THE DOUBLE BUBBLE PROBLEM ON THE FLAT TWO-TORUS

JOSEPH CORNELI, PAUL HOLT, GEORGE LEE, NICHOLAS LEGER,
ERIC SCHOENFELD, AND BENJAMIN STEINHURST

ABSTRACT. We characterize the perimeter-minimizing double bubbles on all flat two-tori and, as corollaries, on the flat infinite cylinder and the flat infinite strip with free boundary. Specifically, we show that there are five distinct types of minimizers on flat two-tori, depending on the areas to be enclosed.

1. INTRODUCTION

Our Main Theorem 6.1 shows that on any flat two-torus, the least-perimeter way to enclose and separate two prescribed areas is a double bubble of one of the five types shown in Figure 1: the standard double bubble, the band lens, a standard chain, the double band, or the standard hexagon tiling. Figure 2 gives "phase portraits," computational plots of the minimizing type for given areas, for four different tori. (Section 8 describes the creation of these phase portraits, using formulas for the perimeter and area of the five minimizers derived in Section 7.)

The strategy of the proof is to separate all candidates into five classes, as described in Proposition 4.6. Proposition 4.6 says that a minimizer must either be the double band or (possibly after relabelling the two interior regions and the exterior) belong to one of four other topological classes: Section 5.1 deals with contractible double bubbles; Section 5.2 deals with double bubbles whose components taken together wrap around the torus, and whose complement is not contractible; Section 5.3 deals with double bubbles with one band adjacent to a contractible set of components; and Section 5.4 deals with double bubbles for which both enclosed regions and the exterior region are contractible.

The last of these topological classes (with both regions and the exterior contractible) we refer to as "tilings" because the three regions taken together lift to a tiling of the plane. Despite escaping our initial conjectures, this class turned out to be one of the most interesting. Whereas the nontiling types tie only for prescribed areas in the transitions between phases (Figure 2), the hexagon tiling and the double band provide distinct minimizers over an open set of areas (Figure 2(b)) — the first example of such a phenomenon in double bubble history. Another candidate tiling, the octagon square, with one region consisting of two four-sided components (Figure 23), was the last candidate to be eliminated (Proposition 5.4.22).
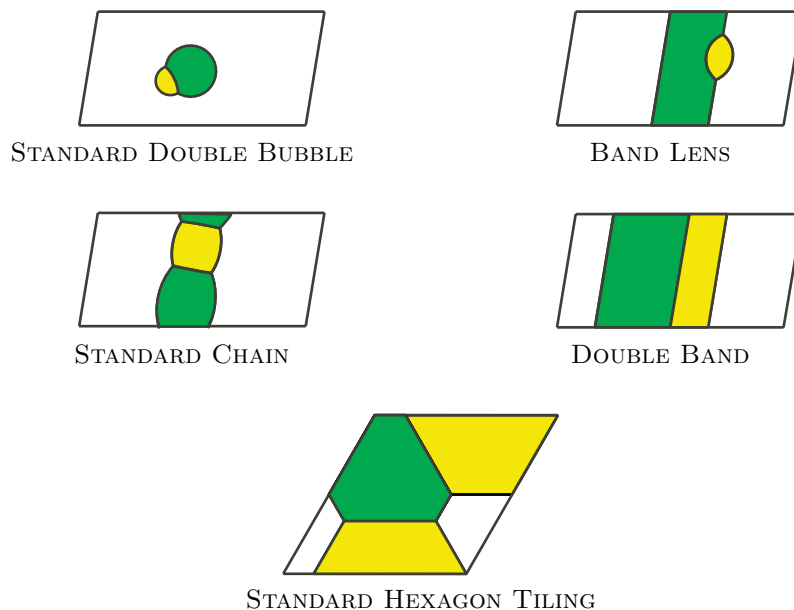
FIGURE 1. There are five types of perimeter-minimizing double bubbles on a flat two-torus. The standard hexagon tiling occurs only on the hexagonal torus.

We also consider the double bubble problem on the flat infinite cylinder, which is simpler because there are no candidate tilings. Corollary 6.3 provides a complete characterization of the minimizers in this space as a trivial consequence of the torus result, and the cylinder result is used in Corollary 6.5 to characterize all minimizers on the flat infinite strip with free boundary.

The torus and cylinder lack the useful symmetries used in other spaces to establish connectedness of the exterior and bounds on the number of components. Luckily, the recent work of Wichiramala [MW] on the triple bubble problem in the plane gave us a very useful component bound (Proposition 3.7).

History and recent developments. The double bubble problem is a generalization of the isoperimetric problem. Given a geometric space, the latter seeks the least-area way to enclose a single volume. (A solution to the isoperimetric problem on the flat two-torus can be found in [HHM].) The double bubble problem seeks the least-area way to enclose and separate two volumes. In recent years, this problem has been solved in a number of spaces, including $\mathbf{R}^2$ [F], $\mathbf{S}^2$ [Ma2], and still more recently $\mathbf{R}^3$ [HMRR] and $\mathbf{R}^4$ [RHLS]. In these cases the minimizers are all standard double bubbles, in contrast to the multiplicity of types on the torus. Some of the major contributions were due to undergraduates affiliated with Williams College and the Williams "SMALL" REU program. Underlying all of these results is geometric measure theory as developed by Federer, Fleming, Almgren, and others. For a more complete discussion of the history of the double bubble problem as well as many related problems on films, foams, and other efficient surfaces, see [M1].

The triple bubble problem in the plane has recently been solved by Wichiramala [W]. On the other hand, even the single bubble problem remains unsolved on a torus
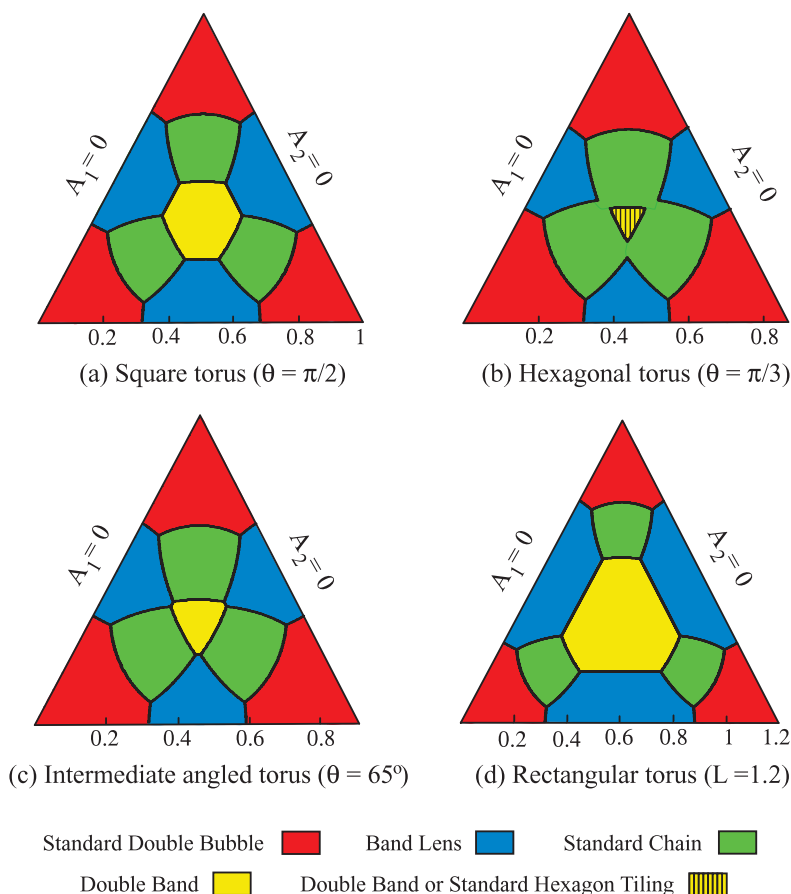
FIGURE 2. These phase portraits show which type of double bubble is minimizing given prescribed area pairs (A_1, A_2) on four specific flat tori. Each torus for (a)–(c) has a rhombic fundamental domain with interior angle θ; the torus for (d) has a rectangular fundamental domain with side lengths 1 and $L = 1.2$.

of revolution. We have contributed to a paper that gives experimental evidence on the solutions to the double bubble problem in a cubic flat three-torus [CCWB], another space where the single bubble problem is still open.

2. DEFINITIONS

This section contains definitions for, among other things, the classes of potential minimizers mentioned above.

2.1. Tori and the cylinder.

Definition 1. A *flat two-dimensional torus* can be represented by a planar parallelogram with opposite sides identified. Throughout the paper the torus is normalized so that each shortest closed geodesic has unit length.

Remark 2.1.1. The parallelogram can be chosen such that one shortest side is a shortest closed geodesic and no interior angle is less than sixty degrees. Then the area of the parallelogram, and hence the area of the torus, is at least $\sqrt{3}/2$ with equality precisely for the hexagonal torus (see Definition 2).

Definition 2. The *hexagonal torus* can be represented either by a regular hexagon with opposite sides identified with the same orientation, or by a 60-degree rhombus with opposite sides identified with the same orientation.

Definition 3. We say that each shortest closed geodesic lies along a *short direction* of the torus. On certain nonhexagonal tori (those that can be represented by a parallelogram with four sides of unit length), there are two short directions; on the hexagonal torus, there are three.

Definition 4. The *flat infinite cylinder* can be represented as the surface contained between two identified infinite parallel lines in the plane, where each segment perpendicular to and connecting the lines is a closed geodesic. The *flat infinite strip with free boundary* is the surface between the same two lines, except that the lines are not identified; portions of the lines may be included in the boundary of a double bubble without contributing to the perimeter of the double bubble.

2.2. **Topological definitions.**

Definition 5. A *double bubble* on a smooth Riemannian surface consists of two disjoint regions (i.e., nonempty open sets) bounded by piecewise smooth curves, such that the exterior of these regions has nonzero area. The boundary of the double bubble divides the surface into three regions: the two regions belonging to the double bubble, which we call the *enclosed* regions, and the complement of the closure of the double bubble, which we call the *exterior* region. Any of the three regions may have multiple components.

Remark 2.2.1. On the torus, by relabelling which of the three regions in a double bubble is the exterior, we can obtain two other double bubbles with the same boundary. On the infinite cylinder, the exterior region has infinite area.

Definition 6. A *tiling* is a double bubble on the torus for which each component, including each component of the exterior, is contractible, such as the standard hexagon tiling of Figure 1. It is called a tiling because the two enclosed regions plus the exterior region lift to a tiling of the plane.

Definition 7. An *octagon-square tiling* is a tiling in which one of the three regions consists of two curvilinear quadrilaterals, and each of the other two regions consists of one curvilinear octagon (see Figure 23).

Definition 8. A *band* is a noncontractible annulus.

Definition 9. A *swath* is a set of contractible components with noncontractible complement and whose union has noncontractible and connected closure, such as the standard chain of Figure 1.

Definition 10. A *chain*, such as the standard chain of Figure 1 or the chain of Figure 3, is a minimal swath — that is, a swath without a proper subset of components that themselves form a swath. Note that every swath contains at least one chain and that every chain is itself a swath.
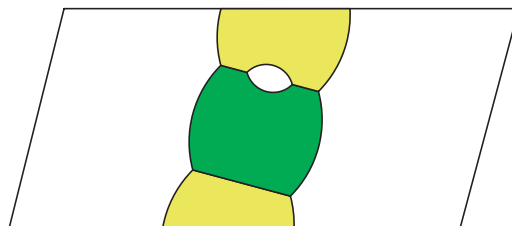

FIGURE 3. The two shaded six-sided components form a *chain*, even though there is a curvilinear digon embedded between them. The two shaded components plus the curvilinear digon do not form a chain, because the swath they form is not minimal.

Definition 11. Two closed curves lie in the same *homology class*, or have the same homology, if they can be smoothly deformed into each other along the surface of the torus. Given two different oriented directions of the torus (i.e., directions of closed geodesics on the torus), the homology class of a closed curve is given by an ordered pair (p, q), where the curve wraps around the torus p times in the first direction and q times in the other direction. (The ordered pair is well defined up to negating both p and q.) We will refer to the homology of a band or chain K, by which we mean the homology of any noncontractible component of the boundary of $\overline{K}$.

