined over some annulus of a definite size surrounding the singular point 0. Note that the outer radius r_2 of this annulus is independent of r.

It is easy to prove from the gradient estimate and the quadratic curvature bound (VI.3.13) that x is contained in a very flat graph $\Gamma_{x,r}$ defined over the annulus

(VI.3.20)
$$\{x_3 = 0, r^2/4 < x_1^2 + x_2^2 < C r^2\},\$$

where the constant $C = C(C_{\text{flat}})$ can be as large as we want for C_{flat} sufficiently small. A priori, one might worry that this would give a multi-valued graph. However, by the usual argument, Σ_j cannot contain a multi-valued graph (since it would otherwise be forced to spiral infinitely because of the stable barrier) and, therefore, neither can Γ . It remains to extend $\Gamma_{x,r}$ as a graph all the way out to ∂D_{r_2} for some fixed r_2 . As long as C is sufficiently large, this can be done using the sublinear growth of the height of the graph and the scale-invariant curvature bound. This sublinear growth is proven in Proposition II.2.12 in [CM04b]. Namely, the sublinear growth of the height gives some $\omega_0 > 1$ so that the graph grows sublinearly out to scale $C r/\omega_0$. By the gradient estimate and the curvature bound, the graph has small gradient out at this larger scale too and we can repeat until we get out to a fixed larger scale (independent of r).

 Γ cannot be too "horizontal" near 0. We will show next that

(VI.3.21)
$$\limsup_{s \to 0} \inf_{\Gamma \cap (B_{2s} \setminus B_s)} |\langle \mathbf{n}, (0, 0, 1) \rangle| < 1,$$

where **n** is the unit normal to the surface Γ . Note that $|\langle \mathbf{n}(x), (0, 0, 1) \rangle|$ is equal to one if and only if the tangent plane at x is horizontal.

108
Proof of (VI.3.21). Suppose first for some
$$s$$
 that

(VI.3.22)
$$\inf_{\Gamma \cap (B_{2s} \setminus B_s)} |\langle \mathbf{n}, (0, 0, 1) \rangle| > 0.$$

It follows that Γ is locally graphical with bounded gradient in $B_{2s} \setminus B_s$. By the usual argument, Σ_j cannot contain a multi-valued graph and, therefore, neither can Γ . Hence, (VI.3.22) would imply that $\Gamma \cap (B_{2s} \setminus B_s)$ is a collection of graphs.

Consequently, if we had a uniform lower bound for $|\langle \mathbf{n}, (0, 0, 1) \rangle|$ in any neighborhood of 0, then standard removable singularity theorems for minimal graphs would imply that Γ has a removable singularity at 0. However, Γ would have to be flat by the strong maximum principle if the singularity at 0 was removable. We conclude therefore that

(VI.3.23)
$$\liminf_{s \to 0} \inf_{\Gamma \cap (B_{2s} \setminus B_s)} |\langle \mathbf{n}, (0, 0, 1) \rangle| = 0.$$

Finally, (VI.3.21) follows easily from "extending flatness" and (VI.3.23).

Using (VI.3.21) to blow up \mathcal{L}' . The point about (VI.3.21) is that any limit of rescalings of \mathcal{L}' about 0 will have a nonflat leaf. More precisely, if s_n is any sequence going to zero, then a subsequence of the rescaled laminations

(VI.3.24)
$$\mathcal{L}_n = \frac{1}{s_n} \left(\mathcal{L}' \right)$$

will converge to a lamination \mathcal{L}_{∞} of $\mathbf{R}^3 \setminus \mathcal{S}(\mathcal{L}_{\infty})$ with the following properties:

- (P1) The origin 0 is still in $S_{neck}(\mathcal{L}_{\infty})$ and $\{x_3 = 0\}$ is the corresponding limit plane.
- (P2) The leaves of \mathcal{L}_{∞} satisfy the quadratic curvature bound (VI.3.13) in all of $\{x_3 > 0\}$ (not just in B_{r_1}), the singular set $\mathcal{S}(\mathcal{L}_{\infty})$ does not intersect the half-space $\{x_3 > 0\}$, and 0 is the only singular point in $\{x_3 = 0\}$ "reachable" from $\{x_3 > 0\}$.
- (P3) \mathcal{L}_{∞} contains a nonhorizontal, and hence nonflat, leaf in $\{x_3 > 0\}$.

The lamination \mathcal{L}_{∞} is given as a limit of a subsequence of rescalings of the Σ_j 's; see the proof of (VI.3.13) for such a diagonal argument. Property (P1) follows immediately from this. Property (P2) follows from immediately from (VI.3.13). Finally, (VI.3.21) implies that \mathcal{L}_{∞} contains a nonhorizontal leaf. This nonhorizontal leaf cannot be flat since it would otherwise intersect $\{x_3 = 0\} \setminus \{0\}$, thus giving (P3).

The key point about the rescaled limit lamination \mathcal{L}_{∞} is that it has all of the same properties that \mathcal{L}' did. Therefore, we can repeat the construction to get that limits of rescalings of \mathcal{L}_{∞} also satisfy (P1), (P2), and (P3). This will be important below, so we record it next: (VI.3.25)

Any limit of rescalings of \mathcal{L}_{∞} will also satisfy (P1), (P2), and (P3).

From now on, we will assume that $\{x_3 > 0\} \cap \mathcal{L}'$ has quadratic curvature decay and

(VI.3.26)
$$\{x_3 > 0\} \cap \mathcal{S} = \emptyset;$$

this can be achieved by rescaling as above.

No stable leaves in $\{x_3 > 0\}$. We will show that a lamination (VI.3.27)

 \mathcal{L}_{∞} satisfying (P1), (P2), and (P3) cannot have a stable leaf in $\{x_3 > 0\}$.

The same argument also rules out a leaf in $\{x_3 > 0\}$ whose oriented double cover is stable.

Proof of no stable leaves in $\{x_3 > 0\}$. Suppose instead that \mathcal{L}_{∞} did contain a stable leaf $\tilde{\Gamma}$ in $\{x_3 > 0\}$. We will show first that $\tilde{\Gamma}$ must be a flat plane $\{x_3 = t\}$ for some t > 0. Namely, if it were not flat, then it would be complete away from 0 by Lemma VI.3.3 (see equation VI.3.5) and then Lemma A.26 in Appendix A would give a contradiction.

Since the leaves of \mathcal{L}_{∞} are — by definition — disjoint, it follows that the nonflat leaf Γ of \mathcal{L}_{∞} must be contained in the open slab $\{0 < x_3 < t\}$. Set $t_0 = \sup_{\Gamma} x_3$ so that

(VI.3.28)
$$\Gamma \subset \{x_3 < t_0\}$$

and

(VI.3.29) Γ intersects every tubular neighborhood of the plane $\{x_3 = t_0\}$.

Moreover, the quadratic curvature bound (VI.3.13) for the leaves of \mathcal{L}' implies that

(VI.3.30)
$$\sup_{\{t_0/2 < x_3 < t_0\} \cap \Gamma} |A|^2 \le 4 C_d t_0^{-2} < \infty.$$

However, the three properties (VI.3.28), (VI.3.29), and (VI.3.30) are impossible by the first paragraph of the proof of Lemma 1.5 in [MR05]. We conclude that (VI.3.27) must hold.

 Γ is proper. The first application of (VI.3.27) will be to show that Γ must be proper in compact subsets of $\{x_3 > 0\}$.

Proof of properness. The starting point is that Γ would otherwise accumulate into a stable leaf $\tilde{\Gamma}$; we have used this argument several times and will omit the details (see, e.g., (e) in the proof of Lemma VI.2.1).⁴⁶ Clearly, $\tilde{\Gamma}$ intersects the open half-space $\{x_3 > 0\}$ and, hence, $\tilde{\Gamma}$ must be contained in

⁴⁶To be precise, either $\tilde{\Gamma}$ is stable or its oriented double cover is stable.

 $\{x_3 > 0\}$ by the strong maximum principle. However, this is impossible by (VI.3.27).

Scale-invariant area bounds. Given any $\alpha > 0$, there exists a constant C_{α} so that

(VI.3.31)
$$\operatorname{Area}\left(\left(B_{2r}\setminus B_r\right)\cap \{x_3>\alpha\,|x|\}\cap\Gamma\right)\leq C_\alpha\,r^2.$$

Proof of (VI.3.31). This will be pretty much the same argument as in the proof of " Γ is proper" combined with a compactness argument. We will argue by contradiction, so suppose that (VI.3.31) fails with $r = r_n$ and $C_{\alpha} = n$ for every integer. By a diagonal argument and rescaling, we get a sequence of embedded minimal planar domains $\tilde{\Sigma}_j$ with

Recall that we have proven in (VI.3.25) that a subsequence of the $\tilde{\Sigma}_j$'s converges to a limit lamination \mathcal{L}_{∞} off of a singular set $\mathcal{S}(\mathcal{L}_{\infty})$ satisfying (P1), (P2), and (P3).

Next, we will use (VI.3.32) to show that \mathcal{L}_{∞} contains a stable leaf in $\{x_3 > 0\}$, contradicting (VI.3.27). This would be obvious if \mathcal{L}_{∞} itself had infinite area in $(B_2 \setminus B_1) \cap \{x_3 > \alpha |x|\}$. On the other hand, if \mathcal{L}_{∞} had finite area in $(B_2 \setminus B_1) \cap \{x_3 > \alpha |x|\}$, then the $\tilde{\Sigma}_j$'s must converge with infinite multiplicity to some leaf $\tilde{\Gamma}$ of \mathcal{L}_{∞} that intersects

(VI.3.33)
$$\overline{(B_2 \setminus B_1)} \cap \{x_3 \ge \alpha \, |x|\}.$$

Note that we used that $\overline{(B_2 \setminus B_1)} \cap \{x_3 \ge \alpha |x|\}$ does not intersect the singular set S_{∞} (by (P2)) to guarantee the convergence of the $\tilde{\Sigma}_j$'s in this set. However, as we have used several times, this convergence with multiplicity implies that the leaf $\tilde{\Gamma}$ is stable; see, e.g., the proof of Corollary A.20 for more details.⁴⁷ Finally, since $\tilde{\Gamma}$ intersects the half-space $\{x_3 \ge \alpha\}$, it must be contained in the open half-space $\{x_3 > 0\}$ by the strong maximum principle.

Low points in Γ are contained in graphs. We will need the following complete version of "extending flatness": There exists $\alpha > 0$ so that if $x \in \Gamma$ is in the "low cone" $\{x_3 < \alpha |x|\}$, then x is contained in a graph $\Gamma_x \subset \Gamma$ defined over (at least)

(VI.3.34)
$$\{x_3 = 0, r^2/4 < x_1^2 + x_2^2 < \infty\},\$$

where r = |x|. Moreover, the graph Γ_x must be asymptotic to a plane or to an upward-sloping half-catenoid. Finally, there is a positive lower bound for the

⁴⁷To be precise, this convergence with multiplicity implies that either $\tilde{\Gamma}$ is stable or its oriented double cover is stable.

height of the graph Γ_x ; i.e.,

(VI.3.35)
$$\inf_{\Gamma_x} x_3 > 0.$$

Proof that low points in Γ are contained in graphs. It follows from the gradient estimate and the quadratic curvature decay of Γ that Γ is "very flat" in a neighborhood of x in the sense of "extending flatness." It then follows from the sublinear growth of the height of the graph that Γ_x can then be extended over $\{x_3 = 0, r^2/4 < x_1^2 + x_2^2 < \infty\}$ as long as α is sufficiently small. The proof of this extension will be left to the reader.

Now that we know that Γ_x is defined over $\{x_3 = 0, r^2/4 < x_1^2 + x_2^2 < \infty\}$, it follows that Γ_x is asymptotic to either a plane, an upward-sloping halfcatenoid, or a downward-sloping half-catenoid. The last is impossible since Γ_x is contained in $\{x_3 > 0\}$.

Finally, (VI.3.35) follows from the maximum principle at infinity of [LR88].

The components of $\Gamma \setminus B_r$ are proper. Given any r > 0, then

(VI.3.36) each component of $\Gamma \setminus B_r$ is proper.

Proof of (VI.3.36). To prove (VI.3.36), we must show that any such component Γ_r cannot accumulate into $\{x_3 = 0\}$; this is because we already know that Γ itself is proper in $\{x_3 > 0\}$. We divide this into two cases.

First, suppose that the boundary $\partial \Gamma_r$ of the component Γ_r intersects the "low cone" — i.e., suppose that

(VI.3.37)
$$\inf_{\partial \Gamma_r} x_3 < \alpha r.$$

In this case, it follows that the entire component Γ_r is a graph and, hence, proper.

Suppose now that (VI.3.37) does not hold. In this case, we will find a low component (for some smaller radius) that extends as a graph underneath Γ_r , thus keeping Γ_r strictly away from $\{x_3 = 0\}$. To get this barrier component, note that Lemma E.2 implies that Γ contains a sequence of points $y_n \to 0$ contained in the low cone $\{x_3 < \alpha |x|\}$. If we choose y_n close enough to zero, then the resulting graph Γ_{y_n} must pass underneath $\partial \Gamma_r$ in ∂B_r . It follows that Γ_r sits above $\Gamma_{y_n} \cup B_r$ and, hence, cannot accumulate into $\{x_3 = 0\}$. This completes the proof of (VI.3.36).

The flux contradiction. We will show that the nonflat leaf Γ must contain a sequence of proper subdomains Γ_n with the following properties:

(top) $\partial\Gamma_n$ contains an increasing sequence of compact "top curves" in a fixed plane $\{x_3=h\}$ for some h. Here, increasing means that $\{x_3=h\}$ $\cap \partial\Gamma_n \subset \partial\Gamma_{n+1}$ for every n.

- (ends) Γ_n has finitely many ends and each end is asymptotic to a plane or an upward-sloping half-catenoid.
- (bottom) The rest of $\partial \Gamma_n$ consists of a finite collection of "bottom curves" whose total length goes to zero as n goes to infinity.

This will give a flux contradiction just as in the last step of the proof of Lemma VI.2.1. Namely, the flux of Γ_n across the top curves in $\{x_3 = h\}$ is strictly positive and nondecreasing in n, the ends have nonnegative flux, and the flux across the bottom curves goes to zero. However, this is impossible since the total flux for each Γ_n must be zero by Stokes' theorem. It remains to construct the Γ_n 's with these properties.

We will start with the "top curve" for $\partial \Gamma_n$. As in the proof of (VI.3.36), Lemma E.2 implies that Γ has infinitely many "low" ends that are asymptotic to either planes or upward-sloping half-catenoids. For simplicity, we will assume that these ends are planar; the catenoid case follows similarly and will be left to the reader. Since [LR88] ensures that the planar ends are asymptotic to different planes, we can choose some h > 0 between two consecutive planar ends so that $\{x_3 = h\}$ intersects Γ transversely. Let Γ_h be a component of $\{x_3 < h\} \cap \Gamma$ containing 0 in its closure and fix some component γ_h of $\partial \Gamma_h \subset \{x_3 = h\}$.

Combining the co-area formula with the area bounds from (VI.3.31), we can choose a sequence $r_n \to 0$ so that

(VI.3.38) Length
$$(\partial B_{r_n} \cap \{x_3 \ge \alpha r_n\} \cap \Gamma) \le C r_n$$

for a uniform constant C independent of n. The point here is that the length of these curves goes to zero as n goes to infinity.