2.3. Geometric definitions.

Definition 12. The *standard double bubble* is comprised of three circular arcs meeting in threes at $120°$, such that the curvature of the arc separating the contractible components is the difference of the curvatures of the outer caps. For prescribed areas, there is a standard double bubble in the Euclidean plane, unique up to congruence [M1, Proposition 14.1]; it may or may not fit on the torus.

Definition 13. The *double band* consists of two adjacent bands, bounded by three shortest closed geodesics.

Definition 14. A *symmetric chain* is a chain that is symmetric about a closed geodesic with the same homology as the chain.

Definition 15. A *standard chain* is a symmetric chain enclosing two four-sided components whose boundary is comprised of six circular arcs that satisfy the regularity conditions in Proposition 3.1. That is, the arcs meet in threes at four common vertices at $120°$ such that the curvature of each arc separating the two components in the chain is the difference of the curvatures of the outer arcs (see Figure 1). The *axis-length* of a standard chain is the length of the geodesic with the same homology as the chain.

Remark 2.3.1. Our parameterization of the standard chain (Section 7.2, see Remark 7.2.3) proves existence of some family of standard chains. Furthermore, there exists

graphical evidence that the only chains that are minimizing are those with the same homology as a shortest closed geodesic (Remark 5.2.16).

We believe the standard chain is also unique up to isometries of the torus for given areas and homology class. Indeed, a parametric plot in *Mathematica* (see Figure 4) suggests that different standard chains with the same homology enclose different area pairs.
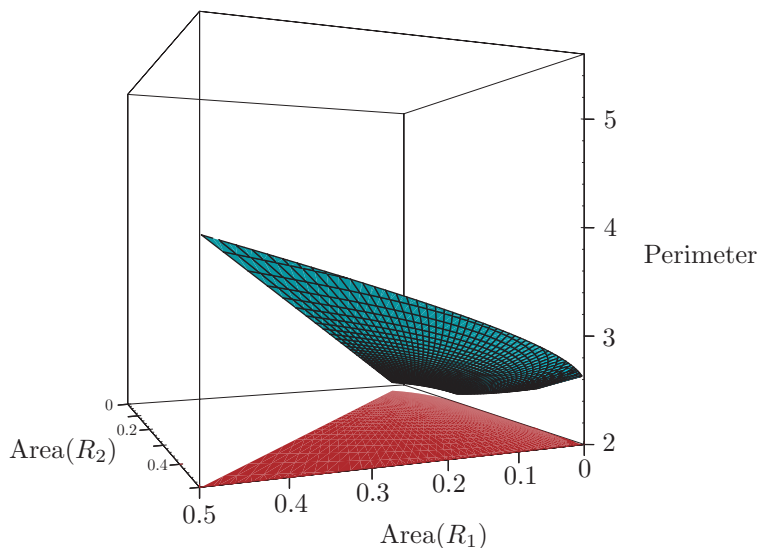

FIGURE 4. This parametric plot from *Mathematica* of perimeter as a function of areas for standard chains with axis-length 1. The plot has no folds in it and contains only one sheet, suggesting uniqueness for prescribed area pairs $A_1 \leq A_2$.

Definition 16. A *lens* is a pair of congruent circular arcs meeting at two vertices at $120°$.

Remark 2.3.2. It is easily shown that if the distance between the two vertices is less than or equal to 1, then the lens can be embedded on any torus with its axis along any short direction.

Definition 17. A *band lens* consists of a band and a lens, such that the boundary curves meet in threes at two vertices at $120°$, as in Figure 1.

Remark 2.3.3. By Lemma 5.3.6, a minimizing band lens on the torus has the same homology as a shortest closed geodesic (that is, the closed geodesic bordering the band has length one). It is easy to see that when a minimizing band lens enclosing two given areas exists, it is essentially unique for those areas (up to isometries of the torus).

Definition 18. A *hexagon tiling* is a tiling in which every component of the double bubble and exterior is a hexagon with interior angles of $120°$.

Definition 19. The *standard hexagon tiling* is a hexagon tiling on the hexagonal torus, which partitions the torus into three hexagonal regions, whose boundary edges lie parallel to the three short directions of the torus. By Corollary 7.4.3 and Lemma 7.4.4, the standard hexagon tiling exists and is unique for a particular range of area pairs.

3. Existence, regularity, and basic properties

This section provides existence and regularity of perimeter-minimizing double bubbles (Proposition 3.1), a bound on the number of components for minimizers (Proposition 3.7), and an easy but important perimeter bound (Proposition 3.8).

Proposition 3.1 (Existence and Regularity Theorem [M2, 2.3 and 2.4])**.** *In a smooth Riemannian surface S with compact quotient by its isometry group, for any two areas A_1 and A_2 (whose sum is less than the area of S), there exists a least-perimeter enclosure of the two areas. This enclosure consists of finitely many smooth constant-curvature curves meeting in threes at 120° angles. No boundary curve separates components of the same region, and all curves separating a specific pair of the three regions have the same (signed) curvature. Moreover, the enclosure satisfies the cocycle condition: the sum of the signed curvatures around any closed path is zero.*

Remark 3.2. In particular, given any minimizing double bubble on the torus, each of the three regions contains finitely many components.

Remark 3.3. It follows that in a minimizing double bubble on the torus, the union of the closures of any two regions is connected. Otherwise, two disjoint pieces of this union can be translated to be made tangent, forming a new minimizer that violates regularity, a contradiction. In particular, each component (in either of the two enclosed regions or the exterior region) is adjacent to components of both other regions.

Remark 3.4. It also follows that alternating sides of a component must border different regions. Hence, any contractible component in a minimizing double bubble or its exterior has an even number of sides.

Remark 3.5. A perimeter-minimizing double bubble is in *equilibrium*: under any smooth variation that preserves areas, the derivative of perimeter with respect to time is initially zero. (Here by "variation" we mean a perturbation of the double bubble over time.) For a piecewise smooth double bubble, being in equilibrium is equivalent to having boundary curves that satisfy the given curvature conditions, satisfy the cocycle condition, and meet in threes at 120° (see, for instance, [HMRR], Lemma 3.1).

Remark 3.6. The curvature and cocycle conditions imply the existence of a well-defined *pressure* for each region with the following property: the difference between the pressures of any two regions in an equilibrium double bubble is the curvature of the separating arcs between those regions, signed so that a curve separating two regions bows into the region of lower pressure. (For instance, we could set the pressure of the exterior region to equal zero and set the pressure of each enclosed region to equal the curvature of the interface separating that region from the exterior.) A least-pressure nonpolygonal contractible component must have more than six sides;

and a highest-pressure nonpolygonal contractible component must have fewer than six sides.

Proposition 3.7 (Wichiramala component bound [MW, Proposition 3.1]). *Unless all three regions have equal pressure, each region of highest pressure in a perimeter-minimizing double bubble in the flat two-torus has at most two contractible components.*

Proof. The proof in [MW], which we sketch here, uses variations that shrink or expand a high-pressure nonpolygonal component at a constant rate along its boundary. If there are at least three such components in some region, then some nontrivial linear combination of the corresponding "shrink and expand" variations will preserve the areas of each region to first order. Each of the three variations will contribute a negative term to the second variation of the combined variation, so that the new variation has negative second variation (i.e., negative second derivative of perimeter with respect to time). In this case, we call the double bubble "unstable," and no such double bubble is perimeter minimizing. The topology of the torus does not affect this proof. □

Proposition 3.8 (Perimeter bound). *The perimeter of a perimeter-minimizing double bubble in the flat two-torus is at most* 3.

Proof. A double band of perimeter three can enclose any pair of prescribed areas, so no minimizer can have greater perimeter. □

4. Topological classification of minimizers

This section proves two important propositions. Proposition 4.5 says that if the three regions in a perimeter-minimizing double bubble all have equal pressure, then it is a hexagon tiling or the double band. Proposition 4.6 is the backbone for the proof of our Main Theorem 6.1. It shows that any minimizing double bubble must be a tiling, a contractible double bubble, a swath, a single band adjacent to a contractible set of components, or the double band. Section 5 will treat in turn each of these topological classes.

Lemma 4.1. *A noncontractible component C in a minimizing double bubble must either be a band or have contractible complement.*

Proof. If every component of the boundary of C is contractible, then the complement of C is contractible.

Otherwise, some component σ of the boundary of C is noncontractible, and any other noncontractible component of the boundary of C has the same homology as σ. Thus, C must be a band with that homology, except possibly with some collection of contractible components in its interior. But such contractible components never exist in a minimizing double bubble (Remark 3.3), so C must indeed be a band. □

Lemma 4.2. *Consider a fundamental domain of the torus, a parallelogram with sides $1, L$ and interior angle $\theta \in [\pi/3, \pi/2]$. Suppose that a chain or band is part of a minimizing double bubble and has homology (p, q), defined with respect to two directions along the sides of the parallelogram. Then $|p|, |q| \leq 1$.*

Proof. Each band or chain has least two disjoint, noncontractible boundary curves. If $|p| \geq 2$ or $|q| \geq 2$, then the length of each of these boundary curves is at least

$\sqrt{3}$, and the total perimeter is at least $2\sqrt{3}$. By the perimeter bound (Proposition 3.8), the double bubble cannot be a minimizer. $\square$

Lemma 4.3. *A perimeter-minimizing double bubble on the torus, together with its exterior, has at most three disjoint bands or chains. If there are three, then the double bubble must be the double band.*

Proof. Each band or chain has at least two noncontractible boundary curves, and each such boundary curve borders at most two bands or chains. Hence, if there are n bands and chains, then there will be at least n such noncontractible boundary curves, each with length at least one. Then by Proposition 3.8, $n \leq 3$. If $n = 3$, then the boundary must consist entirely of three parallel closed geodesics, each of length one. In this case, the double bubble must be the double band. $\square$

Lemma 4.4. *In a perimeter-minimizing double bubble with three equal pressures, every component in the double bubble plus exterior is either a vertex-free band or a hexagon.*

Proof. Consider any component in the double bubble plus exterior, and let $\mathcal{C}$ be the boundary of this component.

Because the double bubble has three equal pressures, any vertex-free component of $\mathcal{C}$ is linear and thus must be a closed geodesic.

Suppose that a component of $\mathcal{C}$ has at least one vertex. Then it has positive total turning angle, whereas any noncontractible closed curve has total turning angle 0. Hence, if a component of $\mathcal{C}$ has at least one vertex, then it is contractible. Also, it consists of straight edges meeting at interior angles measuring $2\pi/3$ — precisely the configuration of a hexagon.

Therefore, every component of $\mathcal{C}$ is a closed geodesic or a contractible hexagon, and the desired result follows easily. $\square$

Proposition 4.5. *A perimeter-minimizing double bubble with three equal pressures is either a hexagon tiling or the double band.*

Proof. By Lemma 4.4, each component in the double bubble plus exterior is a vertex-free band or a hexagon. If every component is a hexagon, then the double bubble is a hexagon tiling. If every component is a band, then because there are at least three components, by Lemma 4.3 the double bubble is a double band. Finally, the double bubble cannot contain both bands and hexagons, because no band can be adjacent to a hexagon without violating regularity (Proposition 3.1). $\square$

Proposition 4.6. *A minimizing double bubble on the torus must be a contractible double bubble, a tiling, a swath, a band adjacent to a set of components such that the closure of their union is contractible, or the double band (after perhaps relabelling the two regions and the exterior).*

Proof. We will categorize the minimizing double bubble by the number and type of noncontractible components in the double bubble plus its exterior. By Lemma 4.1 any noncontractible component must either be a band or have contractible complement. If the double bubble plus its exterior has in any region (say, in the exterior) a component with contractible complement, then the double bubble is contractible. Thus, for our remaining cases we can assume that all noncontractible components are bands.

If the double bubble plus its exterior has no bands, then every component (of the double bubble plus its exterior) is contractible, and by definition the double bubble is a tiling.

If the double bubble plus its exterior has exactly one band, say in the exterior, then the double bubble consists of contractible components. But the closure of the union of this set of components is noncontractible, because it shares a non-contractible boundary curve with the band. Furthermore, any minimizing double bubble is connected (Remark 3.3). Then by definition the double bubble is a swath.