For each n, let Γ_n be the component of $\Gamma_h \setminus B_{r_n}$ with γ_h in its boundary. First, it follows immediately that (top) holds. Second, (VI.3.36) implies that Γ_n is proper. Next, when r_n is sufficiently small, then each point in $\partial B_{r_n} \cap \{x_3 < \alpha r_n\} \cap \Gamma$ is contained in a graphical (planar) end that never intersects $\{x_3 = h\}$. In particular, we must have that

(VI.3.39)
$$\partial B_{r_n} \cap \partial \Gamma_n \subset \{x_3 \ge \alpha r_n\}$$

so that the length bound (VI.3.38) gives (bottom). By construction, each Γ_n has compact boundary, is contained in the slab $\{0 < x_3 < h\}$, and has quadratic curvature decay. Therefore, the gradient estimate implies that each component of Γ_n outside of a cylinder $\{x_1^2 + x_2^2 \leq R^2\}$ must be either an asymptotically planar graph or a multi-valued graph. However, as we have used several times, Γ cannot contain such a multi-valued graph, so we conclude that each component of

(VI.3.40)
$$\Gamma_n \cap \{x_1^2 + x_2^2 > R^2\}$$
 is an asymptotically planar graph.

There are only finitely many such ends for each n because Γ_n is proper. This gives (ends) and, hence, completes the proof.

VI.4. The proofs of Theorems 0.6 and 0.12

We now have all of the necessary ingredients to prove Theorems 0.6 and 0.12.

Proof of Theorem 0.6. We have already established properties (A) and (B) of Theorem 0.6 in Lemma II.1.2 and Definition/Lemma II.1.1, respectively. Therefore, it remains to show that every leaf of the lamination \mathcal{L}' is contained in a horizontal plane. Once we have shown this, then the lamination \mathcal{L} is obtained by taking the union of the horizontal planes in \mathcal{L}' together with a horizontal plane through each point in \mathcal{S} .

We have already proven that the leaves of \mathcal{L}' are planar when the sequence is ULSC in Theorem 0.9. Therefore, by the no mixing theorem, Theorem 0.4, the only remaining case is when $\mathcal{S} = \mathcal{S}_{neck} \neq \emptyset$. However, Lemmas VI.2.1 and VI.3.1 together prove that every leaf of \mathcal{L}' is flat in this case. This completes the proof of the theorem.

Theorem 0.12 now follows immediately.

Proof of Theorem 0.12. Now that we have established Theorem 0.6, it only remains to show that property (C_{neck}) holds. However, property (C_{neck}) was proven in (C1) in Theorem 0.14.

Part VII. Modifications in the positive genus case

As we noted earlier, the main theorems were stated for sequences of planar domains, i.e., for genus zero. In this section, we will give the versions of these theorems for sequences with bounded genus and describe the necessary modifications for the proofs. The main change in the theorems is a change in the definitions of the singular sets S_{neck} and S_{ulsc} . The new definitions of S_{neck} and S_{ulsc} , as well as an example showing why a change is necessary, can be found in Section VII.1.1.

Many aspects of the proofs in the genus zero case were essentially local and will, therefore, extend easily once we have the local structure near S_{neck} and S_{ulsc} . However, there are some global aspects to the proofs and these will require some work. The two main global facts are the existence of planes through each singular point and the flatness of nearby leaves (which we often call "properness"). These are "global" in the sense that they fail to hold in the local example constructed in [CM04a].

The definitions of S and \mathcal{L}' are unchanged since Definition/Lemma II.1.1 (that defines the singular set) and Lemma II.1.2 (that constructs \mathcal{L}') did not assume genus zero.

VII.1. The definitions and statements for positive genus

VII.1.1. The sets S_{neck} and S_{ulsc} for positive genus. We will begin with an example illustrating why we have to change the definitions of S_{neck} and S_{ulsc} in the case of positive genus. Namely, let the sequence Σ_j be a sequence of rescalings ("blow downs") of the genus one helicoid constructed in [WHW09]. Since the genus one helicoid is asymptotic to the standard helicoid, the Σ_j 's converge to a foliation by horizontal planes away from the vertical axis. However, the vertical axis contains both the origin where the injectivity radius goes to zero — since the genus concentrates there — and uniformly locally simply connected points. This was impossible in the case of genus zero because of the no mixing theorem.

This example of rescalings of the genus one helicoid illustrates that even if the injectivity radius goes to zero at a point, the point still might not belong in S_{neck} . It is then reasonable to ask what it was about the injectivity radius going to zero that was useful in the genus zero case. The answer is that this allowed one to use a barrier argument near a point $y \in S_{\text{neck}}$ to find stable graphs disjoint from the Σ_j 's that converge to a punctured plane through y. This motivates the following re-definition of S_{ulsc} and S_{neck} :

• A point y in S is in S_{ulsc} if there exist $r_y > 0$ and a sequence $r_{y,j} \to 0$ so that for any $r \in [r_{y,j}, r_y]$ and any connected component Σ'_j of $B_r(y) \cap \Sigma_j$ we have

(VII.1.1) $\partial \Sigma'_{j}$ is connected.

(VII.1.2) Σ'_{j} has the same genus as *one* of the components of $B_{r_{y,j}}(y) \cap \Sigma'_{j}$.

• A point y in S is in S_{neck} if there exist $r_y > 0$ and a sequence $r_{y,j} \to 0$ so that *some* component of $B_{r_{y,j}}(y) \cap \Sigma_j$ has disconnected boundary and (VII.1.2) holds.

These definitions agree with the earlier ones when the sequence is uniformly locally genus zero, i.e., when the genus of $B_{r_y}(y) \cap \Sigma_j$ is zero for every j. In particular, these definitions agree with the earlier ones when the Σ_j 's have genus zero. These definitions are natural in the sense that they allow us to extend the proofs, but note that a shrinking limit of genus one helicoids would be S_{ulsc} and is most certainly not ULSC in the traditional sense.

In the positive genus case, the set S_{neck} is divided into two subsets:

- A point $z \in S_{neck}$ is in S_{neck}^1 if the locally separating curves in Σ_j that are shrinking to z are either
 - globally separating in Σ_i (like in the genus 0 case) or, more generally,
 - globally separating in Σ_j once we combine them with at most g other shrinking curves at other points of S_{neck} .

• The set $S_{\text{neck}}^2 = S_{\text{neck}} \setminus S_{\text{neck}}^1$ consists of at most g "exceptional points" where this does not happen.

The sets S_{neck} and S_{ulsc} are obviously disjoint subsets of S. It follows from Proposition I.0.19 in [CM04d] that, after passing to a subsequence, we can assume that⁴⁸

(VII.1.3)
$$S = S_{neck} \cup S_{ulsc}.$$

The fact that there are at most g "exceptional points" follows immediately from Lemma III.3.1.

VII.1.2. The statements of the theorems for positive genus. We will next run through the changes to the statements of the five theorems — Theorems 0.4, 0.6, 0.9, 0.12, and 0.14 — when the surfaces have positive genus.

The first theorem is the no-mixing theorem, Theorem 0.4; in the positive genus case, this becomes

THEOREM VII.1.4 (No-mixing theorem in the positive genus case). If $\Sigma_i \subset B_{R_i} = B_{R_i}(0) \subset \mathbf{R}^3$ is a sequence of compact embedded minimal surfaces of genus at most g with $\partial \Sigma_i \subset \partial B_{R_i}$ where $R_i \to \infty$, then there is a subsequence so that S^2_{neck} consists of at most g points and either $S_{\text{ulsc}} = \emptyset$ or $S^1_{\text{neck}} = \emptyset$.

Moreover, if $S_{\text{ulsc}} \neq \emptyset$, then the lamination \mathcal{L}' given by Lemma II.1.2 consists of a foliation of (all of) \mathbf{R}^3 by parallel planes away from a singular set S consisting of either one or two lines perpendicular to the planes together with at most g points of S_{neck}^2 .

Theorem 0.6 applies verbatim to the general case of bounded genus with the new definitions of S_{ulsc} and S_{neck} .

On the other hand, Theorem 0.9 holds also for sequences with fixed genus with one minor change in the conclusion and one in the hypothesis. The change in the hypothesis is that we do not assume (0.5). The change in the conclusion is that there might be either one or two singular curves. Assumption (0.5), which says that the Σ_j 's are "uniformly not-disks," was used in the genus zero case to rule out the possibility of just one singular curve (as occurs both for sequences of disks and for rescalings of the genus one helicoid). However, we cannot rule out the possibility of just one singular curve in the fixed genus case regardless of whether we assume (0.5). For this reason, we will not assume (0.5) and we will allow for the possibility of just one singular curve.

Similarly, Theorems 0.12 and 0.14 require small changes in (C_{neck}) . Recall that for each point y in $\mathcal{S}_{\text{neck}}$, (C_{neck}) gives a sequence of graphs in the Σ_j 's

⁴⁸More precisely, this follows from the proof of Proposition I.0.19 in [CM04d]; that proposition was stated for the complementary case where the inner radius is fixed and the outer radii go to infinity.

that converges to a plane through y away from at most two punctures. In the positive genus case, there are now two types of points in S_{neck} and the results are different for each:

- (\mathcal{S}_{neck}^1) If $y \in \mathcal{S}_{neck}^1$, then there is a sequence of graphs in the Σ_j 's that converges to a plane through y away from at most (g+2) punctures.
- (S_{neck}^2) If $y \in S_{\text{neck}}^2$, then there is a sequence of graphs or multi-valued graphs in the Σ_j 's that converges to a plane through y away from at most (g+2) punctures.

We do not know of any examples where $S_{\text{neck}}^2 \neq \emptyset$.

The remainder of this part will be devoted to sketching the modifications needed to prove the main theorems in the general case of bounded genus.

VII.2. The local structure near points in S_{ulsc} and S_{neck}

The starting point for understanding the sequence Σ_j is to describe the sequence in a neighborhood of each singular point, depending on whether the point is in S_{ulsc} or S_{neck} . Roughly speaking, we will get the same picture as in the case of planar domains. The precise statements are

- (α) Given a point x in S_{ulsc} , there is a ball $B_r(x)$ so that
 - (a) The Σ_j 's contain multi-valued graphs that "collapse" to a punctured graph in $B_r(x)$ with a removable singularity at x.
 - (b) The set S satisfies the cone property with respect to this graph in $B_r(x)$.
 - (c) For j sufficiently large, the Σ_j 's are connected near x.⁴⁹
- (β) Given a point x in S_{neck} , there is a ball $B_r(x)$ and a sequence of graphs in the Σ_j 's that converges (with multiplicity at least two) to a finitely punctured graph in $B_r(x)$ with a removable singularity at x.

We will prove (α) first and then (β) . Properties (a) and (b) in (α) give the same structure that Lemma II.2.3 gave in the genus zero case. The proof will follow the same outline as in the genus zero case, with modifications that are standard by this point in the series of papers.

Sketch of proof of (α) . Suppose that $y \in S_{\text{ulsc}}$. We can assume that there is genus concentrating at y (otherwise the genus 0 argument applies). Thus, (by definition) there exist r > 0, a sequence $r_j \to 0$, and points $y_j \in B_{r_j}(y) \cap \Sigma_j$ so that

• For any $s \in [r_j, r]$, the component Σ_{s,y_j} of $B_s(y) \cap \Sigma_j$ containing y_j has positive genus and has connected boundary.

⁴⁹The precise statement is that there exists C > 1 so that if Cs < r and j is sufficiently large, then there is only one connected component of $B_{Cs}(y) \cap \Sigma_j$ that intersects $B_s(y)$.

• If $r_j \leq s_1 < s_2 \leq r$ (and s_1, s_2 are regular values), then $\sum_{s_2, y_j} \setminus \sum_{s_1, y_j}$ is a topological annulus with one inner boundary component in $\partial B_{s_1}(y)$ and one outer boundary component in $\partial B_{s_2}(y)$.

The proof of (a). The first step in (a) is to prove the existence of small multi-valued graphs near y in the Σ_j 's; when the Σ_j 's had genus zero, this was done in [CM04c] by identifying blow up pairs and working on the scale of the maximum of the curvature. This approach does not work here because we do not have any *a priori* relationship between the radii r_j and the maximum of the curvature near y on Σ_j . This difference is the biggest change in the extension of (a) to the positive genus case. Instead, we argue as follows:

- (1) The first observation is that, by Theorem 1.22 in [CM04c], the area of intrinsic sectors over the inner boundaries in the annuli $\sum_{s_2,y_j} \sum_{s_1,y_j}$ must grow faster than quadratically.
- (2) Observe next that these annuli are scale-invariant simply-connected. Namely, if $x \in \sum_{s_2,y_j} \setminus \sum_{s_1,y_j}$ has intrinsic distance $t > 2s_1$ to the inner boundary (and is also not too close to the outer boundary), then $\mathcal{B}_{t/4}(x)$ is a topological disk.

To prove this, suppose instead that the exponential map from x is not injective on $\mathcal{B}_{t/4}(x)$; this would give two geodesics from x with the same endpoint that combine to give a simple closed curve with two break-points. Using the nonpositive curvature and Gauss-Bonnet, we see that this curve cannot bound a disk and, thus, must be homologous to the inner boundary component. However, using Stokes' theorem (applied to $\Delta |x - y|^2$), this would imply small area growth, which is impossible by (1).

- (3) Using Corollary II.2.10 in [CM04d], we can now divide the intrinsic tubular neighborhood of the inner boundary into sectors whose sides are minimizing geodesics (in fact, even minimizing back to the entire inner boundary). The bases of these sectors will be chosen to have a length comparable to a fixed large multiple of s₁.
- (4) By (1), if we choose s_1 so that $\frac{s_1}{r_j}$ is large enough, then we can make the number of disjoint sectors in (3) as large as we would like.
- (5) By (2) and the intrinsic version of the one-sided curvature estimate from [CM01] (recorded here in Theorem I.1.3), if any two of these sectors are sufficiently (scale-invariant) close extrinsically, then they both satisfy a uniform scale-invariant curvature estimate.
- (6) Combining (4) and (5) with Corollary 2.13 of [CM04c], we can arrange that at least one of these sectors is 1/2-stable (with the width and length of the sector fixed, but as large as we wish).
- (7) Finally, (6) allows us to apply Corollary II.1.45 of [CM04d] to get the desired multi-valued graph on a fixed scale.

The extension of this multi-valued graph now follows from (a slight variation of) Theorem II.0.21 in [CM04b], which showed that stable multi-valued graphs extend. Stability there was used for two different purposes: to get some *a priori* scale-invariant bound on curvature, and then to come back and get a better (global) estimate leading to almost flatness. The only notable difference in the current case is that we only have the 1/2-stability as long as the sheets stay close together; this is easily overcome by using the sublinear growth of the separation (i.e., Proposition II.2.12 in [CM04b]) to keep them together.

The proof of (b). The second property that we need is the local cone property. This follows immediately as in (5) from the intrinsic version of the one-sided curvature estimate from [CM01] together with (2) above.

The proof of (c). Finally, we need the local connecting property. This follows immediately from (b) and a barrier argument. Namely, if there were multiple components, then we could use them as barriers to get a stable (thus very flat) surface between them. Using the intrinsic one-sided curvature estimate and simple-connectivity of (2), the multi-valued graph forming in Σ_j would then be forced to spiral graphically forever. This is impossible since each surface is proper.