If the double bubble plus its exterior has exactly two bands (say, with at least one band in the exterior), then consider the set C of the contractible components — that is, the set of components different from the bands. Suppose some subset of these components has noncontractible and connected closure. Then C contains a swath, and therefore a chain. Now the double bubble plus its exterior has two bands and at least one chain, so by Lemma 4.3 it is not perimeter minimizing. Therefore, C must be a contractible set of components. The two bands must then be adjacent to each other, and hence lie in different regions. Then the double bubble contains exactly one band, together with a subset of C — that is, the double bubble is a band adjacent to contractible set of components.

If the double bubble plus its exterior has three or more bands, then by Lemma 4.3, the double bubble must be the double band. □

5. The four classes with contractible components

By Proposition 4.6, upon relabelling, every minimizing double bubble on the torus must be a contractible double bubble, a tiling, a swath, a band adjacent to a set of components whose union is contractible, or the double band. While the double band describes a specific geometric configuration, each of the remaining possibilities is merely a topological class of configurations. This section identifies the potential minimizer in each of these four remaining topological classes: the standard double bubble among contractible double bubbles (Proposition 5.1.2), a standard chain among swaths (Proposition 5.2.15), the band lens among double bubbles with a single band adjacent to a contractible set of components (Proposition 5.3.5), and the standard hexagon tiling among tilings (Proposition 5.4.23).

5.1. **Contractible double bubbles.** Proposition 5.1.2 uses the result from $\mathbf{R}^2$ that perimeter-minimizing double bubbles are standard [F], to show that the only potential minimizer among contractible double bubbles is the standard double bubble.

Proposition 5.1.1. *The perimeter of any planar standard double bubble is greater than π times its diameter: $P > \pi D$.*

Proof. We obtain a formula for P/D as a function $f(\theta)$, where $\theta \in [0, \pi/3)$ is the angle between the interior arc of the standard double bubble and the chord that connects the endpoints of the arc (see Figure 5). Then we show $f(\theta) > \pi$ for the entire domain $0 < \theta < \pi/3$. (A separate computation easily verifies the inequality $P > \pi D$ when $\theta = 0$.)

Let r_1, r_2, c_1, c_2 be the radii and centers of the longer and shorter exterior arcs of the standard double bubble, respectively. Then the diameter is the sum of r_1, r_2, and the distance between c_1 and c_2:

$$D = r_1 + r_2 + r_1 \cos(\pi/3 - \theta) + r_2 \cos(\pi/3 + \theta).$$
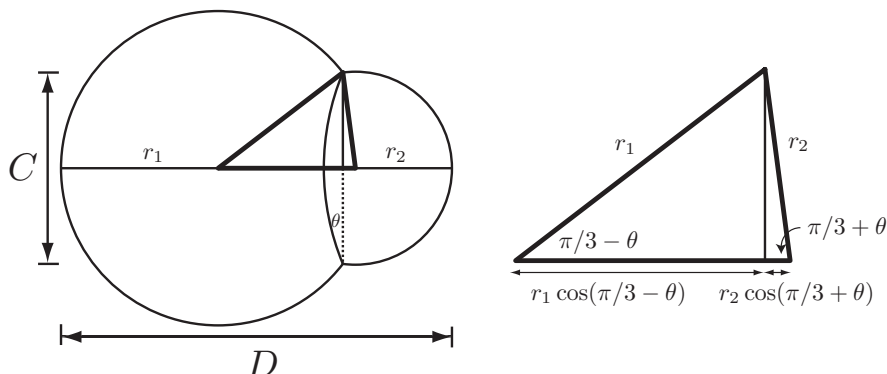
FIGURE 5. The perimeter of the standard double bubble is greater than π times its diameter.

Applying the formula for radii (see (7.3)), we obtain an expression that can be solved for C, the length of the chord connecting the two vertices:

$$D = C \left(\frac{1 + \cos(\pi/3 - \theta)}{2 \sin(2\pi/3 + \theta)} + \frac{1 + \cos(\pi/3 + \theta)}{2 \sin(2\pi/3 - \theta)} \right).$$

Substituting for C in the formula for perimeter (see (7.2)) yields after some simplification the desired function in θ:

$$f(\theta) = P/D = \frac{8\pi \sin\theta \cos\theta + 3\sqrt{3}\theta}{6 \sin\theta \cos\theta + 3\sin\theta}.$$

Now, showing $f(\theta) > \pi$ is equivalent to showing

$$8\pi \sin\theta \cos\theta + 3\sqrt{3}\theta > \pi(6 \sin\theta \cos\theta + 3\sin\theta).$$

Moving the terms to the left-hand side and using the identity $\sin(2\theta) = 2\sin\theta \cos\theta$ gives the equivalent inequality

$$g(\theta) = \pi \sin(2\theta) + 3\sqrt{3}\theta - 3\pi \sin\theta > 0$$

with

$$g''(\theta) = -4\pi \sin(2\theta) + 3\pi \sin\theta = \pi \sin\theta(-8\cos\theta + 3).$$

For θ in $(0, \pi/3)$, $\sin\theta > 0$ while $\cos\theta > \cos(\pi/3) > 3/8$. Therefore, $g''(\theta)$ is strictly negative. Because $g(\theta)$ equals 0 at both endpoints of this interval, it must be positive along the interior, as desired. □

Proposition 5.1.2 (Suggested by Masters [Ma1]). *The standard double bubble is the only contractible double bubble that may be a minimizer.*

Proof. Assume that there is a perimeter-minimizing contractible double bubble Σ enclosing areas A_1 and A_2 that is not the standard double bubble. By [F], the unique perimeter-minimizing solution in $\mathbf{R}^2$ for the same prescribed areas is a standard double bubble, Θ. By the perimeter bound (Proposition 3.8), Σ has perimeter at most three. Since Θ has less perimeter than Σ, the perimeter P of Θ satisfies $P < 3$.

By Proposition 5.1.1, the diameter D of Θ satisfies $D < P/\pi < 1$. Hence, it fits on the torus with its axis of symmetry in any direction (Figure 6). Therefore, Θ is
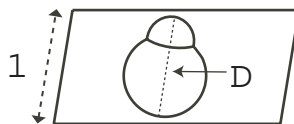
FIGURE 6. A standard double bubble with minimizing perimeter has diameter $D \leq 1$ and hence fits on the torus.
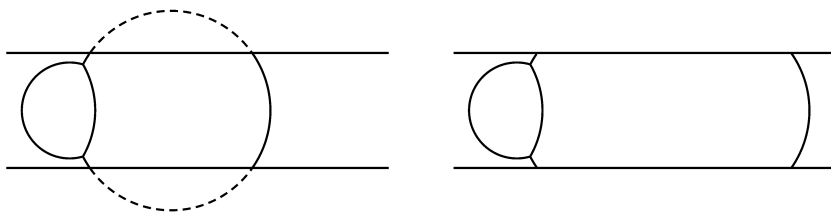

FIGURE 7. If a standard double bubble does not fit on the cylinder, we can easily construct a double bubble with even less perimeter.

a double bubble on the torus enclosing areas A_1 and A_2 with less perimeter than Σ, a contradiction. $\square$

Remark 5.1.3. There is an easier argument on the infinite cylinder: when Θ does not fit with its diameter parallel to the sides of the cylinder, we can modify its boundary to create a new double bubble with less perimeter than Θ and enclosing the same areas, contradicting the assumption that there is a nonstandard contractible minimizer (see Figure 7).

5.2. **Swaths of contractible components.** Recall that by Definition 9 a swath is a set of contractible components with noncontractible and connected closure, and with noncontractible complement, such as the standard chain of Figure 1. This section shows that minimizing chains are standard chains (Definition 15), exactly as in Figure 1. Lemma 5.2.11 shows that minimizing swaths must be chains consisting of curvilinear quadrilateral components. Lemma 5.2.13 shows that asymmetric chains satisfying regularity cannot be minimizers since they have perimeter greater than three. Finally, Proposition 5.2.15 shows that minimizing symmetric chains are standard.

Remark 5.2.1. In this section, we often speak about contractible four-sided components as if they sat in the plane rather than a torus, because it is easier to talk about lines, reflections, and circles in the plane. (Lines can wrap around a torus infinitely many times, reflections across lines on a torus do not always exist, and circles on a torus may overlap themselves.)

Lemma 5.2.2 ([W, Lemma 5.31]). *Given a contractible four-sided component in a regularity-satisfying double bubble or its exterior, the vertices of the component lie on a circle or a line. They lie clockwise on the circle or line in the same or opposite order as they appear clockwise on the boundary of the four-sided component.*

Proof. [W, Lemma 5.31] actually considers four-sided components in regularity-satisfying triple bubbles in the plane. However, the proof looks at each contractible

four-sided component in isolation from the rest of the triple bubble, using only the facts that each side has constant curvature, that the sides do not intersect each other, and that the sides meet at angles of $2\pi/3$. Thus, the proof also applies to contractible four-sided components of double bubbles in the torus.

[W, Lemma 5.31] is also slightly more limited than our result, not proving that if the vertices of the component lie on a line, then they do so in the same order that they appear on the boundary of the component. However, it does offer a proof of the analogous fact about the vertices when they lie on a circle, and this proof can be easily modified to give our complete result. $\qquad \square$

Lemma 5.2.3. *Given a contractible four-sided component with vertices A, B, C, D in a regularity-satisfying double bubble or its exterior, $AB = CD$ unless sides AB and CD of the component lie on the same circle or line.*

Proof. By Lemma 5.2.2, vertices A, B, C, D lie on a circle or line ω. Suppose that sides AB and CD do not lie on the same circle or line. By regularity in double bubbles, sides AB and CD have constant curvature and lie on some circles or lines ω_1 and ω_2, respectively. Also by regularity, sides AB and CD have the same curvature, and they meet side BC at equal angles (namely, $2\pi/3$). Hence, ω_1 and ω_2 are reflections of each other across the perpendicular bisector of $\overline{BC}$. Thus, $\omega_1 \cap \omega = \{A, B\}$ and $\omega_2 \cap \omega = \{C, D\}$ are reflections of each other across the perpendicular bisector of $\overline{BC}$ as well, implying that $AB = CD$. $\qquad \square$

Lemma 5.2.4. *Given a contractible four-sided component in a regularity-satisfying double bubble or its exterior, the vertices of the component do not lie on a line.*

Proof. Let the vertices of the component be A, B, C, D in that order along the boundary of the component. Suppose, for sake of contradiction, that the vertices lie on a line ℓ. By Lemma 5.2.2, A, B, C, D lie in that order along ℓ — without loss of generality, ℓ is horizontal with A at the left and D at the right. Because $BC \neq DA$, Lemma 5.2.3 implies that sides DA and BC of the component lie on the same circle or line. The only circle or line that contains all four vertices A, B, C, D is ℓ, so that sides DA and BC of the component are linear and lie along ℓ. However, this is impossible because then sides DA and BC would overlap, a contradiction. $\qquad \square$

Lemma 5.2.5 (Compare [W, proof of Lemma 5.3]). *Given a contractible four-sided component in a regularity-satisfying double bubble or its exterior, the vertices of the component form a nondegenerate isosceles trapezoid. The vertices appear clockwise on this trapezoid in the same order as they appear clockwise on the boundary of the component. If the trapezoid is not a rectangle, then the two boundary curves subtended by the parallel sides of the trapezoid lie on a circle circumscribing the polygonal quadrilateral.*

Proof. By Lemmas 5.2.2 and 5.2.4, the polygonal quadrilateral $ABCD$ formed by the vertices of the component is inscribed in some circle ω; the vertices appear on the quadrilateral in the same order that they appear on the circle (although perhaps with different orientation).

By regularity, curvilinear sides AB, BC, CD, DA cannot all lie along ω. Without loss of generality, assume that sides AB and CD do not lie on the same circle. By Lemma 5.2.3, we have $AB = CD$. Hence, because quadrilateral $ABCD$ is inscribed in a circle, it must be an isosceles trapezoid with $\overline{BC} \parallel \overline{DA}$.