Sketch of proof of (β) . The structure (β) near a point in S_{neck} follows from a local version of the results of Section IV.1 for the genus zero case. As in the genus zero case, there are three main steps:

- (1) Using the (local) topology to put in a sequence of stable barrier surfaces that converge to a graph through the singular point; see Proposition IV.1.1. This goes through as before, except that the outer radii of the extrinsic balls remains bounded. Hence, the limiting stable graph is defined over a disk and not the entire plane.
- (2) Decomposing Σ_j into ULSC pieces by cutting along "small necks." This goes through as in Section IV.1.2 with only obvious changes.
- (3) Showing that these ULSC pieces contain graphs that converge to the limiting stable graph through the singular point. This goes through as in Section IV.1.3 with obvious changes. □

VII.2.1. Collapsed leaves. The key properties (1), (2), and (3) of collapsed leaves are recorded in Proposition II.3.1. We next extend the proofs of these properties to the positive genus case. The proof of (1) goes through as in Section II.3.1 using the local structure (α) above. The proof of (2) in Sections II.3.2 and II.3.3 goes through with minor modifications that are noted there. (See, e.g., the second paragraph of Lemma II.3.22.) The proof of (3) in Section II.3.4 goes through with the following minor changes:

- In the proof of (3), we used that the "figure eight" curves γ_j were separating in Σ_j. In the genus 0 case, this is automatic since all curves are separating. When the genus is positive, note the infinite multiplicity of the convergence allows us to choose g+1 distinct graphs (all on different sheets) γ_j¹,..., γ_j^{g+1} that are embedded graphs over the curve γ in Γ and all of these have the same orientation (meaning all are on sheets where the "normal points upward"). If none of these is separating on its own, then (since the genus is g and they are all disjoint) Lemma III.3.1 gives a collection of them that together separate in Σ_j. We will use this collection as the inner boundary in the Plateau problem and follow the rest of the argument.
- The other small change is that in the barrier construction, we apply [HS79] rather than [MY82b]. Thus, we do not get an explicit bound on the topology of the stable surface, but this bound was never used in the argument.

VII.3. Part III: When the surfaces are ULSC

This part completed the proof of Theorem 0.9 in the genus zero case, using the tools already developed along with two new ingredients developed there:

- Proposition III.0.2 shows that the closure of a collapsed leaf is a plane.
- Lemma III.1.4 proves "properness."

We follow the same approach in the positive genus case, with minor changes. The first changes are in the statement of Proposition III.0.2, where

- (1) We no longer assume that $\Gamma_{\text{Clos}} \cap S_{\text{neck}} = \emptyset$ but instead make the weaker assumption that $\Gamma_{\text{Clos}} \cap S_{\text{neck}}^1 = \emptyset$. (Weakening this assumption is not necessary for the ULSC results in Part III but will be needed later for the generalization of the no-mixing theorem to the positive genus case.)
- (2) We omit (2) since we make no assumption in the fixed genus case to ensure that there are two axes. As a result, we will need to also consider the case of ULSC, one axis, and finite genus; in the genus zero case this follows already from [CM04e].

Once we have these two things, then the modified Theorem 0.9 will follow as in the genus zero case with one last small change. Namely, we can only apply Meeks' result, [Mee04], at points in the traditional S_{ulsc} (where the sequence is locally simply connected). It follows that the singular set is (one or two) Lipschitz curves in S_{ulsc} and these curves are orthogonal to the planar foliation at all but a finite collection of points; this of course implies that they are orthogonal everywhere.

The proof of (1) in the (modified) Proposition III.0.2 goes through with the following changes:

- Since there are at most g points in S^2_{neck} , we get that $\Gamma_{Clos} \setminus \Gamma$ consists of at most two points in S_{ulsc} together with at most g of the "exceptional points."
- The points in S_{ulsc} are already known to be removable singularities and (a cover of) Γ is already known to be stable. This stability together with (β) allows us to apply Lemma A.35 to conclude that the isolated exceptional singular points are also removable. The claim now follows from a Bernstein theorem as in the genus zero case.

The proof of properness when the genus is zero was given in Lemma III.1.4 using a global flux argument. We will describe the necessary modifications next. Suppose first that the leaf has only one point of S_{ulsc} in it. As in Lemma III.1.4, we need to rule out the possibility of one leaf that spirals into the plane Γ_{Clos} . We would like to appeal to Corollary 0.7 of [CM02b] as in Part I of [CM04e] to get a contradiction, but we will need some modifications:

- (1) We can find the "short curves" (required in [CM02b]) by using the multivalued graph structure that we have already obtained together with Corollary III.3.5 of [CM04d] to get blow up pairs converging down to the singular point from above and then following the argument in [CM04e] (see Corollary IV.0.10 there).
- (2) The leaf is not known to be locally graphical above the plane since we cannot directly apply the one-sided curvature estimate. In particular, as we extend the sheets of the multi-valued graph, we may come to an intrinsic ball that is not scale-invariant simply connected. Because of the closeness to the plane and the sublinear growth proven in [CM04b], we can take the scale-invariant constant to be very small. However, there are at most g of these "bad balls"; otherwise, some combination of curves in these balls would be (globally) separating and we could put in a stable barrier that is forced to "cut the axis" near the singular point; this is a standard variation on the "estimates between the sheets" argument from [CM04b] that we have now employed a number of times.⁵⁰ In particular, the "bad balls" (where it is not simply-connected) can be surrounded by "good balls" and the sheets can be continued globally (with at most g disks removed).
- (3) The last modification is that we may need to start "lower" to ensure that we do not hit any of these "bad balls" as we extend the sheets of the multi-valued graph. Since there are at most g of these and the multi-valued graph

⁵⁰The original "estimate between the sheets" was proven in Theorem I.0.8 in [CM04b]; the version that we use here is essentially (D) in the proof of property (2) in Proposition III.0.2. The difference is that the stable surface may have up to g + 1 inner boundary components, and we use the existence theory of [HS79] instead of Meeks-Yau.

has infinitely many sheets, this is not a problem. The argument now goes through as in Part I on pages 584–593 of [CM04e].

When the leaf contains two points in S_{ulsc} (as was the case in Lemma III.1.4 because of (2) in Proposition III.0.2), the modifications are similar. Namely, the local picture near each singularity is identical and the leaf may fail to be locally graphical over the plane, but only at at most g "bad balls" as in (2). We use the argument in (1) to find the short curves and we argue as in (3) to work "below" these "bad balls" and then follow the proof of Lemma III.1.4.

VII.4. Parts IV and V: When the surfaces are not ULSC

We will next turn to analyzing the structure of non-ULSC singular points, including the proofs — in the positive genus case — of Theorems 0.14 and 0.4 (the no-mixing theorem). To do this, we must prove

- (C1) in Theorem 0.14; this will follow from Proposition VII.4.2 below.
- (C2) and (D) in Theorem 0.14.
- Theorem 0.4.

As in the genus zero case, a key point will be to prove that there is a limit plane through each point in the singular set S. These planes were actually (the closure of) leaves of \mathcal{L}' when the sequence was ULSC, but this was not the case in general. However, these planes were always given as smooth limits of subsets of the Σ_j 's; cf. (C_{neck}) in Theorem 0.12. This is the motivation for the following definition:

Definition VII.4.1. Let Σ_j be a sequence of surfaces with limit lamination \mathcal{L}' and singular set \mathcal{S} . We will say that a surface Γ is a *pseudo-leaf* of \mathcal{L}' if it is *connected* and there is a sequence of subsets $\Sigma_j^{\Gamma} \subset \Sigma_j$ that converges smoothly to Γ . We will also require that Γ is maximal with respect to these properties, so that Γ is not a proper subset of a connected surface that is also a limit of subsets of the Σ_j 's.

Here "converges" means that for each open subset $\Gamma_c \subset \Gamma$ with compact closure in Γ , then the Σ_j 's contain a sequence of graphs — or multi-valued graphs — over Γ_c and these converge smoothly to Γ_c . If we get multi-valued graphs, then we require that the number of sheets goes to infinity as j goes to infinity.

Note that every leaf of \mathcal{L}' is also a pseudo-leaf. We have already come across pseudo-leaves that may not be leaves. Namely, (C_{neck}) in Theorem 0.12 implies that, for each point x in $\mathcal{S}_{\text{neck}}$, we get a flat pseudo-leaf whose closure is a plane through x. This pseudo-leaf is a plane punctured at x and possibly at one other point.

One useful property of a pseudo-leaf is that none of the leaves of \mathcal{L}' can intersect a pseudo-leaf transversely. It then follows from the local structure of nodal sets that the leaves of \mathcal{L}' cannot cross a pseudo-leaf.

The key point for generalizing the main results for non-ULSC sequences from genus zero to fixed genus is to show that

PROPOSITION VII.4.2. For each point $x \in S$, we get a flat pseudo-leaf whose closure is a plane through x. This pseudo-leaf is a plane punctured at at most g + 2 points; each puncture is in S.

We will need one more definition before proving Proposition VII.4.2. Recall that when we studied the leaves of \mathcal{L}' , we began with the collapsed leaves, i.e., the ones "through" a point in \mathcal{S}_{ulsc} . The collapsed leaves were shown to be stable and to have removable singularities at points in \mathcal{S}_{neck} . With this in mind, we will say that a pseudo-leaf Γ is *pinched* if it goes "through" a point in \mathcal{S} . There are two local models for the Σ_j 's near a point x in \mathcal{S} , depending on whether $x \in \mathcal{S}_{ulsc}$ or $x \in \mathcal{S}_{neck}$. First, if $x \in \mathcal{S}_{ulsc}$, then we know that there is a collapsed leaf of \mathcal{L}' through x; see (α) . Second, if $x \in \mathcal{S}_{neck}$, then it follows from (β) that there is a pinched pseudo-leaf through x.

VII.4.1. The local structure (β). We begin by recalling the local structure (β) near points of S_{neck} :

(β) Given a point x in S_{neck} , there is a ball $B_r(x)$ and a sequence of graphs in the Σ_j 's that converges (with multiplicity at least two) to a finitely punctured graph in $B_r(x)$ with a removable singularity at x.

Remark VII.4.3. The structure above is forced to be local because the curves that are shrinking off may not be globally separating in the Σ_j 's. However, if $y \in S^1_{neck}$, then we can argue as in (β) to solve a sequence of global Plateau problems using the Σ_j 's as barriers to get a limiting plane P_y through y so that

- P_y is a smooth limit (of stable graphs disjoint from the Σ_j 's) away from at most g + 1 points in S^2_{neck} . In particular, P_y does not cross any leaves (or pseudo-leaves).
- Observe that P_y cannot contain any points of S_{ulsc} . (If it did, then the multi-valued graphs that developed would be forced to spiral forever contradicting properness of the Σ_i 's.)
- As in (β), we can cut the Σ_j's along a collection of at most g + 1 small necks to get graphs in the Σ_j's that converge to P_y away from at most g + 2 points. This follows as in steps (1) and (2) in Section IV.1.1.
- Finally, as in the genus zero case, we can do this both above and below the stable barriers.

From now on, we will assume that every pseudo-leaf is oriented. This slightly simplifies some of the arguments below involving stability. As we have seen several times, the unoriented case can be dealt with by going to a double cover. We will leave the easy modifications needed for this case to the reader.

VII.5. Planes through S_{ulsc} and the proofs of (C2) and (D)

Suppose now that $x \in S_{ulsc}$ and Γ is the collapsed leaf through it. Generalizing Proposition IV.2.2, we first show that

(0) \mathcal{S}_{neck}^1 does not intersect the closure of the collapsed leaf Γ .

The keys for showing this are the structure result (α) and the following lemma (which generalizes (IV.2.16)):

LEMMA VII.5.1. If Γ is a pinched pseudo-leaf (e.g., a collapsed leaf), y is a point in S^1_{neck} , P_y is the plane through y given by Remark VII.4.3, and $\Gamma_{Clos} \cap P_y \neq \emptyset$, then

(VII.5.2)
$$\Gamma \subset P_u$$
.

Proof. Because of embeddedness of the sequence, it is not hard to see that Γ and P_y cannot cross.⁵¹ We will argue by contradiction and, thus, assume that Γ is above P_y .

The key point will be the following claim:

CLAIM. $\Gamma_{\text{Clos}} \cap S_{\text{neck}}$ is a finite collection of points.

Proof of Claim (proven by a modification of the proof of Lemma VI.3.3). Since Γ cannot cross any of the separating planes through the points in S_{neck}^1 , it follows that $\Gamma_{\text{Clos}} \cap S_{\text{neck}}^1$ is contained in P_y together with at most one other plane parallel to P_y (and above it) and Γ is contained either in the half-space above P_y or in the slab between the two planes. We will show that P_y contains finitely many points in $\Gamma_{\text{Clos}} \cap S_{\text{neck}}^1$; the claim follows from this (together with a similar argument for the second plane in the case of two planes).

We already know that P_y is a smooth limit of stable graphs that are disjoint from the Σ_j 's away from at most g + 1 points and that there are at most g points where the genus is concentrating; let G denote these (at most 2g + 1) "bad points."

We will prove the claim by showing that $P_y \cap \Gamma_{\text{Clos}} \cap S^1_{\text{neck}}$ cannot contain 4g + 3 distinct points. Namely, if it did, then $P_y \cap \Gamma_{\text{Clos}} \cap S^1_{\text{neck}}$ contains two collections

(VII.5.3)
$$\{y_1, \dots, y_{g+1}\}$$
 and $\{z_1, \dots, z_{g+1}\}$

 $^{^{51}}$ In the genus zero case, this is Lemma IV.2.3; the lemma extends easily to the finite genus case.

so that all these points are distinct and disjoint from (the at most 2g + 1 points in) G.

It follows from the one-sided curvature estimate that the injectivity radius of the Σ_j 's is going to zero at the points in Σ_j that are converging to Γ near the y_i 's and z_i 's. Thus, since the genus is at most g, we can choose subcollections of each collection so that they separate in the Σ_j 's and we can put in stable barriers (using the local version of Proposition IV.1.1). This leads to a contradiction as in the end of the proof of Lemma VI.3.3: Namely, we get two distinct stable barriers that separate in space, thus they must be ordered by height, but the limiting surface somehow goes "through" both of them. This contradiction proves the claim.

Once we have shown that Γ is complete away from isolated points on the boundary of a half-space and Γ is contained in this half-space, then Lemma A.35 implies that the isolated exceptional singular points are also removable. The strong maximum principle then gives (VII.5.2).

Here is why (α) and Lemma VII.5.1 imply (0):

Suppose that (0) fails and Γ_{Clos} contains $y \in \mathcal{S}_{\text{neck}}^1$. Let P_y be the limiting plane through y given by Remark VII.4.3, so that Lemma VII.5.1 implies that $\Gamma \subset P_y$.

We get the contradiction from using the barrier graphs to separate the sheets (cf. Remark VII.4.3), which is impossible because of the local connecting property near S_{ulsc} given in part (c) of (α).

Using (0), we can now apply the modified (1) from Proposition III.0.2 to get that the collapsed leaves are all punctured planes, and we can apply the modified Lemma III.1.4 to get that a neighborhood of each point in S_{ulsc} is foliated by collapsed leaves. Thus, (the modified) (C2) and (D) hold.

VII.6. The remaining cases of Proposition VII.4.2

Sketch of the proof of Proposition VII.4.2. Suppose that $x \in S$. We have already dealt with the cases where x is in S_{ulsc} or S_{neck}^1 , so we may assume that $x \in S_{\text{neck}}^2$. Let Γ be the pinched pseudo-leaf through x guaranteed by (β) . We will show next that Γ_x is flat. This follows from stability when Γ_x is complete or if it has only isolated removable singularities (by the usual logarithmic cut-off argument). We will divide into several cases:

• Suppose that Γ_x contains a point of S_{ulsc} in its closure. Since we have already shown that a neighborhood of each point in S_{ulsc} is foliated by flat leaves, we conclude that Γ_x is contained in one of these flat leaves and is, thus, itself a punctured plane. (The bound on the number of punctures has also already been established.)

- Suppose next that Γ_x contains a point y of $\mathcal{S}^1_{\text{neck}}$ in its closure; let P_y be the corresponding plane through y. It follows from Lemma VII.5.1 that $\Gamma \subset P_y$, giving the desired flatness.
- Finally, suppose that Γ_x contains only points in S^2_{neck} in its closure. Since there are at most g of these points and each is a removable singularity (by stability, the structure (β), and Lemma A.35 in Appendix A), we can apply the Bernstein type argument to get flatness.

This completes the proof.

VII.6.1. *Part* V: *The no-mixing theorem*. We will now combine the previous results to extend the no-mixing theorem, Theorem 0.4, to the positive genus case; that is, we will prove Theorem VII.1.4. We must show that

• If $S_{\text{ulsc}} \neq \emptyset$, then the planar collapsed leaves (through S_{ulsc}) foliate all of \mathbb{R}^3 . We have already shown that the foliated region consists of an open set of planes and S_{ulsc} is either one or two straight lines perpendicular to these planes. We must rule out that one of these lines has an endpoint. However, the singular set is closed, so this endpoint would have to be in S_{neck} . Thus (β) gives graphical stable barriers near the endpoint (for j large) that force the spiralling in Σ_j (from the nearby S_{ulsc} points) to continue forever. This contradicts that the Σ_j 's are proper.

VII.6.2. Part VI: The leaves are all flat. The other two global flux arguments are used to show that the leaves are flat in the non-ULSC case, i.e., when $S = S_{\text{neck}}$. This is divided into two cases, depending on whether or not the leaf Γ is complete. The complete leaves were shown to be flat in Lemma VI.2.1 and the incomplete leaves were handled in Lemma VI.3.1. We will next explain how to extend the proofs of these to the positive genus case.

 Γ is complete: Lemma VI.2.1. The point is that Γ must lie in a half-space (since it cannot cross any of the limit planes through \mathcal{S}) and, after a translation and a rotation, we may assume that $\Gamma \subset \{x_3 \geq 0\}$ and

(VII.6.1)
$$\inf_{n \to \infty} x_3 = 0.$$

Arguing as in the claim after (VI.2.3) (using the intrinsic version of the onesided curvature estimate), we get a sequence of points $p_n \in \Gamma$ satisfying

(VII.6.2)
$$i(p_n) \to 0$$

(VII.6.3)
$$x_3(p_n) \to 0,$$

where $i(p_n)$ is the injectivity radius of Γ at p_n . Thus far, there is no difference in the positive genus case.

The contradiction comes from cutting Γ along these "short curves" to get that the flux of Γ is arbitrarily small, which contradicts the strict positivity of

the flux "at the top" that comes from slicing Γ by a plane between two of its ends. This is carried out in steps (a) through (h) of the proof of Lemma VI.2.1. In the positive genus case, we need the following modifications:

- (a) No changes.
- (b) This is where we find the separating curves; this comes almost for free in the genus zero case just because the injectivity radius is going to zero at the p_n 's. When the genus is at most g, then at most g of the balls centered at the p_n 's can have positive genus,⁵² so we throw these out; this allows us to apply the one-sided lemma for nonsimply connected surfaces, Lemma III.4.1, on the remaining balls. There is still another difficulty; namely, the curves in the balls $\mathcal{B}_{5C_1 i(p_n)}(p_n)$ are locally separating, but they may not be globally separating. To deal with this, we group the p_n 's together with (g + 1) of them in each group. We know that some subcollection of each group must be globally separating. This requires obvious changes when we introduce the stable barriers (as we have now done many times).
- (c), (c') No changes.
 - (d) The first part of this is just repeating (b), with the same changes. The second part is to get the properties in (VI.2.16). This follows without change because we only work on the balls that are genus zero and, thus, can still apply Lemma VI.1.1.
 - (e) No changes.
 - (f) Here we use the one-sided curvature estimate and a decomposition into ULSC pieces to show that ends of Γ (above where we cut) are graphical. This is dealt with exactly as in the decomposition around necks. Namely, this can only fail on at most g "bad balls," each of which connects to a finite number of sheets, and each bad ball can be surrounded with graphical pieces.
 - (g) We choose the slicing plane below all of the "bad sheets" from (f).
 - (h) This is where the flux contradiction comes in. The only difference is that instead of one "bottom curve," there may be (g + 1) bottom curves.

 Γ is not complete: Lemma VI.3.1. Suppose instead that $\Gamma \subset \{x_3 > 0\}$ is a nonflat leaf of \mathcal{L}' with $0 \in \Gamma_{\text{Clos}} \cap \mathcal{S}_{\text{neck}}$ and $\{x_3 = 0\}$ is the associated stable limit plane through 0. As in the genus zero case, the argument for incomplete leaves uses short curves to get a flux contradiction. The issue is the construction of the "bottom curves" that required that the injectivity radius was small relative to the distance to the boundary (cf. Remark VI.2.24).

⁵²We are using that we can assume that all the balls $\mathcal{B}_{5C_1 i(p_n)}(p_n)$ are disjoint; cf. (VI.2.6).

• We first modify Lemma VI.3.3 for the positive genus case to get that $\Gamma_{\text{Clos}} \cap \mathcal{S}$ consists of at most 3g + 1 points and there are two possibilities: - $\Gamma_{\text{Clos}} \cap \mathcal{S} \subset \{x_3 = 0\},\$

 $-\Gamma_{\text{Clos}} \cap \mathcal{S} \subset \{x_3 = 0\} \cup \{x_3 = x_3(p) > 0\}$ for some p.

(The proof of this modification follows the original proof of Lemma VI.3.3 with the obvious modifications that we throw away the (at most g) points where the genus is concentrating and then we need to work with two collections of g + 1 points in order to guarantee that they separate globally.) The important point is that Γ fails to be complete only at isolated points.

- Next note that Γ is scale-invariant ULSC near each singular point (cf. (VI.3.12)); this follows as before, except that we may need to throw away g bad balls and work below these.
- Next we show that the curvature blows up at most quadratically near each singular point; see (VI.3.13). We will argue by contradiction, so suppose instead that there is a sequence of points $q_n \in \Gamma$ with $q_n \to 0$ and

(VII.6.4)
$$|q_n|^2 |A|^2(q_n) > n.$$

Thus, the sequence of dilated and translated laminations

(VII.6.5)
$$\mathcal{L}_n = |q_n|^{-1} \left(\mathcal{L}' - q_n \right)$$

satisfies $|A|^2(0) > n$ and the point $\frac{-q_n}{|q_n|} \in \partial B_1$ is in $\mathcal{S}_{\text{neck}}(\mathcal{L}_n)$ (where $\mathcal{S}_{neck}(\mathcal{L}_n)$ is the non-ULSC singular set for the rescaled lamination \mathcal{L}_n). Since \mathcal{L}' is a limit lamination, we can apply these rescalings to a subsequence of the original sequence and use a diagonal argument to get the \mathcal{L}_n 's to converge to a limit \mathcal{L}_∞ . It follows that $\mathcal{S}_\infty = \mathcal{S}(\mathcal{L}_\infty)$ has a ULSC singularity at 0 and the points $\frac{-q_n}{|q_n|}$ converge to a point $q \in \partial B_1$ that is a non-ULSC singularity for \mathcal{L}_{∞} . When the genus is zero, this violates the no-mixing theorem giving the desired contradiction. In the present case of positive genus, this alone is not enough. However, observe that the plane $x_3 = 0$ is the limit of punctured graphs in the Σ_i 's and not multi-valued graphs; this is because $S_{ulsc} = \emptyset$. It follows that the horizontal planes through $\frac{-q_n}{|q_n|}$ are also limits of graphs (in the dilated and translated Σ_j 's). From this, we conclude that the plane $\{x_3 = x_3(q)\}$ is also a limit of graphs and, thus, that $\mathcal{S}_{ulsc} = \mathcal{S}_{ulsc}(\mathcal{L}_{\infty})$ does not intersect this plane. However, this violates the generalized no-mixing theorem, Theorem VII.1.4 (which gives that once $S_{ulsc} \neq \emptyset$, then S_{ulsc} contains a line that intersects every one of the limiting planes), giving the desired contradiction.

• Once we have the quadratic curvature bounds, the rest of the proof follows as in the genus zero case.

Part VIII. Appendices

Appendix A. Surfaces with stable covers

A.1. Going from stability of a covering space to stability of a surface itself. If an oriented minimal surface is stable, then any covering space is also stable. However, the converse may not always be true. The next lemma states that the converse *is true* if in addition the holonomy group of the covering space has sub-exponential growth.

Before showing this, we will need to recall a few elementary properties of groups and covering spaces.

Growth of groups. Suppose that Λ is a finitely generated group, and fix a set of generators. Such a choice of generators induces a natural metric on Λ called the word metric; cf. [Gro81]. Let Λ_n denote the ball of radius *n* about the identity in this metric. The group is said to have sub-exponential growth if we have for every $\varepsilon > 0$ that

(A.1)
$$\lim_{n \to \infty} \frac{|\Lambda_n|}{e^{\varepsilon n}} = 0,$$

where $|\Lambda_n|$ denotes the number of elements of Λ_n .⁵³ Given any fixed integer k, it follows, almost immediately, that sub-exponential growth guarantees that there is a sequence $n_j \to \infty$ with

(A.2)
$$\frac{|\Lambda_{n_j+k} \setminus \Lambda_{n_j}|}{|\Lambda_{n_j}|} \to 0.$$

Covering spaces. Recall that a connected covering space $\hat{\Pi} : \hat{\Gamma} \to \Gamma$ with base point $x \in \Gamma$ is uniquely determined by the holonomy homomorphism Hol from $\pi_1(\Gamma)$ to the automorphisms of the fiber $\hat{\Pi}^{-1}(x)$. To define this homomorphism, suppose that $\gamma : [0,1] \to \Gamma$ is a curve with $\gamma(0) = \gamma(1) = x$ and \hat{x} is a point in $\hat{\Pi}^{-1}(x)$. The lifting property for covering spaces gives a unique lift $\gamma_{\hat{x}} : [0,1] \to \hat{\Gamma}$ of γ with $\gamma_{\hat{x}}(0) = \hat{x}$. We define $\operatorname{Hol}(\gamma)(\hat{x})$ to be the endpoint $\gamma_{\hat{x}}(1)$.

We call the image $\operatorname{Hol}(\pi_1(\Gamma))$ the holonomy group of the covering space; to keep the notation simple, set $\Lambda = \operatorname{Hol}(\pi_1(\Gamma))$.⁵⁴ If we fix a point \hat{x} with $\Pi(\hat{x}) = x$, then we can define a fundamental domain Γ_0 in $\hat{\Gamma}$ by

(A.3) $\Gamma_0 = \{ y \in \hat{\Gamma} | \operatorname{dist}_{\hat{\Gamma}}(y, \hat{x}) \leq \operatorname{dist}_{\hat{\Gamma}}(y, z) \text{ for all } z \in \Pi^{-1}(x) \}.$

Using this, let $\hat{\Gamma}_n = \bigcup_{\lambda \in \Lambda_n} \lambda(\Gamma_0)$ be the covering of Γ corresponding to Λ_n .

 $^{^{53}}$ It is not hard to see that having sub-exponential growth is independent of the choice of generators.

⁵⁴If $\hat{\Gamma}$ is the universal cover, then the holonomy group is exactly the group of deck transformations and, hence, isomorphic to $\pi_1(\Gamma)$. However, the deck group acts transitively on the fiber $\Pi^{-1}(x)$ if and only if $\pi_{\hat{\Gamma}}$ is a normal subgroup of $\pi_1(\Gamma)$; when this is not the case, the holonomy group is bigger than the deck group.

The next property that we will need is a positive lower bound for the distance between $\hat{\Gamma}_n$ and $\hat{\Gamma} \setminus \hat{\Gamma}_{n+k_0}$ when k_0 is sufficiently large. Precisely, if Γ has compact closure (so, in particular, $\pi_1(\Gamma)$ is finitely generated and diam(Γ_0) is finite), then an easy compactness argument gives a constant k_0 so that

(A.4)
$$\operatorname{dist}_{\widehat{\Gamma}}(\Gamma_n, \Gamma \setminus \Gamma_{n+k_0}) > 1.$$

Here k_0 depends on $\hat{\Gamma}$, Γ , and Λ but does not depend on n.

The last fact that we will need is that the holonomy group extends to an action on $\hat{\Gamma}$ when it is abelian; we include a proof for completeness.

LEMMA A.5. If $\hat{\Gamma} \to \Gamma$ is a connected covering space with abelian holonomy group Λ , then Λ extends to an action on $\hat{\Gamma}$ as the group of deck transformations as follows:

Suppose that $\gamma : [0,1] \to \Gamma$ is a curve with $\gamma(0) = \gamma(1) = x$ (where x is the base point in Γ). We have to define the action of $\operatorname{Hol}(\gamma)$ on an arbitrary point \hat{y} in $\hat{\Gamma}$. To do this, choose a curve $\sigma : [0,1] \to \Gamma$ from y to x and define $\operatorname{Hol}(\gamma)(\hat{y})$ to be the second endpoint of the curve starting at \hat{y} that lifts the curve

(A.6)
$$(-\sigma) \circ \gamma \circ \sigma$$

where $(-\sigma)$ denotes the curve σ traversed in the opposite direction.

Proof. The only thing to check is that this definition does not depend on the choice of the curve σ . Suppose therefore that $\mu : [0,1] \to \Gamma$ is a second curve from y to x. It is then easy to see that σ and μ give the same endpoint in (A.6) if and only if

(A.7)
$$\mu \circ (-\sigma) \circ \gamma \circ \sigma \circ (-\mu) \circ (-\gamma)$$

lifts to a closed curve in $\hat{\Gamma}$ starting at \hat{x} . However, the second endpoint of the curve in (A.7) is nothing more than

(A.8) Hol
$$(\mu \circ (-\sigma) \circ \gamma \circ \sigma \circ (-\mu) \circ (-\gamma))(\hat{x})$$

= Hol $(\mu \circ (-\sigma)) \circ$ Hol $(\gamma) \circ$ Hol $(\sigma \circ (-\mu)) \circ$ Hol $(-\gamma)(\hat{x}) = \hat{x}$

since the holonomy group is abelian. This completes the proof.

Stability of covering spaces. We will next show that if a cover of a minimal surface is stable and its holonomy group has sub-exponential growth, then the surface itself is stable. (This simple fact was also later observed independently in [MPR06].) This would be obvious for finite covers; in that case, any compactly supported function on Γ lifts to a compactly supported function on $\hat{\Gamma}$. When the holonomy group is infinite, the lift of a compactly supported function on Γ no longer has compact support. To deal with this, we have to introduce a second cutoff function.

LEMMA A.9. Suppose that $\Gamma \subset \mathbf{R}^3$ is an oriented minimal surface with compact closure, possibly with boundary, and $\hat{\Gamma}$ is a covering space of Γ . If $\hat{\Gamma}$ is stable and its holonomy group has sub-exponential growth, then Γ itself is stable.

Proof. We will show that, for each function $0 \le \phi \le 1$ compactly supported on $\Gamma \setminus \partial \Gamma$, we have the following stability inequality:

(A.10)
$$\int_{\Gamma} |A|^2 \phi^2 \le \int_{\Gamma} |\nabla \phi|^2.$$

Since the holonomy group Λ of the covering space has sub-exponential growth, (A.2) gives a sequence $n_j \to \infty$ with

(A.11)
$$\frac{|\Lambda_{n_j+k_0} \setminus \Lambda_{n_j}|}{|\Lambda_{n_j}|} \to 0,$$

where k_0 is given by (A.4).

Define a sequence of functions ψ_j on $\hat{\Gamma}$ by

(A.12)
$$\psi_j = \begin{cases} 1 & \text{on } \hat{\Gamma}_{n_j}, \\ 1 - \operatorname{dist}_{\hat{\Gamma}}(\hat{\Gamma}_{n_j}, \cdot) & \text{on } \{0 < \operatorname{dist}_{\hat{\Gamma}}(\hat{\Gamma}_{n_j}, \cdot) < 1\}, \\ 0 & \text{otherwise.} \end{cases}$$

In particular, ψ_j is one on $\hat{\Gamma}_{n_j}$, zero outside the 1-tubular neighborhood of $\hat{\Gamma}_{n_j}$ and hence zero outside $\hat{\Gamma}_{n_j+k_0}$ by (A.4). Moreover, ψ_j decays linearly in the distance to $\hat{\Gamma}_{n_j}$ and hence satisfies

(A.13)
$$|\nabla \psi_j| \le 1.$$

Below, we will identify the functions ϕ and $|A|^2$ on Γ with their lifts to the cover $\hat{\Gamma}$.