If sides BC and DA of the curvilinear component do not lie on the same circle, then by the same analysis we have $BC = DA$, so that quadrilateral $ABCD$ is a rectangle (Figure 8(a)). Also, if trapezoid $ABCD$ is not a rectangle, say because $BC \neq DA$, then sides BC and DA of the curvilinear component *do* lie on the same circle (Figure 8(b)).
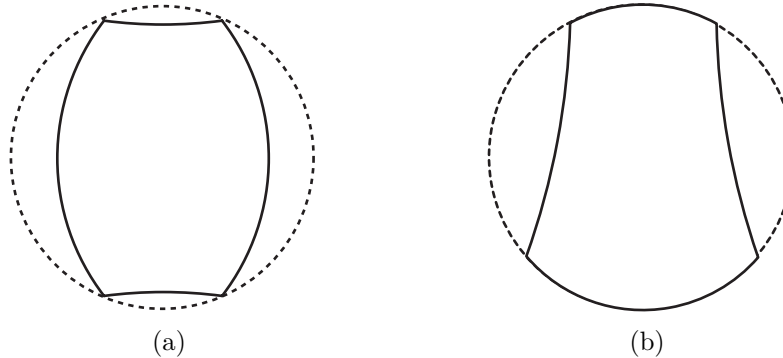

(a)   (b)

FIGURE 8. The vertices of any four-sided component lie on a circle. Either (a) the vertices form a rectangle, or (b) two opposite sides of the component lie along the same circle.

Finally, it remains to show that the vertices appear clockwise on the boundary of the component in the same order as they appear clockwise on the trapezoid. Because the vertices form a trapezoid (say, with $\overline{BC} \parallel \overline{DA}$), the four-sided component is symmetric about the perpendicular bisector m of sides BC, DA of the component. It follows easily that neither side AB nor side CD intersects m, implying the desired fact about the order of the vertices.   □

Lemma 5.2.6. *In a contractible four-sided component of a highest-pressure region in a regularity satisfying double bubble or its exterior, the arcs bordering a least-pressure region are subtended by parallel chords, and the arcs bordering a medium-pressure region are subtended by congruent chords.*

Proof. By Lemma 5.2.5, the chords subtending the boundary curves of the component form an isosceles trapezoid inscribed in a circle ω. If the trapezoid is a rectangle, we are done. Otherwise, the two boundary curves σ, σ' subtended by parallel chords lie on ω, and the other two boundary curves μ, μ' are subtended by congruent chords. Because σ, σ' meet the remainder of ω at angles *greater* than $2\pi/3$ (namely, π), the other boundary curves μ, μ' of the component bow within ω. In other words, the curvature of μ is less than the curvature of ω, which equals the curvature of σ. It follows that σ, σ' bound a region of less pressure than μ, μ' do.   □

Lemma 5.2.7. *For a minimizing tiling on a torus, the average number of edges per component is exactly six.*

Proof. The Euler characteristic on the torus is $v - e + c = 0$, where v is the number of vertices, e is the number of edges, and c is the number of components. By regularity, (Proposition 3.1 and Remark 3.2), v, e and c are all finite. Now, let a be the average number of edges per component. Then $c = \frac{2}{a}e$, since each edge is

adjacent to two components. Similarly, $v = \frac{2}{3}e$, since each edge is adjacent to two vertices and (by regularity, Proposition 3.1) each vertex is adjacent to three edges. Substituting into the Euler formula, we obtain $(\frac{2}{3} - 1 + \frac{2}{a})e = 0$, from which we obtain $a = 6$. $\hspace{1cm}\square$

Proposition 5.2.8. *A minimizing double bubble with a contractible component in a region of least pressure must be either an octagon-square tiling or a hexagon tiling.*

Proof. If the pressures of the three regions are equal, then by Proposition 4.5 either the double bubble is a hexagon tiling or the double band; the latter case is impossible because the double bubble must have a contractible component by assumption. Hence, we may assume that the three regions do not all have the same pressure.

If any component in the double bubble plus exterior is a vertex-free band, then its boundary curves are closed geodesics. This either implies that both regions and the exterior have equal pressure, contradicting our previous assumption, or we can slide the band until it collides with some other boundary curve, for a violation of regularity (Remark 3.3). Thus, no component is a vertex-free band.

If any region has contractible complement, then by Proposition 5.1.2, the double bubble is the standard double bubble. Again, the least-pressure region consists solely of noncontractible components, a contradiction. Hence, we may assume that every region has noncontractible complement.

Let R_1 be a region of highest pressure. From the previous paragraph and Lemma 4.1, any noncontractible component of R_1 is a band. Because R_1 is of highest pressure, such a band cannot have any vertices; if it did, then each boundary edge of the band would have positive turning angle, a contradiction. However, we have already shown that no region contains a vertex-free band. Therefore, any component in R_1 is contractible. By Proposition 3.7, R_1 has at most two contractible components. Also, each component of R_1 contains at most four vertices (Remark 3.6). Hence, R_1 contains at most eight vertices.

Let R_2 be a region of lowest pressure, containing a contractible component with n sides (and perhaps other contractible, or noncontractible, components). This component has least pressure but is not polygonal (because the three pressures are not all equal), so $n > 6$ (Remark 3.6). By Remark 3.4, n is even, implying that $n \geq 8$.

Because every vertex of R_1 is a vertex of R_2 (and vice versa), and because neither R_1 nor R_2 contains a vertex-free band, it follows that R_1 consists of two contractible four-sided components and that R_2 consists of a single contractible curvilinear octagon.

Because no region has contractible complement, R_2 and some four-sided component C in R_1 must form a chain that wraps around the torus. We now consider two cases.

Case 1: C borders two opposite sides of R_2. Then just as R_2 and C form a chain wrapping around the torus, so too do R_2 and the *other* four-sided component of R_1 — wrapping around the torus in a different direction. In this case, the double bubble is a tiling. Because R_1 and R_2 each have eight vertices, R_3 must have eight vertices as well, for a total count of 24 vertices. But by Lemma 5.2.7, the average number of vertices per component must be exactly six. Therefore, there are exactly four components, implying that R_3 consists of one curvilinear octagon. Therefore, as shown in Figure 9(a), the configuration is an octagon-square tiling.
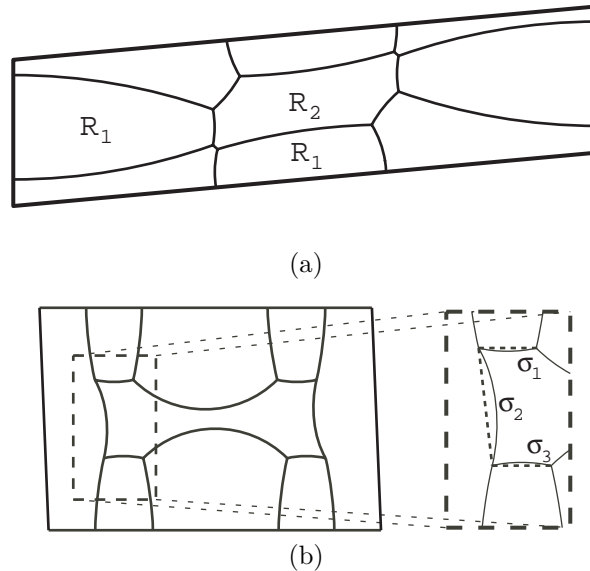
(a)

(b)

FIGURE 9. Two double bubbles with a purported contractible low-pressure component.

Case 2: C borders two sides σ_1, σ_3 of R_2 as depicted in Figure 9(b). Let σ_2 be the side of the eight-sided component between them. By Lemma 5.2.6, the chords subtending σ_1 and σ_3 are parallel. Therefore, one of these chords must meet the chord subtending σ_2 at a nonobtuse angle. But then the corresponding arcs meet at an angle less than $2\pi/3$, a contradiction.

Therefore, if a least-pressure region contains a contractible component, then the double bubble is a hexagon tiling or an octagon-square tiling. $\square$

Corollary 5.2.9. *A perimeter-minimizing tiling is either a hexagon tiling or an octagon-square tiling.*

Lemma 5.2.10. *Let γ be a boundary edge separating two components X_1 and X_2 in a minimizing double bubble, where both X_1 and X_2 have different pressures from the third region of the double bubble. Let the two boundary arcs of X_1 adjacent to γ be σ_1 and σ_2, and let the two boundary arcs of X_2 adjacent to γ be μ_1 and μ_2. (The σ_i coincide if X_1 has two sides, and the μ_i coincide if X_2 has two sides.) Suppose that σ_1 and σ_2 have the following property: they lie on a common arc, and this arc forms an immersed curvilinear digon with γ. Then μ_1 and μ_2 have this property as well.*

Proof. Consider the curvilinear digon bordered by γ and the arc containing σ_1 and σ_2. We may add a third arc μ to the curvilinear digon to form an immersed standard double bubble. Because the cocycle condition holds for both the given double bubble and for the immersed standard double bubble, μ_1, μ_2, and μ have the same (signed) curvatures. Therefore, μ_1 and μ_2 lie on the same circular arc μ, which forms a curvilinear digon with γ. $\square$

Lemma 5.2.11. *A minimizing swath must be a chain of four-sided components. The exterior is a band with less pressure than either enclosed region of the swath.*

Proof. By Proposition 5.2.8, every component of a region of least pressure is non-contractible. This implies two facts: first, because the swath contains contractible components from two regions, the exterior region is the only region of least pressure, as claimed. Second, because the exterior is a region of least pressure, each component of the exterior region is noncontractible. Also, the *complement* of each component in the exterior is noncontractible, because this complement contains the swath formed by the two enclosed regions. Thus, by Lemma 4.1, each component of the exterior is a band. Because the swath formed by the two enclosed regions is connected (Remark 3.3), the exterior region consists of exactly one band, as claimed. Also, the swath cannot contain a contractible cycle of components, because otherwise the cycle would encircle a contractible component of the exterior.

Consider a chain K in the swath. The components in K must alternate between two regions R_1 and R_2 with higher pressure than the exterior. We will show that the swath contains no additional component of R_1 or R_2 attached to the chain, as in Figure 10.
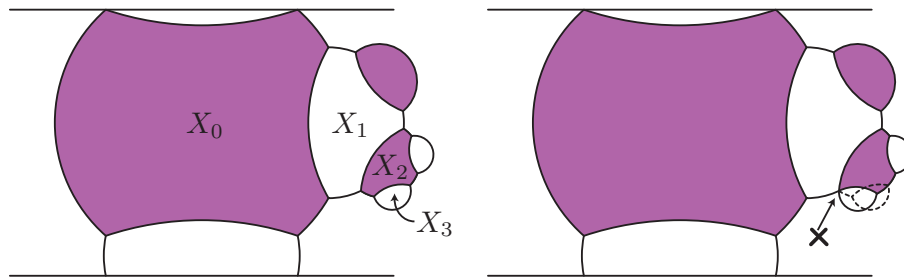

FIGURE 10. Extra components attached to a chain could be slid to create a double bubble violating regularity.

Suppose, for sake of contradiction, that the swath contained an additional component X_1 attached to some component X_0 in the chain. Either X_1 has two sides, or another component X_2 attaches to it. In this case, either X_2 has two sides, or another component X_3 attaches to it. Because the swath contains no contractible cycle, this process terminates after yielding some sequence of components $X_0, X_1, \ldots, X_k$, where X_k has two sides. By Lemma 5.2.10, the boundary edges of X_{k-1} meeting the curvilinear digon X_k lie on the same circle. Therefore, X_k can be slid along the boundary of this circle until it bumps into the boundary of another component, creating a minimizer with an illegal singularity (Proposition 3.1), a contradiction. Therefore, our original assumption was false, and no additional components of R_1 or R_2 attach to the chain.

Thus, the double bubble is a chain. Because the exterior consists of a single band, each component in the chain is enclosed by four boundary curves: two separating it from the other components of the chain, and two separating it from the exterior. Therefore, the double bubble is a chain of four-sided components. □

Lemma 5.2.12. *Consider a chain of four-sided components in a perimeter-minimizing double bubble. If the arcs separating each component from the exterior of the chain lie on the same circle, then the perimeter of the chain is greater than 3.*
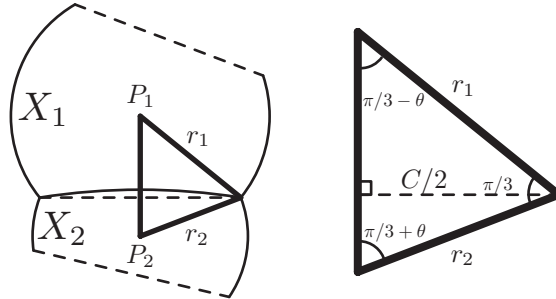

FIGURE 11. Triangulating chains in which opposite arcs are cocircular leads to a perimeter estimate.

Proof. Consider two adjacent components X_1 and X_2 of the chain, as drawn in Figure 11 so that the chord C of the arc separating them is horizontal and lies below X_1. Consider the four sides of the polygonal quadrilateral inscribed on the vertices of X_1. By Lemma 5.2.11, the chain neighbors a least-pressure region; hence, by Lemma 5.2.6, C and the side opposite it in the quadrilateral are congruent. It follows that C lies below P_1, the center of the circle passing through the vertices of X_1. Similarly, C lies above the center P_2 of the circle passing through the vertices of X_2.

We draw the triangle whose vertices are P_1, P_2, and an endpoint P_3 of the chord C. The triangle has angles $\pi/3 - \theta, \pi/3, \pi/3 + \theta$ and side lengths $r_1, r_2, P_1 P_2$ for some θ, r_1, r_2. Because P_1 and P_2 lie on opposite sides of C, angles $\angle P_1 P_2 P_3$ and $\angle P_2 P_1 P_3$ are acute, implying that $\theta \in (-\pi/6, \pi/6)$. Using the Law of Sines in this triangle, we have $r_1 = P_1 P_2 \frac{\sin(\pi/3+\theta)}{\sin(\pi/3)}$. Also, $|C|/2 = r_1 \sin(\pi/3 - \theta)$. Hence,

$$|C| = \frac{4}{\sqrt{3}} P_1 P_2 \sin(\pi/3 + \theta) \sin(\pi/3 - \theta).$$

Applying the standard transformation $2 \sin x \sin y = \cos(x - y) - \cos(x + y)$ shows that

$$|C| = \frac{2}{\sqrt{3}} P_1 P_2 (\cos(2\theta) - \cos(2\pi/3)).$$

With the restriction $\theta \in (-\pi/6, \pi/6)$, we see that $|C| \geq \frac{2}{\sqrt{3}} P_1 P_2$.

As we vary over pairs of adjacent components, the segments analogous to $\overline{P_1 P_2}$ form a closed noncontractible curve with total length greater than or equal to one. Hence, the total length of the chords analogous to C is at least $\frac{2}{\sqrt{3}} > 1$, implying that the total length of the corresponding arcs is also greater than one. Furthermore, the arcs separating the exterior and the chain form two closed noncontractible curves with total length greater than or equal to two. Therefore, the total perimeter of the chain is greater than three. □

Lemma 5.2.13. *No asymmetric chain is perimeter minimizing.*

Proof. Suppose, for sake of contradiction, that an asymmetric chain is minimizing. By Lemma 5.2.11, it consists of four-sided components, and the exterior is a least-pressure region consisting of a single band. By Lemma 5.2.5, each component of the chain has vertices that form an isosceles trapezoid. If the trapezoid corresponding to each component is a rectangle, then the rectangles (and hence the chain) are all symmetric about a closed geodesic (with the same homology as the chain), a contradiction. Therefore, one of the components has vertices that form a nonrectangular isosceles trapezoid. By Lemma 5.2.5 and Lemma 5.2.6, the arcs separating this component from the least-pressure exterior lie on a single circle. By Lemma 5.2.10, the same is true for all other components in the chain. Applying Lemma 5.2.12, the perimeter of the chain is greater than 3. By Proposition 3.8, the asymmetric chain cannot be minimizing, a contradiction.   □

Lemma 5.2.14. *A minimizing symmetric chain must be a standard chain.*

Proof. By Lemma 5.2.11, such a chain consists of four-sided components, where alternate components lie in each of two regions. Thus, the chain contains an even number of components. By Lemma 5.2.5, each four-sided component is circumscribed about an isosceles trapezoid. Because the chain is symmetric, and using Lemma 5.2.6, we find that each four-sided component must be circumscribed about a rectangle. It easily follows that any pair of adjacent components in the chain is congruent to any other pair.
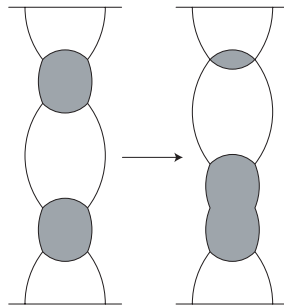

FIGURE 12. Rearranging the components of a symmetric chain with more than two components brings about a violation of regularity, proving that a minimizing chain has two components.

Now if the chain contains more than two components, then rearranging four adjacent components as in Figure 12 maintains area and perimeter while violating regularity, so that the original chain cannot be perimeter minimizing. Therefore, the chain contains exactly two components. No curvilinear quadrilateral lies in a least-pressure region, so the two components of the chain belong to the highest- and medium-pressure regions. Applying the cocycle condition (Proposition 3.1) shows that the chain is standard.   □

Proposition 5.2.15. *Every perimeter-minimizing swath is a standard chain.*

Proof. By Lemma 5.2.11, a minimizing swath must be a chain of four-sided components. By Lemma 5.2.13 a minimizing chain must be symmetric. By Lemma 5.2.14, a minimizing symmetric chain must be a standard chain.   □

Remark 5.2.16. While we will show that a minimizing band lens must have the same homology as a shortest closed geodesic (Lemma 5.3.6), we do not have a proof of the analogous fact for minimizing standard chains. However, *Mathematica* plots for perimeter suggest that this property holds at least for the particular tori described by Figure 2.

5.3. **A single band adjacent to a contractible set of components.** Proposition 5.3.5 uses regularity (Proposition 3.1) and Proposition 5.2.8 to show that the only potential minimizer in the class of double bubbles with a single band adjacent to a contractible set of components is the band lens.

Throughout this section, we will consider a proposed minimizer with regions R_1, R_2 and the exterior, such that R_1 is contractible and R_2 contains a single band (and possibly some contractible components).

Lemma 5.3.1. *In a minimizing double bubble with some contractible region R_1, and another region R_2 containing a single band, the pressure difference between R_2 and the third (exterior) region must be zero, that is, all curves separating R_2 from the exterior must be straight line segments.*

Proof. Consider a curve separating the band in R_2 from the exterior. If it is a straight line segment, the proof is complete, since all curves separating R_2 from the exterior have the same curvature by regularity (Proposition 3.1).

If the curve has nonzero curvature, continuously straighten the curve to reduce perimeter, while sliding one component of the boundary of the band in R_2 in order to maintain area. Either we will be able to straighten all such pieces while reducing perimeter, or we will first cause a violation of regularity. $\square$

Lemma 5.3.2. *In a minimizing double bubble with some contractible region R_1, and another region R_2 containing a single band, R_1 must have higher pressure than each of R_2 and the exterior region.*

Proof. By Lemma 5.3.1, R_2 and the exterior have the same pressure. Since the double bubble is neither a tiling nor the double band, R_1 cannot also have the same pressure (Proposition 4.5). Moreover, by Proposition 5.2.8, since R_1 is contractible it cannot be a region of least pressure. Thus, R_1 must have higher pressure than both R_2 and the exterior. $\square$

Corollary 5.3.3. *All contractible components of the double bubble plus exterior belong to R_1, and R_2 consists of a single band.*

Proof. By Lemma 5.3.1 and Lemma 5.3.2, R_2 and the exterior have equal pressure, less than that of R_1. Thus, by Proposition 5.2.8, their components must be noncontractible. Hence only components of R_1 can be contractible, and R_2 consists of the only noncontractible component of the double bubble — a band. $\square$

Lemma 5.3.4. *In a minimizing double bubble with some contractible region R_1, and another region R_2 containing a single band, all components of R_1 are lenses.*

Proof. By Lemma 5.3.2, R_1 has higher pressure than both R_2 and the exterior. By Remarks 3.6 and 3.4, all components of R_1 have either two or four sides. We will assume that some component of R_1 is a curvilinear quadrilateral Z, and show that this leads to a contradiction. Let the four sides of Z be labelled clockwise x_1, y_1,

x_2, y_2, where x_1, x_2 border R_2 (a single band, by Corollary 5.3.3) and y_1, y_2 border Y_1, Y_2 in the exterior.

If x_1 and x_2 lie on the same side of R_2, then either y_1 or y_2 lies along that same side of the band; the corresponding component Y_1 or Y_2 is contractible, contradicting Corollary 5.3.3. Otherwise, x_1 and x_2 lie on opposite sides of R_2. In this case, the closure of $R_2 \cup Z$ has contractible complement, so that both Y_1 and Y_2 are contractible — again contradicting Corollary 5.3.3.

Thus, all components of R_1 are curvilinear digons. Now, by Lemma 5.3.1, R_2 and the exterior have equal pressure. Therefore, all curves bounding components of R_1 have equal, nonzero curvature, and so all curvilinear digons in R_1 must be lenses. $\square$

Proposition 5.3.5. *A perimeter-minimizing double bubble with a single band adjacent to a contractible set of components must be the band lens.*
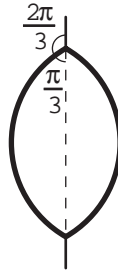
FIGURE 13. Any lens from R_1 must be embedded in a closed geodesic.

Proof. By Corollary 5.3.3 and Lemma 5.3.4, one region R_1 consists only of lenses and the other region R_2 consists of a single band. Corollary 5.3.3 also implies that the exterior contains no contractible components. It follows from Lemma 4.1 that the exterior consists solely of bands. Because the double bubble is connected (Remark 3.3), the exterior consists of a single band.

Each lens in R_1 has two vertices, each the endpoint of some curve separating R_2 and the exterior. By Lemma 5.3.1, these curves are straight line segments. Furthermore, both segments must lie on the line passing through the diameter of the lens, as depicted in Figure 13. Hence, the diameters of the lenses, together with the segments separating R_2 from the exterior, form closed geodesics. Since R_2 and the exterior each consist of exactly one band, the boundary of the double bubble consists of exactly two closed geodesics, with lenses from R_1 embedded in at least one of them. Each of these two closed geodesics has length $L' \geq 1$.

Let the lenses in R_1 have total perimeter p and let their diameters have total length d. If $d \geq 1$, then the perimeter of the double bubble is at least $(2L'-d)+p > (2L' - d) + 2d \geq 3$, so by Proposition 3.8 the double bubble is not perimeter minimizing. Otherwise, we can move all the lenses so that they are adjacent to each other and embedded in the same closed geodesic, creating a new double bubble enclosing and separating the same areas. If there is more than one lens, then the new double bubble violates regularity (Figure 14), a contradiction. Therefore, the original minimizing double bubble contains exactly one lens, and so must be the band lens. $\square$
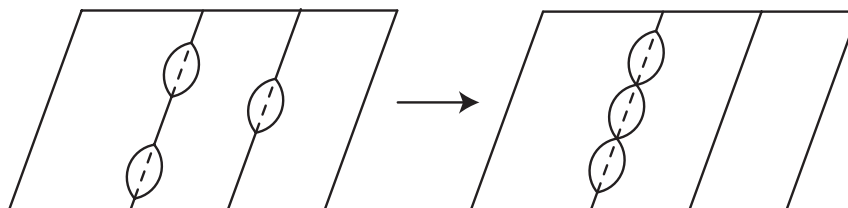
FIGURE 14. Multiple lenses can be slid together to violate regularity. Therefore, there can only be one lens.

Lemma 5.3.6. *In a minimizing band lens, the band must have the same homology as a shortest closed geodesic.*

Proof. We show that any pair of areas that can be enclosed by a band lens wrapping around a longer direction of the torus can be enclosed by an embedded band lens wrapping around a short direction. Suppose not. By Proposition 3.8, the perimeter of the original band lens is less than or equal to three, so the perimeter of the lens is less than two. Therefore, the diameter of the lens is less than one, and the lens can be embedded on the torus with its axis along a short direction (Remark 2.3.2). Thus, if the band lens wrapping around the short direction does not exist, the closed geodesic not meeting the lens cannot get close enough to make the band thin enough. But if this happens in the short direction, then it will certainly happen for the band lens wrapping around the longer direction, a contradiction.