Although the function ψ_j does not vanish on all of $\partial \hat{\Gamma}$, the function $\psi_j \phi$ does. We can therefore use $\psi_j \phi$ in the stability inequality for $\hat{\Gamma}$ to get (A.14)

$$\begin{split} |\Lambda_{n_j}| \int_{\Gamma} |A|^2 \phi^2 &= \int_{\hat{\Gamma}_{n_j}} |A|^2 \phi^2 \leq \int_{\hat{\Gamma}_{n_j}+k_0} |A|^2 (\psi_j \phi)^2 \\ &\leq \int_{\hat{\Gamma}_{n_j}+k_0} |\nabla(\psi_j \phi)|^2 = \int_{\hat{\Gamma}_{n_j}} |\nabla \phi|^2 + \int_{\hat{\Gamma}_{n_j}+k_0} \langle \hat{\Gamma}_{n_j}} |\nabla(\psi_j \phi)|^2 \\ &= |\Lambda_{n_j}| \int_{\Gamma} |\nabla \phi|^2 + \int_{\hat{\Gamma}_{n_j}+k_0} \langle \hat{\Gamma}_{n_j}} |\nabla(\psi_j \phi)|^2. \end{split}$$

Since ϕ is smooth and has compact support, there is a constant C_{ϕ} so that $2 |\nabla \phi|^2 + 2 \leq C_{\phi}$; hence

(A.15)
$$|\nabla(\psi_j \phi)|^2 \le 2 (|\nabla \phi|^2 + |\nabla \psi_j|^2) \le 2 |\nabla \phi|^2 + 2 \le C_{\phi}.$$

We can use this to bound the last term in (A.14) as follows: (A.16)

$$\int_{\hat{\Gamma}_{n_j+k_0}\setminus\hat{\Gamma}_{n_j}} |\nabla(\psi_j \phi)|^2 \le C_{\phi} \operatorname{Area}(\hat{\Gamma}_{n_j+k_0}\setminus\hat{\Gamma}_{n_j}) = C_{\phi} \operatorname{Area}(\Gamma) |\Lambda_{n_j+k_0}\setminus\Lambda_{n_j}|.$$

Substituting (A.16) into (A.14) gives

(A.17)
$$\int_{\Gamma} |A|^2 \phi^2 \leq \int_{\Gamma} |\nabla \phi|^2 + C_{\phi} \operatorname{Area}(\Gamma) \frac{|\Lambda_{n_j+k_0} \setminus \Lambda_{n_j}|}{|\Lambda_{n_j}|}$$

Finally, (A.11) implies that (A.17) goes to (A.10) as $j \to \infty$, completing the proof.

A.2. A surface and stable cover with cyclic holonomy group where the previous lemma applies. We will show, in Corollary A.20 below, that a certain minimal surface Γ given as a limit of embedded minimal multi-valued graphs Σ_j must be stable. This will follow from Lemma A.9 once we show that there is a connected covering space $\hat{\Gamma}$ satisfying the following two properties:

- The holonomy group Λ of the covering space is cyclic (and, hence, has sub-exponential growth).
- The cover $\tilde{\Gamma}$ is stable.

Throughout this subsection, $\Gamma \subset \mathbf{R}^3$ will be an oriented minimal surface with compact closure, possibly with boundary, and $\Pi : \hat{\Gamma} \to \Gamma$ a covering map with holonomy group \mathbf{Z} (in fact, abelian is sufficient) with the following properties:

- (G1) Σ_j is a sequence of embedded minimal multi-valued (normal exponential) graphs over Γ .
- (G2) There is a sequence $K_1 \subset K_2 \subset \cdots \subset \hat{\Gamma}$ of open domains exhausting $\hat{\Gamma}$ and functions $u_j : K_j \to \mathbf{R}$ with

$$(A.18) |u_j| + |\nabla u_j| \le 1/j$$

so that there is a bijection from from K_j to Σ_j given by

(A.19)
$$x \to \Pi(x) + u_j(x) \mathbf{n}_{\Gamma}(\Pi(x)).$$

Condition (G2) says that the Σ_j 's can be thought of as one to one graphs over the domains K_j in the cover $\hat{\Gamma}$.

COROLLARY A.20. If $\hat{\Gamma} \to \Gamma$ satisfies (G1) and (G2), then the surface Γ is stable.

Proof. By assumption, the holonomy group Λ is cyclic and, thus, has subexponential growth. Therefore, to apply Lemma A.9, we must show that the cover $\hat{\Gamma}$ is stable. We will prove the stability of $\hat{\Gamma}$ by constructing a positive solution w of the Jacobi equation on $\hat{\Gamma}$.

First, since the holonomy group Λ is abelian, Lemma A.5 implies that it acts as the deck group of $\hat{\Gamma}$.

Next, define a sequence of subsets $\tilde{K}_i^o \subset \hat{\Gamma}$ by

(A.21)
$$\tilde{K}_j^o = \{x \in K_j \mid h(1)(x) \in K_j\} = K_j \cap h(1)^{-1}(K_j),$$

where $h(1) \in \Lambda$ is the generator of the infinite cyclic subgroup $\Lambda = \mathbf{Z}$. Fix a point $p \in \tilde{K}_1$, and let \tilde{K}_j be the connected component of \tilde{K}_j^o containing p.

We will need below that the \tilde{K}_j 's are nested, open, connected sets that exhaust $\hat{\Gamma}$. The only point to check is that they exhaust $\hat{\Gamma}$. To see this, suppose that $y \in \hat{\Gamma}$ and choose a path $\sigma : [0,1] \to \hat{\Gamma}$ from p to y. Since the K_j 's are open and exhaust $\hat{\Gamma}$, the compact set $\sigma([0,1]) \cup h(1)(\sigma([0,1]))$ is entirely contained in some K_j for j sufficiently large and, in particular, $\sigma([0,1]) \subset \tilde{K}_j$.

Given $x \in \tilde{K}_j$, both x and h(1)(x) are in K_j and, therefore, we can define functions w_j on \tilde{K}_j by

(A.22)
$$w_j(x) = u_j(h(1)(x)) - u_j(x).$$

Since the bijection (A.19) takes x and h(1)(x) to distinct points in the embedded surface Σ_j and these distinct points have the same projection to Γ , we conclude that

Therefore, we may as well assume that w_j is positive on the connected set K_j . Since u_j and $|\nabla u_j|$ are going to zero by (A.18), a standard calculation (cf. Lemma 2.4 in [CM04c]) gives that u_j almost satisfies the Jacobi equation.⁵⁵ Likewise, the positive function w_j is almost a solution of the Jacobi equation. In particular, if we define normalized functions

(A.24)
$$\tilde{w}_j = \frac{w_j}{w_j(p)},$$

then a subsequence of the \tilde{w}_j 's converges to a positive solution w of the Jacobi equation on $\hat{\Gamma}$ and, thus, $\hat{\Gamma}$ is stable.

A.3. A Bernstein theorem for incomplete surfaces. The results of the previous subsections will be used show that certain *incomplete* minimal surfaces must be stable. We will next prove a Bernstein theorem showing that such a stable surface Γ must then be flat, as long as it is "complete away from a single point." This generalizes the well-known Bernstein theorem for complete stable surfaces of [FCS80], [dCP79].

More precisely, we will assume that the closure Γ_{Clos} of Γ is equal to the union of Γ and a single point. Recall that the closure Γ_{Clos} , defined in (II.2.10),

⁵⁵Precisely, $\Delta u_j + |A|^2 u_j = Q(u_j)$, where the nonlinear term $Q(u_j)$ is at least quadratic in u_j and ∇u_j .

is given by

(A.25)
$$\Gamma_{\text{Clos}} = \bigcup_{r} \overline{\mathcal{B}_r(x_{\Gamma})}$$

The flatness of such a Γ follows from an argument of Gulliver and Lawson, [GL86]; for completeness, we recall this in the next lemma.

LEMMA A.26. Suppose that $\Gamma \subset \mathbf{R}^3$ is a connected stable minimal surface without boundary and with trivial normal bundle. If

(A.27)
$$\Gamma_{\text{Clos}} \setminus \Gamma = \{0\},\$$

then Γ is a (punctured) plane.

Proof. We will use an argument of Gulliver and Lawson, [GL86], to conformally change the metric ds^2 on Γ so that

- (1) The universal cover Γ_U of Γ is complete in the new metric $d\tilde{s}^2$.
- (2) The operator $\tilde{L} = \tilde{\Delta} 2\tilde{K}$ is nonnegative on Γ_U ; i.e., if ϕ is any compactly supported function on Γ_U , then

(A.28)
$$\int \phi \,\tilde{L}\,\phi \le 0.$$

Note that the sign convention here may be the opposite of what one would expect.

Once we have done this, it follows from [FCS80] that $(\Gamma_U, d\tilde{s}^2)$ is conformal to \mathbf{R}^2 with the standard flat metric. Translating back to the original metric ds^2 will then imply that the original Γ was flat.

Following [GL86], we make the conformal change of metric

(A.29)
$$d\tilde{s}^2 = \frac{ds^2}{|x|^2}$$

Since the covering map from Γ_U to Γ is an immersion, the metric $d\tilde{s}^2$ on Γ pulls back to give a metric on Γ_U ; we will also use $d\tilde{s}^2$ to denote this pullback metric. It follows immediately that Γ_U is complete in the new metric $d\tilde{s}^2$. Set $\tilde{L} = \tilde{\Delta} - 2\tilde{K}$ where the Laplacian $\tilde{\Delta}$ and the curvature \tilde{K} are computed with respect to the metric $d\tilde{s}^2$. Corollary 2.13 in [GL86]⁵⁶ gives that

(A.30)
$$\tilde{L} = |x|^2 L - 4(1 - |\nabla |x||^2).$$

⁵⁶Note that our operator L has the opposite sign convention from the operator L_2 in [GL86].

Combining (A.30) with the stability of Γ_U gives for any compactly supported ϕ that (A.31)

$$\int_{\Gamma_U} \phi \, \tilde{L}\phi \, d\tilde{\mu} = \int_{\Gamma_U} \phi \, \left(|x|^2 \, L\phi - 4(1 - |\nabla |x||^2)\phi \right) \, |x|^{-2} \, d\mu \le \int_{\Gamma_U} \phi \, L\phi \, d\mu \le 0 \, ;$$

that is, the operator \tilde{L} is nonnegative on Γ_U with the complete metric $d\tilde{s}^2$. However, Theorem 2 in [FCS80] states that, for any complete surface conformal to the disk, the intrinsically defined operator $\Delta - 2K$ must be negative.⁵⁷ Therefore, since the plane is the only other possible conformal type, we conclude that $(\Gamma_U, d\tilde{s}^2)$ — and hence also (Γ_U, ds^2) — is conformally equivalent to \mathbf{R}^2 . In particular, there is a sequence of compactly-supported logarithmic cutoff functions ϕ_i defined on Γ_U with

(A.32)
$$\phi_j \leq \phi_{j+1} \text{ for every } j \text{ and } \phi_j(x) \to 1 \text{ for every } x \in \Gamma_U,$$

(A.33)
$$\lim_{j \to \infty} \int_{\Gamma_U} |\tilde{\nabla}\phi_j|^2 d\tilde{\mu} = \lim_{j \to \infty} \int_{\Gamma_U} |\nabla\phi_j|^2 d\mu = 0.$$

Using the functions ϕ_j in the stability inequality for L on Γ_U gives

(A.34)
$$-2\int_{\Gamma_U} K \phi_j^2 d\mu \le \int_{\Gamma_U} |\nabla \phi_j|^2 d\mu \to 0.$$

Since $K \leq 0$ and the functions ϕ_j go to 1, we conclude that Γ_U is flat. This completes the proof.

We will also need a local version of this:

LEMMA A.35. If Γ is a connected embedded minimal surface with trivial normal bundle, Γ (or a cover) is stable, and $B_1 \cap \Gamma_{\text{Clos}} \setminus \Gamma = \{0\}$, then Γ has a removable singularity at 0.

Proof. Using estimates for two-sided stable surfaces, the fact that stability is preserved under limits, and Lemma A.26, a compactness argument gives that

$$(A.36) |x| |A| \to 0 \text{ as } |x| \to 0.$$

To show that this scale-invariant curvature decay implies removability, the main point is

(A.37)
$$\langle \mathbf{n}, \frac{x}{|x|} \rangle \to 0 \text{ as } x \to 0.$$

To prove (A.37) by contradiction, we use the scale-invariance of (A.36) and (A.37) to assume that

 $\bullet \ |x|=1, \, |\langle x, {\bf n}\rangle|>\delta>0 \ {\rm and} \ |x|\,|A|<\varepsilon \ {\rm on} \ B_{1/\varepsilon}$

⁵⁷Note that [FCS80] does not assume a sign on the curvature K.

for some very small $\varepsilon > 0$ (which we will take much less than δ). Next, we let γ be an intrinsic geodesic parametrized by arc length with $\gamma(0) = x$ and set $f(t) = |\gamma(t)|^2$. Note that f(0) = 1 and $f'(t) = 2\langle \gamma, \gamma' \rangle$, and $f''(t) = 2 + 2 \langle \gamma, \gamma'' \rangle$. Since $|\langle \gamma(0), \mathbf{n} \rangle| > \delta$, it follows that

$$|f'(0)|^2 \le 2 - 2\,\delta^2.$$

Moreover, (A.36) implies that (as long as γ is defined)

$$|f'' - 2| \le \varepsilon.$$

By comparing f with a quadratic polynomial, we see that f has a positive lower bound for all t as long as ε is small enough. Since this is valid for any geodesic, we see that we cannot reach 0 within distance $1/\varepsilon$ which is impossible since xmust connect in this ball by assumption. This completes the proof of (A.37).

It is easy to finish the proof. Namely, (A.36) and (A.37) imply that each point near 0 is almost contained in a flat plane *through* 0. If these regions closed up as we circled the origin, then 0 is removable by standard theory. Thus, the regions spiral. Any limit leaf that they accumulate on is a very flat graph through 0. Pick a point above this graph, and look at how its height decays as we follow it towards 0. It is easy to see that (A.37) implies that the height decays at least linearly. On the other hand, the sublinear growth of the separation implies that it decays sublinearly, giving the desired contradiction.

Appendix B. An extension of [CM08]

We need slight modifications of several results in [CM08]. We will give these results in this appendix and explain whatever modifications are needed for their proofs.

B.1. *Chord-arc bounds for* ULSC *surfaces*. The next lemma extends the chord-arc bounds of [CM08] from disks to ULSC surfaces.

LEMMA B.1. Given a constant r, there exists R > r so that if Σ is an embedded minimal surface with $\mathcal{B}_R(x_0) \subset \Sigma \setminus \partial \Sigma$ and

(B.2) $\mathcal{B}_1(x)$ is a disk for each $x \in \mathcal{B}_R(x_0)$,

then the connected component of $B_r(x_0) \cap \mathcal{B}_R(x_0)$ containing x_0 has boundary in $\partial B_r(x_0)$.

Proof. The proof follows the proof of Lemma 2.23 given in Appendix B of [CM08] with one modification. (The statement of Lemma 2.23 from [CM08] is recalled below in Lemma B.3.) The difference is that Lemma 2.23 assumed a curvature bound and used this to show that two disjoint intrinsic balls whose centers were close (in \mathbf{R}^3) could be written as graphs over each other. In the

current case, the required curvature bound is not assumed but rather comes from the intrinsic version of the one-sided curvature estimate (Corollary 0.8 in [CM08]).

B.2. Chord-arc and area bounds for surfaces with bounded curvature. We also needed following lemma from [CM08], which gives chord-arc bounds for surfaces with bounded curvature:⁵⁸

LEMMA B.3 (Lemma 2.23 in [CM08]). There exists $C_0 > 1$ so that given a constant C_a , we get another constant C_b such that the following holds: If $\Sigma \subset \mathbf{R}^3$ is an embedded minimal surface with $0 \in \Sigma \subset B_{C_0 R}$ and $\partial \Sigma \subset \partial B_{C_0 R}$ and in addition

(B.4)
$$\sup_{B_{C_0 R} \cap \Sigma} |A|^2 \le C_a R^{-2},$$

then the component $\Sigma_{0,R}$ of $B_R \cap \Sigma$ containing 0 satisfies

(B.5)
$$\Sigma_{0,R} \subset \mathcal{B}_{C_b R}(0)$$

In particular, we also get a constant C_c (depending only on C_a) so that

(B.6)
$$\operatorname{Area}\left(\Sigma_{0,R}\right) \leq C_c R^2.$$

Proof. The first claim (B.5) follows precisely from the proof of Lemma 2.23 in [CM08], which is given in Appendix B in [CM08].⁵⁹

Since $|A|^2$ is bounded on $\Sigma_{0,R}$ by assumption, (B.5) and standard comparison theorems give the area bound (B.6).

A key point in Lemma B.3 is that the constant C_0 does not depend on the constant C_a in the curvature bound (B.4).

Appendix C. Estimates for stable surfaces

Throughout this section, Γ will be a stable surface with connected "interior boundary" γ . We will use $\operatorname{An}_r(\gamma)$ to denote the intrinsic tubular neighborhood of radius r about the curve γ ; i.e.,

(C.1)
$$\operatorname{An}_{r}(\gamma) = \{ x \in \Gamma \, | \, \operatorname{dist}_{\Gamma}(x, \gamma) < r \}.$$

Similarly, we will write $\operatorname{An}_{s,t}(\gamma)$ for the "annulus" $\operatorname{An}_t(\gamma) \setminus \operatorname{An}_s(\gamma)$.

⁵⁸Of course, any surface with bounded curvature is also ULSC and is therefore already covered by Lemma B.1. The usefulness of Lemma B.3 is that it makes the dependence very precise.

⁵⁹We should point out that we have slightly modified the statement of Lemma 2.23 from [CM08]; in particular, the statement in [CM08] assumes that Σ is a disk. However, this was not used in the proof of the lemma given in Appendix B in [CM08].

The main result of this appendix is the following "stable graph" proposition. This proposition shows that a stable embedded minimal surface with a single interior boundary curve γ and an area bound near γ is graphical away from its boundary.

PROPOSITION C.2. Given a constant C, there exists $\omega > 4$ so that if $R > \omega^2$ and $\Gamma \subset B_R$ is a stable embedded minimal surface whose "interior boundary"⁶⁰ $\partial \Gamma \setminus \partial B_R$ is a simple closed curve $\gamma \subset B_4$ satisfying

(C.3)
$$\operatorname{Area}(\operatorname{An}_2(\gamma)) \le C$$

then each component of $B_{R/\omega} \cap \Gamma \setminus B_\omega$ is a graph with gradient bounded by one.

C.1. The regularity of the distance function to the interior boundary. In proving the proposition, we will need some basic results on the level sets of the distance function to an interior boundary curve. Before stating these results, it will be helpful to recall the Gauss-Bonnet theorem with corners and set the notation.

The Gauss-Bonnet theorem with corners implies that a compact surface Σ with piecewise smooth boundary $\partial \Sigma$ satisfies

(C.4)
$$\int_{\partial \Sigma} k_g + \int_{\Sigma} K_{\Sigma} + \sum \alpha_i = 2\pi \, \chi(\Sigma).$$

Here K_{Σ} is the Gauss curvature of Σ , $\chi(\Sigma)$ is its Euler characteristic, and k_g is the geodesic curvature of $\partial \Sigma$. The sign convention of k_g is such that it is *positive* on the boundary of the unit disk in the plane. Finally, the α_i 's are the "jump angles" at the corners of $\partial \Sigma$; see Figure 49. By convention, α_i is positive at a corner where Σ is locally convex. For instance, on each corner of a square, α_i is $\pi/2$.

The next lemma of Shiohama and Tanaka contains the main results that we will need (cf. the proof of Theorem 1 in [Ros03]).

LEMMA C.5 ([ST89], [ST93]). Suppose that Γ is a complete noncompact oriented surface whose boundary $\partial\Gamma$ is a smooth simple closed curve. The set

(C.6)
$$S(t) = \{x \in \Gamma \mid \operatorname{dist}_{\Gamma}(x, \partial \Gamma) = t\}$$

satisfies the following properties:

(*1) For almost every t, the set S(t) is a finite union of piecewise smooth curves with length $\ell(t)$. Let $\alpha_i(t)$ be the "jump angle" at each corner; for consistency with [Fia41] and [Har64], set $\theta_i(t) = -\alpha_i(t)$.

⁶⁰The boundary of Γ is contained in ∂B_R except for a single simple closed curve γ in B_4 ; we call γ the inner boundary and $\partial B_R \cap \partial \Gamma$ the outer boundary, respectively.



Figure 49. The jump angle α_i at a corner.

Concave corners can develop (see $(\star 1)$).



The distance level sets are initially smooth.

Figure 50. The level sets S(t) of the distance function to a curve.

(*2) For almost every t, the derivative $\ell'(t)$ exists and satisfies

(C.7)
$$\ell'(t) = \int_{S(t)} k_g - \sum_{\{i|\theta_i(t)<0\}} \left[2 \tan\left(\frac{\theta_i(t)}{2}\right) - \theta_i\right] - \sum_i \theta_i(t).$$

(*3) Given any $s > r \ge 0$, we get

(C.8)
$$\ell(s) - \ell(r) \le \int_r^s \ell'(t) \, dt.$$

(*4) The area of the "annulus" $\operatorname{An}_{r,s}(\partial\Gamma) = \{x \in \Gamma \mid r \leq \operatorname{dist}_{\Gamma}(x,\partial\Gamma) < s\}$ is

(C.9)
$$\operatorname{Area}\left(\operatorname{An}_{r,s}(\partial\Gamma)\right) = \int_{r}^{s} \ell(t) \, dt.$$

Remark C.10. The papers [ST89] and [ST93] extend earlier results of Fiala, [Fia41], for analytic surfaces and Hartman, [Har64], for simply connected surfaces. Since our surfaces are minimal in \mathbf{R}^3 and, thus, analytic, the classical results of Fiala could be applied here. However, it is useful not to require analyticity so that the results easily generalize to local ones in a Riemannian 3-manifold.

The claim $(\star 1)$ was proven in [ST89], while the claims $(\star 2)$, $(\star 3)$, and $(\star 4)$ appear in [ST93]. Note also that $(\star 4)$ follows from the co-area formula. We should note that the formula (C.7) does not appear explicitly in [ST93] but is implicit there and can also be found in Section 9.6 of [Fia41].

We will need three additional properties of the level sets S(t) that hold if in addition Γ is stable: (*2') If Γ is embedded minimal and *stable*, then each $\theta_i(t) > 0$ and $\ell'(t)$ from (*2) satisfies

(C.11)
$$\ell'(t) = \int_{S(t)} k_g - \sum_i \theta_i(t) \, .$$

Notice that the right-hand side of (C.11) is exactly the boundary term corresponding to S(t) in the Gauss-Bonnet formula with corners.

(*5) There is a constant C_g so that if Γ is embedded minimal and *stable*, then we get the upper bound

(C.12)
$$\sup_{S(t)} k_g \le C_g t^{-1}$$

Recall that given our sign convention for k_g , (C.12) means that S(t) cannot be "too convex" when it is thought of as part of the boundary of $\operatorname{An}_t(\partial\Gamma)$.

(*6) There is a constant ε_g so that if Γ is embedded minimal and *stable* and $\sigma \subset S(t)$ is a closed curve with

(C.13)
$$\operatorname{Length}(\sigma) \le \varepsilon_q t,$$

then σ bounds a disk $\Gamma_{\sigma} \subset \Gamma$ and $\Gamma_{\sigma} \subset \operatorname{An}_{3t/4, 4t/3}(\partial \Gamma)$.

Proof of
$$(\star 2')$$
, $(\star 5)$ and $(\star 6)$. Given $p \in S(t)$, let

(C.14)
$$\gamma_p: [0,t] \to \Gamma$$

be a minimizing geodesic connecting $\partial\Gamma$ to p. The triangle inequality then implies that S(t) does not intersect the interior of the geodesic ball $\mathcal{B}_s(\gamma_p(t-s))$ for any s between 0 and t. Standard comparison theorems and the curvature estimate for stable surfaces then give some a > 0 and C_g so that $\partial \mathcal{B}_{at}(\gamma_p(t-at))$ is a smooth curve with geodesic curvature at most $C_g t^{-1}$. Since the p is in the boundaries of these balls, we see that

- The jump angles $\alpha_i(t)$ at the corners of S(t) can only be negative as claimed in $(\star 2')$.
- If p is a smooth point of S(t), then $k_g(p) \le C_g t^{-1}$, giving (*5).

Note that we do not claim a lower bound for k_g . (In fact, easy examples show that k_g can go to $-\infty$; see [Har64].)

To finish $(\star 2')$, note that (C.7) implies (C.11) since $\theta_i(t)$ is between 0 and $\pi/2$.

To see (*6), fix a point $p \in \sigma$ and note that the entire curve σ is contained in the intrinsic geodesic ball $\mathcal{B}_{\varepsilon_g t}(p)$ and this ball stays away from $\partial \Gamma$. Taking ε_g small, the interior curvature estimates for stable surfaces imply that $\mathcal{B}_{\varepsilon_g t}(p)$ is a graph over some plane. In particular, the curve σ is contractible in $\mathcal{B}_{\varepsilon_g t}(p)$, giving the desired Γ_{σ} .

Remark C.15. We will actually use a very slight generalization of these results. Namely, in applications, Γ will not be complete but will rather be allowed to have other boundary components. This does not matter since we will always work with level sets S(t) where t is less than the distance to any of the other boundary components. It's easy to see that the above results extend to this case.

C.2. The proof of the "stable graph" proposition. The key point for proving Proposition C.2 will be to show that Γ has quadratic area growth. This quadratic area estimate formally follows from the argument in [CM02a], but we need the results of the previous subsection to deal with technical difficulties that arise from the lack of regularity of the level sets of the distance function.

Proof of Proposition C.2. The key point is to prove that the intrinsic annuli $\operatorname{An}_r(\gamma)$ have quadratic area growth:

(C.16)
$$\operatorname{Area}\left(\operatorname{An}_{r}(\gamma)\right) \leq C_{1} r^{2} + C,$$

where the constant C_1 depends on the constant C in (C.3). Once we have (C.16), the lemma follows easily from the proof of Lemma II.1.34 in [CM04d]. For the reader's convenience, next we will sketch the proof of the lemma assuming (C.16):

First, (C.16) allows us to use a logarithmic cutoff function to get sub-annuli with small total curvature. Since these sub-annuli have small total curvature and are stable, the mean value inequality gives a small scale-invariant pointwise curvature estimate. Here the scale refers roughly to the distance to γ . In particular, integrating this curvature bound implies that each component of a level set of the distance to γ is itself a graph over (a curve in) *some* plane. Moreover, Proposition 1.12 in [CM01] uses the fact that the Gauss map is conformal to piece these together and get a graph over one fixed plane, as desired.

It remains therefore to establish (C.16). Note that Proposition II.1.3 in [CM04d] gives (C.16) directly under the additional assumption that Γ is a topological annulus. We will see that the general case follows similarly if we also use the regularity of the length of level sets of the distance function from γ given by Lemma C.5.

The proof of (C.16). There are two steps in the proof of (C.16):

- (1) The stability inequality allows us to bound the total curvature in terms of the energy of a cutoff and this in turn is bounded by the area.
- (2) The area growth is always controlled by the total curvature. This follows easily from Gauss-Bonnet when the exponential map is smooth but holds more generally by Lemma C.5.

Step (1). Set $d(\cdot) = \text{dist}_{\Gamma}(\gamma, \cdot)$, and define a (radial) cut-off function ϕ by

(C.17)
$$\phi = \begin{cases} d & \text{on } \operatorname{An}_1(\gamma), \\ (r-d)/(r-1) & \text{on } \operatorname{An}_{1,r}(\gamma), \\ 0 & \text{otherwise }. \end{cases}$$

By the stability inequality applied to ϕ , we get

(C.18)

$$\int_{\operatorname{An}_{1,r}(\gamma)} |A|^2 \left[(r-d)/(r-1) \right]^2 \leq \int |A|^2 \phi^2 \leq \int |\nabla \phi|^2$$

$$\leq \operatorname{Area} \left(\operatorname{An}_1(\gamma) \right) + (r-1)^{-2} \operatorname{Area} \left(\operatorname{An}_{1,r}(\gamma) \right).$$

If we set $K(s) = \int_{\operatorname{An}_{1,s}(\gamma)} |A|^2$, then the co-area formula gives

(C.19)
$$K(s) = \int_{\operatorname{An}_{1,s}(\gamma)} |A|^2 = \int_1^s \left[\int_{\{d=t\}} |A|^2 \right] dt.$$

In particular, we can integrate by parts twice to get (C.20)

$$2(r-1)^{-2} \int_{1}^{r} \int_{1}^{t} K(s) ds \, dt = 2/(r-1) \int_{1}^{r} K(s)(r-s)/(r-1) ds$$

= $\int_{1}^{r} K'(s) ((r-s)/(r-1))^{2} ds$
 $\leq \operatorname{Area} (\operatorname{An}_{1}(\gamma)) + (r-1)^{-2} \operatorname{Area} (\operatorname{An}_{1,r}(\gamma)),$

where the last inequality is (C.18).

Step (2). We will now use Lemma C.5 to estimate the area by the total curvature. Set $\ell(t)$ equal to the length of the level set $\{d = t\}$. The key will be to prove the following estimate for $\ell(t)$ for $t \ge 1$:

(C.21)
$$\ell(t) \le C_2 (1+t) + \frac{1}{2} \int_1^t \int_{\operatorname{An}_{1,s}(\gamma)} |A|^2 \, ds = C_2 (1+t) + \frac{1}{2} \int_1^t K(s) \, ds,$$

where C_2 depends only on the constant C in (C.3).

The proof of the proposition assuming (C.21). Integrating the length bound (C.21) gives the area bound

(C.22)
$$\operatorname{Area}\left(\operatorname{An}_{r}(\gamma)\right) \leq \operatorname{Area}\left(\operatorname{An}_{1}(\gamma)\right) + \int_{1}^{r} \ell(t) dt$$
$$\leq C + C_{2} r + C_{2} r^{2}/2 + \int_{1}^{r} \int_{1}^{t} \frac{K(s)}{2} ds dt.$$

Combining (C.20) and (C.22) gives the needed bound (C.16).

The proof of (C.21). We will prove (C.21) by integrating a bound on $\ell'(t)$. There will be two steps — namely, bounding $\ell'(t)$ and then finding some value of t where $\ell(t)$ is bounded. (This is where we will integrate the bound on $\ell'(t)$ from.)