Hence, if there is a minimizing band lens wrapping around a longer direction of the torus — say, parallel to a closed geodesic of length $L_0 > 1$ — there is a band lens in a short direction of the torus enclosing the same areas. However, the new band lens has perimeter $2(L_0 - 1)$ less than the initial band lens, a contradiction. Therefore, any minimizing band lens has the same homology as a shortest closed geodesic. □

5.4. **Tilings.** This section will show that every perimeter-minimizing tiling is a standard hexagon tiling (Proposition 5.4.23). Corollary 5.2.9 showed that if a tiling is minimizing, then it is a hexagon tiling or an octagon-square tiling. Proposition 5.4.13 shows that every minimizing hexagon tiling is standard, and Proposition 5.4.22 shows that no octagon-square tiling is perimeter minimizing.

Hexagon tilings. Lemmas 5.4.3, 5.4.5, and 5.4.6 show that a minimizing hexagon tiling divides the torus into exactly three hexagonal components. Lemma 5.4.9 shows that such a tiling can be transformed without changing perimeter to make the three hexagons congruent translations of each other. By Lemma 5.4.3, the perimeter of the new hexagon tiling is greater than 3 except on the hexagonal torus, implying that the original minimizing hexagon tiling lies on the hexagonal torus. Finally, Proposition 5.4.13 shows that any such hexagon tiling is the standard hexagon tiling.

Lemma 5.4.1. *The number of components in a perimeter-minimizing hexagon tiling plus exterior is a multiple of three.*

Proof. By Proposition 3.1, there are finitely many vertices in the tiling — say, v — each bordering three distinct regions. Consider any of the three regions R_i. Each

of its components borders six distinct vertices in the tiling, and each vertex in the tiling borders exactly one component of R_i. Hence, R_i contains $v/6$ components. Therefore, $v/6$ is an integer, and there are $3(v/6)$ components in all. $\qquad\square$

Lemma 5.4.2. *Given $n \geq 3$ real numbers $a_1, \ldots, a_n \in [0, A/3]$ such that $\sum_{i=1}^{n} a_i = A$, we have $\sum_{i=1}^{n} \sqrt{a_i} \geq 3\sqrt{A/3}$.*

Proof. Because $\sqrt{A}$ is strictly concave, $\sum_{i=1}^{n} \sqrt{a_i}$ attains its minimum at some vertex of the given domain. At each vertex, three of the a_i equal $A/3$, and the other a_i equal 0. The result follows immediately. $\qquad\square$

Lemma 5.4.3. *In a hexagon tiling on a torus with area A such that each component has area at most $A/3$, the perimeter of the tiling is greater than or equal to 3. Equality can hold only for a tiling on a hexagonal torus with three components in the tiling plus exterior.*

Proof. By Lemma 5.4.1, if there are n components in the tiling plus exterior, then $n \geq 3$. Let $a_1, \ldots, a_n$ be the areas of the n components $C_1, \ldots, C_n$, and let $p_1, \ldots, p_n$ be their perimeters. By the isoperimetric inequality for hexagons, $p_i \geq \sqrt{8\sqrt{3}}\sqrt{a_i}$ for each i, with equality if and only if C_i is a regular hexagon.

The total perimeter of the tiling, $\frac{1}{2}\sum_{i=1}^{n} p_i$, is thus greater than or equal to $\sqrt{2\sqrt{3}}\sum_{i=1}^{n} \sqrt{a_i}$. By Lemma 5.4.2, $\sqrt{2\sqrt{3}}\sum_{i=1}^{n} \sqrt{a_i}$ is in turn greater than or equal to $\sqrt{2\sqrt{3}} \cdot 3\sqrt{A/3} = \sqrt{6\sqrt{3}A}$.

Because $A \geq \sqrt{3}/2$ with equality only for the hexagonal torus (Remark 2.1.1), the perimeter of the tiling is at least $\sqrt{6\sqrt{3}A} \geq \sqrt{6\sqrt{3}\frac{\sqrt{3}}{2}} = 3$. The perimeter is exactly 3 only if equality holds in all the intermediate inequalities. In other words, the torus must have area $\sqrt{3}/2$, which only holds for the hexagonal torus; each component must be a regular hexagon and there must be 3 components each with area $A/3$. $\qquad\square$

Remark 5.4.4. This method of comparing hexagon tilings to tilings of three equal areas was inspired by the result of Hales that the regular hexagon is the most efficient way to tile the plane — not necessarily with polygons — into unit areas [H]. Earlier, Fejes-Tóth proved the result for polygonal tilings ([FT1, Chapter III, Section 9, p. 84] or [FT2, Section 26, Corollary, p. 183] after [FT3]), a result more comparable than Hales' to ours.

Lemma 5.4.5. *Consider a hexagon tiling with $n \geq 6$ components in the tiling plus exterior. For each component H in any of the three regions, the perimeter of the tiling exceeds 1 plus the perimeter of H.*

Proof. We orient the torus so that the edges of the tiling form angles of 0 and $\pm\pi/3$ with the horizontal. Consider H along with the hexagons neighboring it, labelled $H_1, \ldots, H_6$ as in Figure 15.

Although some of $H_1, \ldots, H_6$ may coincide (if H and some component share more than one edge), we claim that H_2, H_4, and H_6 cannot *all* coincide. Assume not. H_3 borders H_2 from below, and H_5 borders H_6 from below; because $H_2 = H_6$, we have $H_3 = H_5$. We can likewise show that H_1, H_3, and H_5 are all the same component. The top edge, bottom-right, and bottom-left edges of this component border $H_2 =$
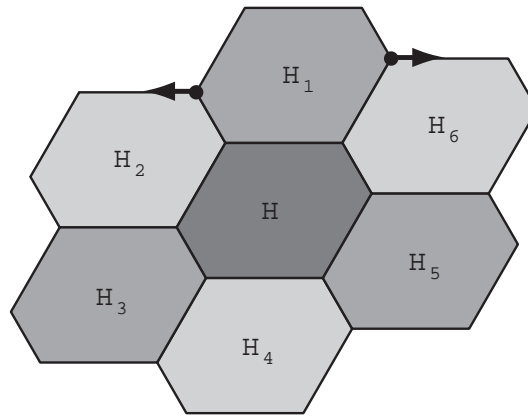
FIGURE 15. A hexagon H on the torus and its six neighbors, shaded according to which regions they lie in. (Some of the six hexagons neighboring H in this picture may actually represent the same component.) In order to find a noncontractible curve not intersecting H, we avoid travelling along a downward edge from either of the two marked vertices.

$H_4 = H_6$ (since, as shown in Figure 15, the top edge of H_3 borders H_2, the bottom-right edge of H_1 borders H_6, and the bottom-left edge of H_5 borders H_4); similarly, the other edges of this component border H. Thus, H_1 only neighbors two distinct components, H and H_2. Likewise, H and H_2 must only neighbor components in $\{H, H_1, H_2\}$. However, this is impossible; at least one of the $n - 3$ hexagons not in $\{H, H_1, H_2\}$ must neighbor some hexagon in $\{H, H_1, H_2\}$. Therefore, our original assumption was false, and we do not have $H_2 = H_4 = H_6$.

Without loss of generality, assume that $H_2 \neq H_6$. Notice also that $H_2 \neq H_5$ and $H_6 \neq H_3$, since alternate sides of H must border different regions (Remark 3.4).

The tiling contains $6n/3$ vertices, which is at least twelve, so we may pick a vertex v not on the boundary of H. Starting at v, we travel as follows: if we are at one of the vertices marked in Figure 15, travel horizontally as shown in Figure 15; otherwise, travel downward along an edge. It follows that we never travel downward along the edges depicted in Figure 16(a). Furthermore, because $H_6 \notin \{H_2, H_3\}$, we never travel to the right along the edges depicted in Figure 16(b). Similarly, we never travel to the left along the edges depicted in Figure 16(c). It easily follows that we never travel along a horizontal edge only to immediately backtrack along it, and that we never travel along the boundary of H or along any of the edges stemming from H.

There are finitely many vertices in the tiling (Proposition 3.1), implying that some portion of our journey begins and ends at the same vertex. Because we never backtrack along a horizontal edge, and because we never travel upward along an edge, this portion of our journey occurs along a closed, noncontractible curve (Figure 17). This curve has length at least 1, and it does not overlap the boundary of H, nor does it overlap any of the edges stemming from H, which have positive length. Hence, the boundary curves not enclosing H have total length greater than 1, as desired. □
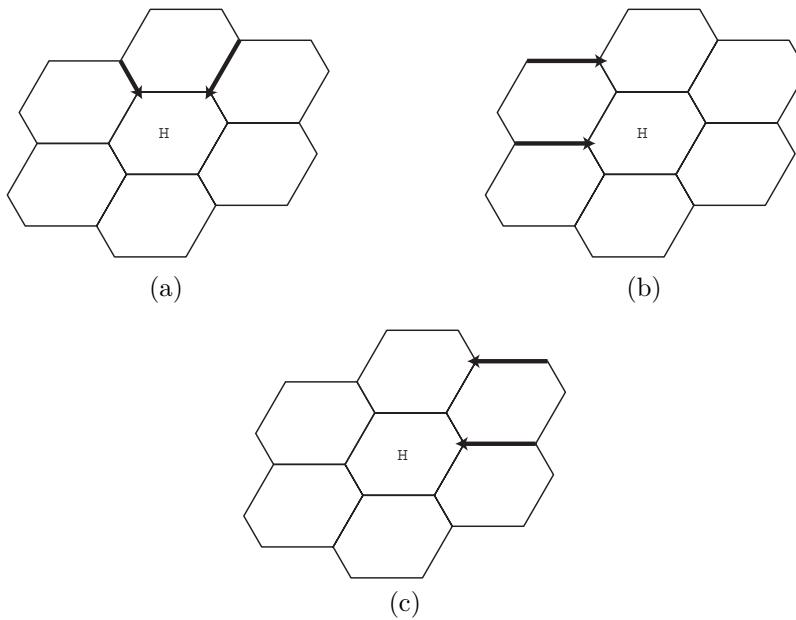
FIGURE 16. In order to bound the perimeter outside any component, we travel along the boundary in an attempt to form a closed noncontractible path. These diagrams show various movements that we do not perform during our journey.
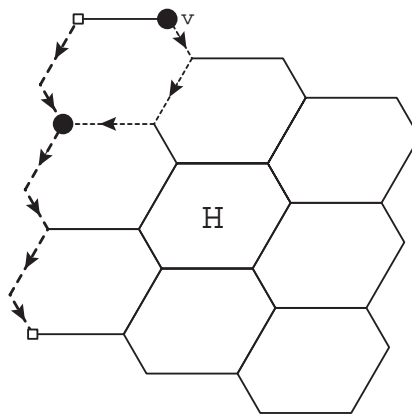

FIGURE 17. Starting at some vertex v and travelling along downward edges — except in those cases where doing so would lead us to a vertex of H — yields a closed, noncontractible curve not passing through the boundary of H, for a total perimeter greater than 3 when H is large.

Lemma 5.4.6. *No hexagon tiling with six or more components in the double bubble plus exterior is perimeter minimizing.*

Proof. We consider a hexagon tiling with six or more components. We first consider the case when the area α of some component is at least one-third the total area A of the torus, where $A \geq \sqrt{3}/2$ (Remark 2.1.1). By the isoperimetric inequality for hexagons, the perimeter of H is at least $\sqrt{8\sqrt{3}\alpha} \geq \sqrt{8\sqrt{3}(A/3)} \geq 2$. Therefore, by Lemma 5.4.5, the perimeter of the tiling is greater than $2 + 1 = 3$.

Next consider the case in which each component has area less than $A/3$. Because there are more than three components, by Lemma 5.4.3, the tiling has perimeter greater than 3.

In both cases, it follows by our perimeter bound (Proposition 3.8) that the hexagon tiling is not perimeter minimizing. $\qquad\square$

Remark 5.4.7. We first attempted to prove Lemma 5.4.6 by shrinking and expanding hexagons to preserve perimeter and areas and eventually contradict regularity. Although the perimeter of the tiling remains the same under such variations, the problem — as Wacharin Wichiramala explained to us — is that it is not always possible to preserve the regions' areas in order to reach such a contradiction.