The bound on $\ell'(t)$. Roughly speaking, we will bound $\ell'(t)$ in terms of the total curvature by using the Gauss-Bonnet theorem in the "annulus" $\operatorname{An}_{t_0,t}$ for a specific choice of t_0 . Recall that $(\star 2')$ implies that $\ell'(t)$ is bounded by the Gauss-Bonnet terms corresponding to S(t). To get the desired upper bound, we will need to control the contributions from the geodesic curvature of the "inner boundary" $S(t_0)$ as well as the Euler characteristic of the "annulus." We will do this next.

First, the area bound (C.3) and (*4) imply that there must exist some $t_0 \in (1/3, 2/3)$ with

(C.23)
$$\ell(t_0) \le 3 \operatorname{Area} (\operatorname{An}_{1/3,2/3}(\gamma)) \le 3 C.$$

Moreover, by the regularity property $(\star 1)$, we may assume that the level set $S(t_0)$ is a finite union of simple closed piecewise smooth curves. We will sort these curves into two groups, depending on their length. Namely, let $\sigma_1^{\text{long}}, \ldots, \sigma_n^{\text{long}}$ be the components of $S(t_0)$ with

(C.24)
$$\operatorname{Length}(\sigma_i^{\operatorname{long}}) \ge \varepsilon_g/3,$$

where ε_g is given by (*6). Let $\sigma_1^{\text{short}}, \ldots, \sigma_m^{\text{short}}$ be the remaining components. Combining (C.24) with the upper bound on the total length of $S(t_0)$ from (C.23) immediately gives the bound

(C.25)
$$n \le n(C),$$

where n(C) depends only on the area bound (C.3).⁶¹

We will not actually apply the Gauss-Bonnet theorem to all of $\operatorname{An}_{t_0,t}(\gamma)$, but rather to the subset Γ_t that "sees the outer boundary" S(t). To be precise, define Γ_t to be the union of all connected components of $\operatorname{An}_{t_0,t}(\gamma)$ whose boundaries intersect S(t); see Figure 51.

By construction, we have

(C.26)
$$S(t) \subset \partial \Gamma_t \text{ and } \partial \Gamma_t \setminus S(t) \subset S(t_0).$$

Consequently, combining the length bound (C.23) with the pointwise geodesic curvature bound (\star 5), we get a total (geodesic) curvature bound for $\partial \Gamma_t \setminus S(t)$

(C.27)
$$\int_{\partial \Gamma_t \setminus S(t)} k_g \le \ell(t_0) \sup_{S(t_0)} k_g \le k(C),$$

⁶¹We are not claiming a bound on the total number m + n of components of $S(t_0)$.



Figure 51. An example illustrating Γ_t in a case where $\Gamma_t \neq \operatorname{An}_{t_0,t}(\gamma)$.

where k(C) depends only on the area bound (C.3). We should make two remarks about (C.27):

- The integration in (C.27) is over only the smooth part of $\partial \Gamma_t \setminus S(t)$.
- The sign convention on k_g in (C.27) is as part of the boundary of $\operatorname{An}_{t_0}(\gamma)$; this is the opposite as it would be as part of the boundary of Γ_t . This is important later when we apply the Gauss-Bonnet theorem.

The last ingredient that we will need to bound $\ell'(t)$ is a bound on the Euler characteristic $\chi(\Gamma_t)$ that depends only on the area bound (C.3). This bound follows immediately from the bound (C.25) on the number of long components of $S(t_0)$ together with the following claim:

CLAIM. For $t \geq 3/4$, each connected component of Γ_t , i.e., each component of $\operatorname{An}_{t_0,t}(\gamma)$ whose boundary touches S(t), contains at least one long component $\sigma_i^{\operatorname{long}}$ in its boundary.

The point here is that the short components of $S(t_0)$ are contractible near $S(t_0)$, so S(t) never sees them. More precisely, (*6) implies that each σ_i^{short} bounds a disk

(C.28)
$$\Gamma_i^{\text{disk}} \subset \operatorname{An}_{1/4,3/4}(\gamma).$$

Therefore, if p is an arbitrary point in S(t), then we know that p and γ are in the same connected component of

(C.29)
$$\operatorname{An}_t(\gamma) \setminus \bigcup_{i=1}^m \sigma_i^{\operatorname{short}}.$$

Note that we have used here that $\operatorname{An}_t(\gamma)$ is itself connected. Since $S(t_0)$ separates $\gamma = S(0)$ from S(t), we conclude that it must be $\cup_i \sigma_i^{\operatorname{long}}$ that separates p and γ . In particular, the component of $\operatorname{An}_{t_0,t}$ with p in its boundary also contains at least one $\sigma_i^{\operatorname{long}}$ in its boundary. This completes the proof of the claim.

We can now bound $\ell'(t)$ for $t \geq 3/4$. Namely, $(\star 2')$ implies that $\ell'(t)$ is bounded by the Gauss-Bonnet integrand along S(t) so the Gauss-Bonnet
theorem gives for almost every t that

(C.30)
$$\ell'(t) \le \int_{S(t)} k_g - \sum_i \theta_i(t) \le \frac{1}{2} \int_{\Gamma_t} |A|^2 + 2\pi \chi(\Gamma_t) + \int_{\partial \Gamma_t \setminus S(t)} k_g$$

We have thrown away the angle contributions at the corners of $\partial \Gamma_t \setminus S(t)$ in (C.30) since these are all negative by $(\star 2')$. Since $\Gamma_t \subset \operatorname{An}_{1/3,t}(\gamma)$, we can use interior curvature estimates for stable surfaces and the area bound on $\operatorname{An}_1(\gamma)$ to get

$$\int_{\Gamma_t} |A|^2 \le \int_{\operatorname{An}_{1,t}(\gamma)} |A|^2 + \operatorname{Area}(\operatorname{An}_1(\gamma)) \sup_{\operatorname{An}_{1/3,1}(\gamma)} |A|^2 \le \int_{\operatorname{An}_{1,t}(\gamma)} |A|^2 + C_3,$$

where C_3 depends only on the initial area bound (C.3). Substituting the above bounds into (C.30), we get for almost every $t \ge 3/4$ that

(C.32)
$$\ell'(t) \le \frac{1}{2} \int_{\operatorname{An}_{1,t}(\gamma)} |A|^2 + C_4,$$

where C_4 depends only on the initial area bound (C.3).

To complete the proof, use the area bound and $(\star 4)$ again to find t_1 between 3/4 and 1 with $\ell(t_1) \leq 4C$. Given $t \geq 1$, we can then use $(\star 3)$ to integrate (C.32):

(C.33)
$$\ell(t) \le \ell(t_1) + \int_{t_1}^t \ell'(s) \, ds$$
$$\le 4 \, C + C_4(t - t_1) + \frac{1}{2} \, \int_{t_1}^t \int_{\operatorname{An}_{1,s}(\gamma)} |A|^2 \, ds.$$

This gives (C.21), thus completing the proof.

We should point out that we have actually shown only that the components coming from the tubular neighborhood $\operatorname{An}_r(\gamma)$ are graphs. However, the other components are easily also seen to be graphs by combining the curvature estimate and embeddedness. Namely, any other component is intrinsically far from the boundary and hence graphical over some plane. By embeddedness, these graphs do not cross, and we can take these planes to be parallel. \Box

Appendix D. Blowing up intrinsically on the scale of nontrivial topology

The next lemma uses a standard blowup argument to locate the smallest scale of nontrivial topology:

LEMMA D.1. Suppose that $\Sigma \subset \mathbf{R}^3$ is a smooth minimal surface, possibly with boundary $\partial \Sigma$. If the ball $\mathcal{B}_{5C_1r_1}(y_0) \subset \Sigma$ is disjoint from $\partial \Sigma$ for some $C_1 > 1 \, and$

(D.2)
$$\mathcal{B}_{r_1}(y_0)$$
 is not a disk.

then there exists a sub-ball $\mathcal{B}_{C_1s}(y_1) \subset \mathcal{B}_{4C_1r_1}(y_0)$ so that

(D.3)
$$\mathcal{B}_{4s}(y_1)$$
 is not a disk,

(D.4)
$$\mathcal{B}_s(y)$$
 is a disk for any $y \in \mathcal{B}_{C_1s}(y_1)$.

Proof. After rescaling, we can assume that $r_1 = 1$. The lemma will follow from a simple rescaling argument as in Lemma 5.1 of [CM04c], except we define F intrinsically on $\mathcal{B}_{4C_1}(y_0)$ by

(D.5)
$$F(x) = d^2(x) i^{-2}(x),$$

where i(x) is the injectivity radius of Σ at x and

(D.6)
$$d(x) = 4C_1 - \text{dist}_{\Sigma}(x, y_0)$$

is the distance to $\partial \mathcal{B}_{4C_1}(y_0)$. It follows that F = 0 on $\partial \mathcal{B}_{4C_1}(y_0)$ and $F(y_0) \ge 16 C_1^2$. Also, since $\mathcal{B}_{5C_1}(y_0)$ is smooth, it follows that F is bounded from above on $\mathcal{B}_{4C_1}(y_0)$. We can therefore choose a point y_1 where $F(y_1)$ is at least half of its supremum,⁶² i.e.,

(D.7)
$$F(y_1) > 1/2 \sup_{\mathcal{B}_{4C_1}(y_0)} F.$$

Set $s^2 = i^2(y_1)/8$.

To see that (D.3) holds, first note that $4s > i(y_1)$. In particular, there must be two distinct geodesics, γ_1 and γ_2 , contained in $\mathcal{B}_{4s}(y_1)$ with

(D.8)
$$\gamma_1(0) = \gamma_2(0) = y_1 \text{ and } \gamma_1(i(y_1)) = \gamma_2(i(y_1)).$$

Since Γ has nonpositive curvature, it follows immediately from the Gauss-Bonnet theorem with corners that the closed curve $\gamma_1 \cup \gamma_2$ cannot bound a disk in Γ , thus giving (D.3).⁶³

 $^{^{62}}$ In the proof above, we are not using that *i*, or *F*, is continuous (and, in particular, not using that a maximum is attained). Instead, we use only that there must be points where *F* is at least half of its supremum. However, one can show continuity; see Proposition 88 in [Ber03].

⁶³The Gauss-Bonnet theorem with corners implies that a disk D has $\int_{\partial D} k_g + \int_D K + \sum \alpha_i = 2\pi$, where ∂D has jump angles α_i at the corners. In this case, both integrals are nonpositive and there are only two corners with each contributing *less* than π , so no such disk can exist.

We will use (D.7) twice to prove (D.4). First, since $d \ge d(y_1)/2$ on $\mathcal{B}_{\underline{d(y_1)}}(y_1)$, (D.7) implies that

(D.9)
$$\sup_{\mathcal{B}_{\frac{d(y_1)}{2}}(y_1)} i^{-2} \le \frac{4}{d^2(y_1)} \sup_{\mathcal{B}_{\frac{d(y_1)}{2}}(y_1)} F < \frac{8F(y_1)}{d^2(y_1)} = s^{-2},$$

so that i > s on $\mathcal{B}_{\frac{d(y_1)}{2}}(y_1)$. Second, using $F(y_0)$ as a lower bound for the sup of F in (D.7) implies that

(D.10)
$$8C_1^2 \le F(y_1) = \frac{d^2(y_1)}{8s^2}$$

so that $d(y_1)/2 > C_1 s$.

Appendix E. Minimal surfaces with a quadratic curvature bound in a half-space

The next lemma deals with a minimal surface whose curvature blows up at most quadratically at a point z in its closure. The lemma shows that the surface must come arbitrarily scale-invariant close to *any* plane through z. Roughly speaking, this means that the surface does not lie in any strictly mean convex cone through z; see Figure 52.

To state the lemma precisely, given a plane through z, we will define a scale-invariant function $\beta(s)$ that measures how close in the sphere $\partial B_s(z)$ a surface $\tilde{\Gamma} \subset \mathbf{R}^3 \setminus \{z\}$ comes to the plane. After a rotation, we can assume that the plane is the horizontal plane $\{x_3 = x_3(z)\}$. Define the function $\beta(s)$ by setting

(E.1)
$$\beta(s) = \frac{\inf_{\partial B_s(z) \cap \tilde{\Gamma}} |x_3 - x_3(z)|}{s}$$

The next lemma shows that the limit of $\beta(s)$ is zero, so that $\tilde{\Gamma}$ comes arbitrarily scale-invariant close to the plane as we approach z.

LEMMA E.2. Let $\tilde{\Gamma} \subset \mathbf{R}^3 \setminus \{z\}$ be a minimal surface (embedded or not) with z in its closure. Suppose that, for each $\varepsilon > 0$, each component of $\tilde{\Gamma} \setminus B_{\varepsilon}(z)$ is complete and has boundary in $\partial B_{\varepsilon}(z)$.

If there exist constants $r_0 > 0$ and C so that for $x \in B_{r_0}(z) \cap \tilde{\Gamma}$ we have

(E.3)
$$|A|^2(x) \le C |x-z|^{-2}$$

then the function $\beta(s)$ defined in (E.1) satisfies

(E.4)
$$\liminf_{s \to 0} \beta(s) = 0$$



Strictly convex cone with vertex z.

Figure 52. We prove Lemma E.2 by contradiction, so suppose that $\tilde{\Gamma}$ lies in a strictly mean convex cone.

Proof. We will prove (E.4) by contradiction (see Figure 52), so suppose that

(E.5)
$$\liminf_{s \to 0} \beta(s) = \beta_0 > 0.$$

In particular, given any $\delta > 0$, equation (E.5) implies that there exists $s_0 > 0$ so that $\beta(s) > \beta_0 - \delta$ for every $s < s_0$ and, hence, $B_{s_0}(z) \cap \tilde{\Gamma}$ lies inside a strictly mean convex (double) cone:

(E.6)
$$B_{s_0}(z) \cap \tilde{\Gamma} \subset \{ |x_3 - x_3(z)| > (\beta_0 - \delta) |x - z| \}.$$

On the other hand, (E.5) also implies that there is an $s < s_0/2$ with $\beta(s) < \beta_0 + \delta$ and, consequently, there is a point $y_s \in \partial B_s(z) \cap \tilde{\Gamma}$ close to the strictly mean convex cone:

(E.7)
$$y_s \in \partial B_s(z) \cap \{|x_3 - x_3(z)| < (\beta_0 + \delta) s\} \cap \tilde{\Gamma}.$$

Note that (E.6) and (E.7) imply that the intrinsic ball $\mathcal{B}_{s/2}(y_s)$ stays inside, but comes close to, the *strictly* mean convex cone

(E.8)
$$\{|x_3 - x_3(z)| = (\beta_0 - \delta) |x - z|\}.$$

We will assume below that $\delta < \beta_0$.

We will see that (E.6) and (E.7) lead to a contradiction for δ sufficiently small, thus proving (E.4).