Lemma 5.4.8. *In a hexagon with interior angles of $2\pi/3$, the sum of the lengths of any two adjacent sides equals the sum of the lengths of the opposite pair of adjacent sides.*

Proof. Orient the hexagon so that the two sides p_1, p_2 not being summed are horizontal, with one pair of remaining sides q_1, q_2 on the left and the other pair of remaining sides r_1, r_2 on the right. Let d be the distance between p_1 and p_2. The sides q_1, q_2 make an angle of $\pi/6$ with the vertical, so the sum of the lengths of q_1, q_2 is $d/\sin(\pi/3)$. Similarly, the sum of the lengths of r_1, r_2 is $d/\sin(\pi/3)$. Hence, the two sums are equal, as desired. $\qquad\square$

Lemma 5.4.9. *Given a hexagon tiling with three components in the tiling plus exterior, there exists a second hexagon tiling with the following properties: the edges of the new tiling are parallel to the edges of the original tiling; the perimeters of the two tilings are equal; and any two parallel edges of the new tiling are congruent. In particular, there exists a second hexagon tiling with the same perimeter, which still contains three components (now of the same area) in the tiling plus exterior.*

Proof. To get the desired result, we will find a variation that shrinks and expands hexagons as in Figure 18 to make the three areas equal. Perimeter is unchanged, since $dP/dA = \text{curvature} = 0$ through the whole family of double bubbles in the variation. We show that such a variation exists, and we also give an alternative geometric explanation for why perimeter remains unchanged.

Let a_i, b_i, c_i $(1 \leq i \leq 3)$ be the lengths of the edges in the hexagon tiling as in Figure 19; for each $i \in \{1, 2, 3\}$, the sides of some hexagon H_i are assigned the lengths $a_{i-1}, b_{i-1}, c_{i-1}, a_{i+1}, b_{i+1}, c_{i+1}$ (where indices are taken modulo 3). In this proof, as we modify the hexagon tiling, the values of the a_i, b_i, c_i will change accordingly.

Fix some subscript i_0. As shown in Figure 20, for some range of real numbers ϵ (which we will describe shortly), we may transform our hexagon tiling into another hexagon tiling, so that the edge lengths change as follows: $a_{i_0}, b_{i_0}, c_{i_0}$ change by
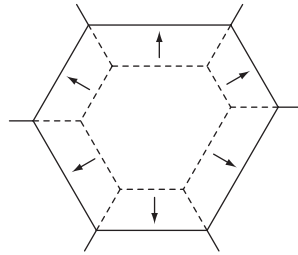
FIGURE 18. Expanding any component in a hexagon tiling or its exterior generates a double bubble enclosing different areas, but does not change perimeter. Hexagons can also be shrunk in a similar fashion.
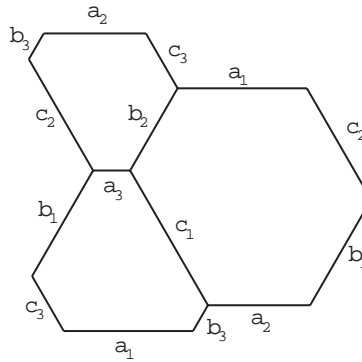

FIGURE 19. A hexagon tiling with three components in the tiling plus exterior.

-2ϵ, and the other a_i, b_i, c_i change by ϵ. More specifically, the range of such ϵ are those ϵ for which the side lengths of the resulting tiling are positive; namely, such that $a_i + \epsilon$, $b_i + \epsilon$, and $c_i + \epsilon$ are positive for $i \neq i_0$, and such that $a_{i_0} - 2\epsilon$, $b_{i_0} - 2\epsilon$, and $c_{i_0} - 2\epsilon$ are positive. Adding up the increases and decreases in perimeter shows that under this transformation, the perimeter of the hexagon tiling remains unchanged.

Without loss of generality, assume that $a_1 \geq a_2 \geq a_3$ initially. By Lemma 5.4.8, we have $(a_2 - a_3)/3 = (b_2 - b_3)/3 = (c_2 - c_3)/3$; set ϵ equal to this common value. As described in the last paragraph, we may shrink H_3 so that its edge lengths (a_i, b_i, c_i for $i \neq 3$) decrease by ϵ and so that a_3, b_3, c_3 increase by 2ϵ. It is easy to check that this ϵ is in the allowed range. Furthermore, after this transformation, we have $(a_2, b_2, c_2) = (a_3, b_3, c_3)$.

At this point, we now have $a_1 \geq a_2 = a_3$. We may now set ϵ equal to $\frac{a_1 - a_2}{3}$ and expand H_1 so that its edge's lengths increase by ϵ. In the resulting hexagon tiling, we have $(a_1, b_1, c_1) = (a_2, b_2, c_2) = (a_3, b_3, c_3)$, as desired. Furthermore, because the two transformations we used did not change perimeter, the resulting tiling and the initial tiling have the same perimeter. This completes the proof. □
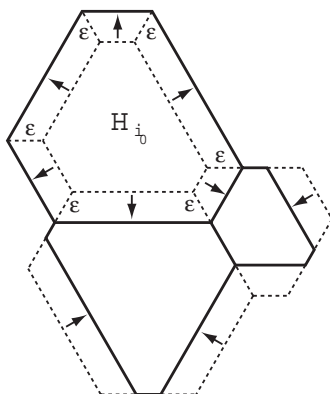
FIGURE 20. In a hexagon tiling with three components in the dou-
ble bubble plus exterior, we may expand (or shrink) one hexagon
so that six edges increase by ϵ while three edges decrease by 2ϵ.

Lemma 5.4.10. *Suppose that the hexagonal torus is partitioned into three hexagons
with interior angles of $2\pi/3$, where the hexagons are congruent translations of each
other. Suppose that the perimeter of the partition is at most 3, or that the sides
of the hexagon are parallel to the three short directions. Then the hexagons are
regular hexagons of side length $\frac{1}{3}$, with sides parallel to the three short directions of
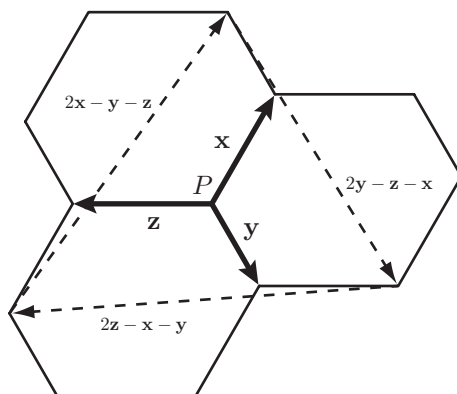the torus. The total perimeter of the partition is equal to three.*

FIGURE 21. In partitions of a torus into three hexagons that
are congruent translations of each other, examining three closed
geodesics (represented by the dotted vectors) shows that if the
hexagons are not regular, then the torus is not hexagonal.

Proof. Based on the partition, we draw a fundamental domain of the torus as in
Figure 21. Let P be the vertex in the center of the figure where the three hexagons
meet; let $\mathbf{x}$, $\mathbf{y}$, $\mathbf{z}$ be the vectors from P outward along the edges of the hexagon
tiling, with magnitudes x, y, z.

The three vectors $3\mathbf{x} - 3\mathbf{y}$, $3\mathbf{y} - 3\mathbf{z}$, $3\mathbf{z} - 3\mathbf{x}$ that start at P on the torus also end at P. If in addition $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are parallel to the short directions, it follows easily that x, y, z are multiples of $1/3$. In fact, x, y, z must then equal $1/3$: if any side length were longer, then the area of the three hexagons would exceed that of the torus, which is impossible.

Otherwise, suppose that the perimeter of the partition is at most 3. The three dotted vectors in Figure 21 all begin and end at the same point on the torus. On the hexagonal torus, any such vector has magnitude equal to 1, equal to $\sqrt{3}$, or greater than or equal to 2.

By assumption, the perimeter $3(x + y + z)$ of the partition satisfies

$$
\begin{aligned}
3 &\geq 3(x + y + z) \\
&= \frac{3}{4}((2x + y + z) + (2y + z + x) + (2z + x + y)) \\
&\geq \frac{3}{4}(\|2\mathbf{x} - \mathbf{y} - \mathbf{z}\| + \|2\mathbf{y} - \mathbf{z} - \mathbf{x}\| + \|2\mathbf{z} - \mathbf{x} - \mathbf{y}\|).
\end{aligned}
$$

It follows that the three magnitudes of the three dotted vectors equal one of the following triples: $1, 1, 1$; $1, 1, \sqrt{3}$; or $1, 1, 2$. The third case is impossible by the triangle inequality.

Without loss of generality, $2\mathbf{z} - \mathbf{x} - \mathbf{y}$ has maximal magnitude among the three vectors (i.e., either 1 or $\sqrt{3}$). Observe that

$$
\begin{aligned}
1 &= (2\mathbf{x} - \mathbf{y} - \mathbf{z}) \cdot (2\mathbf{x} - \mathbf{y} - \mathbf{z}) \\
&= 4\mathbf{x} \cdot \mathbf{x} + \mathbf{y} \cdot \mathbf{y} + \mathbf{z} \cdot \mathbf{z} - 4\mathbf{x} \cdot \mathbf{y} - 4\mathbf{x} \cdot \mathbf{z} + 2\mathbf{y} \cdot \mathbf{z} \\
&= 4x^2 + y^2 + z^2 - 4xy\cos(2\pi/3) - 4xz\cos(2\pi/3) + 2yz\cos(2\pi/3) \\
&= (x + y + z)^2 + 3x^2 - 3yz,
\end{aligned}
$$

or

$$
1 - (x + y + z)^2 = 3(x^2 - yz).
$$

Solving this and the two analogous equations gives $(x, y, z) = (1/3, 1/3, 1/3)$ or $(x, y, z) = (\sqrt{3}/9, \sqrt{3}/9, 4\sqrt{3}/9)$, depending on whether $2\mathbf{z} - \mathbf{x} - \mathbf{y}$ has magnitude 1 or $\sqrt{3}$. In the latter case, the partition has perimeter $3(x + y + z) = 2\sqrt{3} > 3$, a contradiction. In the former case, the partition consists of three regular hexagons of side length $\frac{1}{3}$. It easily follows that the sides of the hexagons are parallel to the three short directions of the torus, and that the perimeter of the partition is three. $\square$

Corollary 5.4.11. *If a perimeter-minimizing hexagon tiling on the hexagonal torus contains three components in the tiling plus exterior, then the tiling is a standard hexagon tiling.*

Proof. By the perimeter bound (Proposition 3.8), the perimeter of the tiling is at most 3. By this result and Lemma 5.4.9, the edges of the hexagon tiling are parallel to the edges of a hexagon tiling satisfying the hypotheses of Lemma 5.4.10. By Lemma 5.4.10, the edges of the new tiling are parallel to the three short directions of the torus. Therefore, so are the edges of the original tiling. $\square$

Corollary 5.4.12. *Any standard hexagon tiling has perimeter three.*

Proof. By Lemma 5.4.9, a standard hexagon tiling has the same perimeter as a partition of the torus satisfying the conditions of Lemma 5.4.10. By Lemma 5.4.10, this

partition has perimeter three, so the original standard hexagon tiling has perimeter three as well.   □

Proposition 5.4.13. *If a hexagon tiling is perimeter minimizing, then it is the standard hexagon tiling.*

Proof. By Lemma 5.4.1, the number of components in a minimizing hexagon tiling plus exterior is divisible by three. By Lemma 5.4.6, a hexagon tiling with six or more components is not perimeter minimizing. Thus, a perimeter-minimizing hexagon tiling must have exactly three components.

By Lemma 5.4.9, a perimeter-minimizing three-component hexagon tiling has the same perimeter as a hexagon tiling in which the area of each component is one-third the area of the torus. By Lemma 5.4.3, this perimeter is greater than three except on the hexagonal torus. Hence, by Proposition 3.8, any perimeter-minimizing three-component hexagon tiling must lie on the hexagonal torus. By Corollary 5.4.11, any such tiling is the standard hexagon tiling.   □
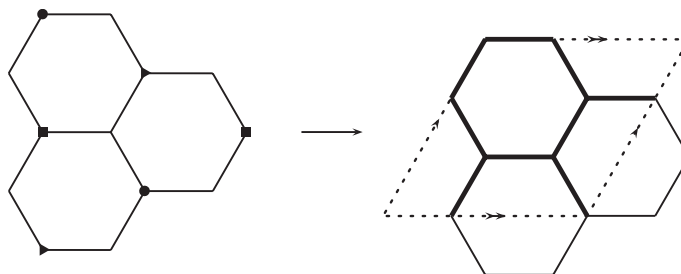

FIGURE 22. This tiling of three congruent hexagons fits on the hexagonal torus and is one of many perimeter-minimizing standard hexagon tilings.
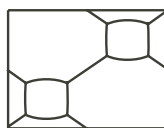

FIGURE 23. An octagon-square tiling always loses to some other double bubble.