First, recall that the quadratic curvature bound (E.3) gives an $\alpha > 0$ so that the component $\tilde{\Gamma}_{\alpha s}$ of $B_{\alpha s}(y_s) \cap \tilde{\Gamma}$ containing y_s is a graph with gradient bounded by one (see, e.g., Lemma 2.2 in [CM99]). After possibly reducing α , we can therefore assume that

(E.9)
$$\tilde{\Gamma}_{\alpha s} \subset \mathcal{B}_{s/2}(y_s).$$

Since $\Gamma_{\alpha s}$ is connected and does not intersect the (double) cone (E.8), it must be in one of the two components of $\{|x_3 - x_3(z)| > (\beta_0 - \delta) |x - z|\}$. After

possibly reflecting, we can assume that

(E.10)
$$\tilde{\Gamma}_{\alpha s} \subset \{x_3 - x_3(z) > (\beta_0 - \delta) |x - z|\}.$$

Define a function f that vanishes on the cone $\{x_3 - x_3(z) = (\beta_0 - \delta) |x - z|\}$ by setting

(E.11)
$$f(x) = x_3 - x_3(z) - (\beta_0 - \delta) |x - z|.$$

Note that (E.6) and (E.7) imply that

(E.12)
$$0 \le \inf_{\tilde{\Gamma}_{\alpha s}} f \le f(y_s) < 2 \,\delta \,s.$$

Using that $\tilde{\Gamma}_{\alpha s}$ is minimal, we have that

(E.13)
$$\Delta f = -(\beta_0 - \delta) \,\Delta |x - z| < -\frac{\beta_0 - \delta}{|x - z|}.$$

Define a function g on $\tilde{\Gamma}_{\alpha s}$ by setting

(E.14)
$$g = f + |x - y_s|^2 \frac{\beta_0 - \delta}{6s}.$$

Using that |x-z| < 3s/2 on $\tilde{\Gamma}_{\alpha s}$, we get that g is superharmonic since

(E.15)
$$\Delta g < -\frac{\beta_0 - \delta}{|x - z|} + 4 \frac{\beta_0 - \delta}{6s} < 0.$$

Therefore, the minimum of g is achieved on $\partial \tilde{\Gamma}_{\alpha s}$ and thus

(E.16)
$$\min_{\partial \tilde{\Gamma}_{\alpha s}} \left[f + (\alpha s)^2 \frac{\beta_0 - \delta}{6s} \right] = \min_{\tilde{\Gamma}_{\alpha s}} g < g(y_s) = f(y_s) < 2 \,\delta \,s,$$

where the last inequality is from (E.12). Combining the first inequality from (E.12) and (E.16) gives

(E.17)
$$0 \le \min_{\partial \tilde{\Gamma}_{\alpha s}} f < 2 \,\delta \, s - \alpha^2 \, s \, (\beta_0 - \delta)/6.$$

This gives the desired contradiction for δ sufficiently small.

References

- [Ber03] M. BERGER, A Panoramic View of Riemannian Geometry, Springer-Verlag, New York, 2003. MR 2002701. Zbl 1038.53002. http://dx.doi. org/10.1007/978-3-642-18245-7.
- [CM99] T. H. COLDING and W. P. MINICOZZI, II, Minimal surfaces, Courant Lecture Notes in Math. 4, New York University, Courant Institute of Mathematical Sciences, New York, 1999. MR 1683966. Zbl 0987.49025.
- [CM01] T. H. COLDING and W. P. MINICOZZI, II, Minimal annuli with and without slits, J. Symplectic Geom. 1 (2001), 47–61. MR 1959578. Zbl 1087.
 53008. Available at http://projecteuclid.org/euclid.jsg/1092316298.

TOBIAS H. COLDING and WILLIAM P. MINICOZZI II

- [CM02a] T. H. COLDING and W. P. MINICOZZI, II, Estimates for parametric elliptic integrands, *Int. Math. Res. Not.* 2002, no. 6, 291–297. MR 1877004.
 Zbl 1002.53035. http://dx.doi.org/10.1155/S1073792802106106.
- [CM02b] T. H. COLDING and W. P. MINICOZZI, II, Multi-valued minimal graphs and properness of disks, *Internat. Math. Res. Notes* 2002, 1111–1127. MR 1904463. Zbl 1008.58012. http://dx.doi.org/10.1155/ S107379280211107X.
- [CM02c] T. H. COLDING and W. P. MINICOZZI, II, On the structure of embedded minimal annuli, *Int. Math. Res. Not.* 2002, no. 29, 1539–1552. MR 1907204. Zbl 1122.53300. http://dx.doi.org/10.1155/ S1073792802112128.
- [CM03] T. H. COLDING and W. P. MINICOZZI, II, Disks that are double spiral staircases, Notices Amer. Math. Soc. 50 (2003), 327–339. MR 1954009. Zbl 1100.53008.
- [CM04a] T. H. COLDING and W. P. MINICOZZI, II, Embedded minimal disks: Proper versus nonproper—global versus local, *Trans. Amer. Math. Soc.* **356** (2004), 283–289. MR 2020033. Zbl 1046.53001. http://dx.doi.org/ 10.1090/S0002-9947-03-03230-6.
- [CM04b] T. H. COLDING and W. P. MINICOZZI, II, The space of embedded minimal surfaces of fixed genus in a 3-manifold. I. Estimates off the axis for disks, Ann. of Math. 160 (2004), 27–68. MR 2119717. Zbl 1070.53031. http://dx.doi.org/10.4007/annals.2004.160.27.
- [CM04c] T. H. COLDING and W. P. MINICOZZI, II, The space of embedded minimal surfaces of fixed genus in a 3-manifold. II. Multi-valued graphs in disks, Ann. of Math. 160 (2004), 69–92. MR 2119718. Zbl 1070.53032. http://dx.doi.org/10.4007/annals.2004.160.69.
- [CM04d] T. H. COLDING and W. P. MINICOZZI, II, The space of embedded minimal surfaces of fixed genus in a 3-manifold. III. Planar domains, Ann. of Math. 160 (2004), 523–572. MR 2123932. Zbl 1076.53068. http://dx.doi.org/10.4007/annals.2004.160.523.
- [CM04e] T. H. COLDING and W. P. MINICOZZI, II, The space of embedded minimal surfaces of fixed genus in a 3-manifold. IV. Locally simply connected, Ann. of Math. 160 (2004), 573–615. MR 2123933. Zbl 1076.53069. http://dx.doi.org/10.4007/annals.2004.160.573.
- [CM08] T. H. COLDING and W. P. MINICOZZI, II, The Calabi-Yau conjectures for embedded surfaces, Ann. of Math. 167 (2008), 211–243. MR 2373154.
 Zbl 1142.53012. http://dx.doi.org/10.4007/annals.2008.167.211.
- [dCP79] M. DO CARMO and C. K. PENG, Stable complete minimal surfaces in **R**³ are planes, *Bull. Amer. Math. Soc.* **1** (1979), 903–906. MR 0546314. Zbl 0442.53013. http://dx.doi.org/10.1090/S0273-0979-1979-14689-5.
- [Fia41] F. FIALA, Le problème des isopérimètres sur les surfaces ouvertes à courbure positive, Comment. Math. Helv. 13 (1941), 293–346. MR 0006422.
 Zbl 0025.23003. http://dx.doi.org/10.1007/BF01378068.

- [FCS80] D. FISCHER-COLBRIE and R. SCHOEN, The structure of complete stable minimal surfaces in 3-manifolds of nonnegative scalar curvature, *Comm. Pure Appl. Math.* 33 (1980), 199–211. MR 0562550. Zbl 0439.53060. http: //dx.doi.org/10.1002/cpa.3160330206.
- [Gro81] M. GROMOV, Groups of polynomial growth and expanding maps, Inst. Hautes Études Sci. Publ. Math. 53 (1981), 53-73. MR 0623534. Zbl 0474.
 20018. http://dx.doi.org/10.1007/BF02698687.
- [GL86] R. GULLIVER and H. B. LAWSON, JR., The structure of stable minimal hypersurfaces near a singularity, in *Geometric Measure Theory and the Calculus of Variations* (Arcata, Calif., 1984), *Proc. Sympos. Pure Math.* 44, Amer. Math. Soc., Providence, RI, 1986, pp. 213–237. MR 0840275. Zbl 0592.53005. http://dx.doi.org/10.1090/pspum/044/840275.
- [HS79] R. HARDT and L. SIMON, Boundary regularity and embedded solutions for the oriented Plateau problem, Ann. of Math. 110 (1979), 439–486.
 MR 0554379. Zbl 0457.49029. http://dx.doi.org/10.2307/1971233.
- [Har64] P. HARTMAN, Geodesic parallel coordinates in the large, Amer. J. Math.
 86 (1964), 705–727. MR 0173222. Zbl 0128.16105. http://dx.doi.org/10.
 2307/2373154.
- [Hil69] S. HILDEBRANDT, Boundary behavior of minimal surfaces, Arch. Rational Mech. Anal. 35 (1969), 47–82. MR 0248650. Zbl 0183.39402. http://dx. doi.org/10.1007/BF00248494.
- [HM90] D. HOFFMAN and W. H. MEEKS, III, The strong halfspace theorem for minimal surfaces, *Invent. Math.* 101 (1990), 373–377. MR 1062966.
 Zbl 0722.53054. http://dx.doi.org/10.1007/BF01231506.
- [HW11] D. HOFFMAN and B. WHITE, Sequences of embedded minimal disks whose curvatures blow up on a prescribed subset of a line, *Comm. Anal. Geom.* 19 (2011), 487–502. MR 2843239. Zbl 1244.53010. http://dx.doi.org/10. 4310/CAG.2011.v19.n3.a2.
- [Kle12] S. J. KLEENE, A minimal lamination with Cantor set-like singularities, *Proc. Amer. Math. Soc.* 140 (2012), 1423–1436. MR 2869127. Zbl 1239.
 53006. http://dx.doi.org/10.1090/S0002-9939-2011-10971-7.
- [LR88] R. LANGEVIN and H. ROSENBERG, A maximum principle at infinity for minimal surfaces and applications, *Duke Math. J.* 57 (1988), 819–828. MR 0975123. Zbl 0667.49024. http://dx.doi.org/10.1215/ S0012-7094-88-05736-5.
- [Mee04] W. H. MEEKS, III, Regularity of the singular set in the Colding-Minicozzi lamination theorem, *Duke Math. J.* 123 (2004), 329–334. MR 2066941.
 Zbl 1086.53005. http://dx.doi.org/10.1215/S0012-7094-04-12324-3.
- [MP04] W. H. MEEKS, III and J. PÉREZ, Conformal properties in classical minimal surface theory, in Surveys in Differential Geometry. Vol. IX, Surv. Differ. Geom. IX, Int. Press, Somerville, MA, 2004, pp. 275–335.
 MR 2195411. Zbl 1086.53007. http://dx.doi.org/10.4310/SDG.2004.v9. n1.a8.

TOBIAS H. COLDING and WILLIAM P. MINICOZZI II

- [MPR04a] W. H. MEEKS, III, J. PÉREZ, and A. ROS, The geometry of minimal surfaces of finite genus. I. Curvature estimates and quasiperiodicity, J. Differential Geom. 66 (2004), 1–45. MR 2128712. Zbl 1068.53012. Available at http://projecteuclid.org/euclid.jdg/1090415028.
- [MPR04b] W. H. MEEKS, III, J. PÉREZ, and A. Ros, The geometry of minimal surfaces of finite genus. II. Nonexistence of one limit end examples, *Invent. Math.* 158 (2004), 323–341. MR 2096796. Zbl 1070.53003. http://dx.doi. org/10.1007/s00222-004-0374-3.
- [MPR06] W. H. MEEKS, III, J. PÉREZ, and A. Ros, Liouville-type properties for embedded minimal surfaces, *Comm. Anal. Geom.* 14 (2006), 703–723. MR 2273291. Zbl 1117.53009. http://dx.doi.org/10.4310/CAG.2006.v14. n4.a5.
- [MPR10] W. H. MEEKS, III, J. PÉREZ, and A. ROS, Limit leaves of an H lamination are stable, J. Differential Geom. 84 (2010), 179–189. MR 2629513. Zbl 05724752. Available at http://projecteuclid.org/euclid. jdg/1271271797.
- [MPR] W. H. MEEKS, III, J. PÉREZ, and A. Ros, The geometry of minimal surfaces of finite genus III; bounds on the topology and index of classical minimal surfaces, preprint.
- [MPR15] W. H. MEEKS, III, J. PÉREZ, and A. Ros, Properly embedded minimal planar domains, to appear in *Ann. of Math.* **181** (2015). arXiv **1306.1690**.
- [MR05] W. H. MEEKS, III and H. ROSENBERG, The uniqueness of the helicoid, Ann. of Math. 161 (2005), 727–758. MR 2153399. Zbl 1102.53005. http: //dx.doi.org/10.4007/annals.2005.161.727.
- [MR06] W. H. MEEKS, III and H. ROSENBERG, The minimal lamination closure theorem, *Duke Math. J.* 133 (2006), 467–497. MR 2228460. Zbl 1098.
 53007. http://dx.doi.org/10.1215/S0012-7094-06-13332-X.
- [MY82a] W. H. MEEKS, III and S. T. YAU, The classical Plateau problem and the topology of three-dimensional manifolds. The embedding of the solution given by Douglas-Morrey and an analytic proof of Dehn's lemma, *Topology* 21 (1982), 409–442. MR 0670745. Zbl 0489.57002. http://dx.doi.org/10. 1016/0040-9383(82)90021-0.
- [MY82b] W. W. MEEKS, III and S. T. YAU, The existence of embedded minimal surfaces and the problem of uniqueness, *Math. Z.* **179** (1982), 151–168. MR 0645492. Zbl 0479.49026. http://dx.doi.org/10.1007/BF01214308.
- [NO61] K. NOMIZU and H. OZEKI, The existence of complete Riemannian metrics, Proc. Amer. Math. Soc. 12 (1961), 889–891. MR 0133785. Zbl 0102.16401. http://dx.doi.org/10.2307/2034383.
- [P05] J. PÉREZ, Limits by rescalings of minimal surfaces: Minimal laminations, curvature decay and local pictures, notes for the workshop "Moduli Spaces of Properly Embedded Minimal Surfaces", American Institute of Mathematics, Palo Alto, California (2005).
- [Ros03] H. ROSENBERG, Some recent developments in the theory of minimal surfaces in 3-manifolds, in 24° Colóquio Brasileiro de Matemática. [24th

Brazilian Mathematics Colloquium], Publicações Matemáticas do IMPA. [IMPA Mathematical Publications], Instituto de Matemática Pura e Aplicada (IMPA), Rio de Janeiro, 2003, pp. iv+48. MR 2028922. Zbl 1064. 53007.

- [Sch83a] R. SCHOEN, Estimates for stable minimal surfaces in three-dimensional manifolds, in *Seminar on Minimal Submanifolds*, Ann. of Math. Stud. 103, Princeton Univ. Press, Princeton, NJ, 1983, pp. 111–126. MR 0795231. Zbl 0532.53042.
- [Sch83b] R. M. SCHOEN, Uniqueness, symmetry, and embeddedness of minimal surfaces, J. Differential Geom. 18 (1983), 791-809 (1984).
 MR 0730928. Zbl 0575.53037. Available at http://projecteuclid.org/ euclid.jdg/1214438183.
- [ST89] K. SHIOHAMA and M. TANAKA, An isoperimetric problem for infinitely connected complete open surfaces, in *Geometry of Manifolds* (Matsumoto, 1988), *Perspect. Math.* 8, Academic Press, Boston, 1989, pp. 317–343. MR 1040533. Zbl 0697.53040.
- [ST93] K. SHIOHAMA and M. TANAKA, The length function of geodesic parallel circles, in *Progress in Differential Geometry*, Adv. Stud. Pure Math. 22, Math. Soc. Japan, Tokyo, 1993, pp. 299–308. MR 1274955. Zbl 0799. 53052.
- [Sol86] B. SOLOMON, On foliations of Rⁿ⁺¹ by minimal hypersurfaces, Comment. Math. Helv. 61 (1986), 67–83. MR 0847521. Zbl 0601.53025. http://dx. doi.org/10.1007/BF02621903.
- [WHW09] M. WEBER, D. HOFFMAN, and M. WOLF, An embedded genus-one helicoid, Ann. of Math. 169 (2009), 347–448. MR 2480608. Zbl 1213.49049. http://dx.doi.org/10.4007/annals.2009.169.347.

(Received: December 12, 2005) (Revised: May 1, 2014)

COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK, NY Current address: MIT, CAMBRIDGE, MA E-mail: colding@math.mit.edu

JOHNS HOPKINS UNIVERSITY, BALTIMORE, MD Current address: MIT, CAMBRIDGE, MA E-mail: minicozz@math.mit.edu