Octagon-square tilings. Recall that by Definition 7, an octagon-square tiling is a tiling in which one of the three regions consists of two curvilinear quadrilaterals, and each of the other two regions consists of one curvilinear octagon (see Figure 23). The proof that an octagon-square tiling is not perimeter minimizing generalizes an argument described to us by Gary Lawlor eliminating most octagon-square tilings on rectangular tori. Lemma 5.4.14 and Corollary 5.4.15 compare a purported minimizing octagon-square tiling to the band lens in order to establish an upper bound on the length of the interface between the two eight-sided regions. Then, Proposition 5.4.19 and Lemma 5.4.21 establish a lower bound on the length of the same interface by examining projections of the boundary curves of the tiling

onto four different directions. Proposition 5.4.22 shows that these upper and lower bounds cannot be satisfied at the same time, implying that no octagon-square tiling is perimeter minimizing.

Lemma 5.4.14. *Let $A_1 \leq A_2 \leq A_0$ be three areas whose sum is equal to the area of the torus. Consider a lens enclosing area A_1 and with perimeter P and diameter D. For any pair of areas from A_0, A_1, A_2, a minimizing double bubble enclosing those areas has perimeter less than or equal to*

$$2 + P - D = 2 + \sqrt{A_1 \left(\frac{8\pi}{3} - 2\sqrt{3} \right)},$$

with equality only if there is a band lens that encloses area A_1 in the lens and A_2 in the band.

Proof. The expression for $P - D$ is easily derived, and we do not include the proof here (see, for instance, the derivations in Proposition 7.3.1).

Since all of the double bubbles that enclose two of the areas from $\{A_0, A_1, A_2\}$ have the same perimeter, it suffices to prove the claim for a minimizing double bubble enclosing the areas A_1, A_2.

If $D \geq 1$, then $2 + P - D > 2 + 2D - D \geq 3$. Thus, by the perimeter bound (Proposition 3.8), $2 + P - D$ exceeds the perimeter of a minimizing double bubble enclosing areas A_1, A_2.

Suppose instead that $D < 1$. Then the lens can be embedded on the torus so that its axis lies along a shortest closed geodesic γ_1 (Remark 2.3.2). There exists another shortest closed geodesic γ_2 such that γ_1 and γ_2 divide the complement of the lens into one part with area A_0 and another part with area A_2. The portions of these geodesics that lie outside the lens, along with the boundary of the lens, thus form the boundary of a double bubble enclosing areas A_1, A_2. The perimeter of the double bubble is at most the total perimeter of γ_2, the lens, and the portion of γ_1 lying outside the lens: $1 + P + (1 - D) = 2 + P - D$. Equality holds only if γ_2 does not intersect the interior of the lens — i.e., only if the double bubble is a band lens. Hence we have the desired result. $\qquad\square$

Corollary 5.4.15. *Let $A_1 \leq A_2 \leq A_0$ be three areas whose sum is equal to the area of the torus. A minimizing double bubble for any pair of areas from A_0, A_1, A_2 has perimeter less than*

$$2 + 2\sqrt{\pi A_1}.$$

Remark 5.4.16. Rather than examining partitions of the torus into a lens and two bands as in Lemma 5.4.14, one could prove the corollary directly by examining double bubbles for which the region of area A_1 is a circle and the other two regions are bands.

Lemma 5.4.17. *If two closed geodesics with lengths L_1, L_2 meet at an angle $\phi \in (0, \pi/2]$, then $L_1 + L_2 \geq 1 + 2\cos\phi$. If $L_1, L_2 \in [1, 3/2]$, then $\phi \geq \cos^{-1}(7/9)$.*

Proof. The torus can be naturally viewed as the plane with various points identified. In the plane, choose a point P that represents the intersection of the two closed geodesics. Draw a segment $\overline{PQ}$ of length L_1 representing the corresponding closed geodesic (on the torus). Similarly, draw a segment $\overline{PR}$ of length L_2 representing the other closed geodesic, such that segments $\overline{PQ}$ and $\overline{PR}$ meet at an angle ϕ.

Because Q and R represent the same point on the torus, $\overline{QR}$ represents a closed geodesic with length at least 1. Thus, applying the Law of Cosines to triangle PQR,

$$1 \le QR^2 = L_1^2 + L_2^2 - 2L_1L_2\cos\phi = (L_1 + L_2)^2 - (2 + 2\cos\phi)L_1L_2.$$

From $0 \le (L_1 - 1)(L_2 - 1)$, we derive $L_1 + L_2 - 1 \le L_1L_2$, so that the above inequality becomes

$$1 \le (L_1 + L_2)^2 - (2 + 2\cos\phi)(L_1 + L_2 - 1)$$

or

$$0 \le (L_1 + L_2 - 1 - 2\cos\phi)(L_1 + L_2 - 1).$$

Because $L_1 + L_2$ is greater than 1, it must be greater than or equal to $1 + 2\cos\phi$, as desired.

Suppose now that $L_1, L_2 \in [1, 3/2]$. For a fixed ϕ, $L_1^2 + L_2^2 - 2L_1L_2\cos\phi$ is convex in each variable L_1, L_2, so it is maximized for some (L_1, L_2) with $L_1, L_2 \in \{1, 3/2\}$. For the four pairs (L_1, L_2) satisfying this condition, $L_1^2 + L_2^2 - 2L_1L_2\cos\phi < 1$ when $\cos\phi > 7/9$. Therefore, $\phi \ge \cos^{-1}(7/9)$. $\qquad\square$

Remark 5.4.18. The next result, Proposition 5.4.19, generalizes a result about "taxicab distances" in the plane. The taxicab distance between any two points, defined with respect to two given orthogonal directions, is the minimum length of a piecewise linear path between them that travels only along the given orthogonal directions (analogous to the streets a taxicab travels on). The length of a segment between two points is at least $\frac{1}{\sqrt{2}}$ times the taxicab distance between them. (Setting $\phi = \pi/2$ in Proposition 5.4.19 gives this estimate.) Given two lines that are *not* orthogonal (to take an extreme example, lines parallel to two short directions on the hexagonal torus), it is useful to have an analogous estimate. We derive such an estimate relating length to the sum of two taxicab distances, one corresponding to each of the two given lines.

Proposition 5.4.19. *Given two pairs of orthogonal lines in the plane, suppose that each line in the first pair meets each line in the second pair at an angle ϕ or $\pi/2 - \phi$, with $\phi \in [\pi/4, \pi/2]$. For any segment γ, let $T(\gamma)$ be the sum of the lengths of the projections of γ onto the four lines. Then*

$$T(\gamma) \le 2(\sin(\phi/2) + \cos(\phi/2))\,\mathrm{Length}(\gamma).$$

Proof. Without loss of generality, assume that one pair of lines are the coordinate axes, and that the other lines lie clockwise angles ϕ from the positive x- and y-axes. Without loss of generality, further assume that γ has one endpoint at the origin and points into the first quadrant of the plane.

Let $\mathbf{u}$ be a vector parallel to γ, whose magnitude equals the length of γ. Let $\mathbf{v}_1$, $\mathbf{w}_1$, $\mathbf{v}_2$, $\mathbf{w}_2$ be unit vectors pointing along the four given lines, chosen as in Figure 24. Then

$$T(\gamma) = \mathbf{u} \cdot (\mathbf{v}_1 + \mathbf{w}_1 + \mathbf{v}_2) + |\mathbf{u} \cdot \mathbf{w}_2|.$$

If $\mathbf{u} \cdot \mathbf{w}_2$ is negative, then

$$T(\gamma) = \mathbf{u} \cdot (\mathbf{v}_1 + \mathbf{w}_1 + \mathbf{v}_2 - \mathbf{w}_2) \le \|\mathbf{u}\|\|\mathbf{v}_1 + \mathbf{w}_1 + \mathbf{v}_2 - \mathbf{w}_2\|,$$

with equality when $\mathbf{u}$ is parallel to $\mathbf{v}_1 + \mathbf{w}_1 + \mathbf{v}_2 - \mathbf{w}_2$. When this is true, $\mathbf{u}$ bisects the angle formed by $\mathbf{v}_1$ and $\mathbf{v}_2$, and it is easy to compute $\mathbf{u} \cdot \mathbf{v}_1 = \mathbf{u} \cdot \mathbf{v}_2 = \|\mathbf{u}\|\cos(\phi/2)$ and $\mathbf{u} \cdot \mathbf{w}_1 = -\mathbf{u} \cdot \mathbf{w}_2 = \|\mathbf{u}\|\sin(\phi/2)$. Hence, when $\mathbf{u} \cdot \mathbf{w}_2$ is negative,

$$T(\gamma) \le \|\mathbf{u}\|(2\cos(\phi/2) + 2\sin(\phi/2)) = 2(\cos(\phi/2) + \sin(\phi/2))\mathrm{Length}(\gamma).$$
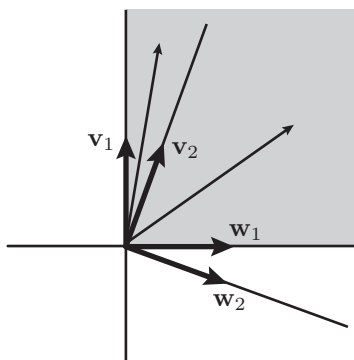
FIGURE 24. If a vector in the first quadrant has fixed magnitude, then the sum of its projections onto the four directions $\mathbf{v}_i$, $\mathbf{w}_i$ drawn above attains a local maximum in two places: when it bisects the angle between $\mathbf{v}_1$ and $\mathbf{v}_2$, and when it bisects the angle between $\mathbf{w}_1$ and $\mathbf{v}_2$.

Similarly, if $\mathbf{u} \cdot \mathbf{w}_2$ is nonnegative, then $T(\gamma)/\|\mathbf{u}\|$ is maximized when $\mathbf{u}$ bisects the angle formed by $\mathbf{w}_1$ and $\mathbf{v}_2$. We then find that

$$
\begin{aligned}
T(\gamma) & \leq 2(\cos(\pi/4 - \phi/2) + \sin(\pi/4 - \phi/2))\text{Length}(\gamma) \\
& \leq 2(\cos(\phi/2) + \sin(\phi/2))\text{Length}(\gamma),
\end{aligned}
$$

where the second inequality is true because $\phi \in [\pi/4, \pi/2]$. $\qquad \square$

Lemma 5.4.20. *If* $\cos^{-1}(7/9) \leq \phi \leq \pi/2$, *then*

$$
\frac{\max\{2, 1 + 2\cos\phi\}(1 + \sin\phi) - 3}{\sin\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2}\right) + \cos\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2}\right) - 1} \geq 2.
$$

Proof. The denominator of the left-hand side is at most $\sqrt{2} - 1$. For $\cos^{-1}(7/9) \leq \phi \leq \pi/4$, we have $\cos\phi \geq \sqrt{2}/2$ and $\sin\phi > 2 - \sqrt{2}$. Hence, the numerator of the left-hand side is greater than $2\sqrt{2} - 2$. Therefore, the left-hand side is greater than 2 when $\phi \leq \pi/4$.

For $\pi/4 \leq \phi \leq \pi/3$, $\max\{2, 1 + 2\cos\phi\} = 1 + 2\cos\phi$ and $\max\{\phi, \pi/2 - \phi\} = \phi$. It is easy to check that on the interval $\phi \in [\pi/4, \pi/3]$, the numerator of the left-hand side is decreasing while the denominator is increasing. Because the inequality holds when $\phi = \pi/3$ (with equality), it holds along the entire interval.

For $\pi/3 \leq \phi \leq \pi/2$, $\max\{2, 1 + 2\cos\phi\} = 2$ and $\max\{\phi, \pi/2 - \phi\} = \phi$. Clearing the denominator of the desired inequality, then simplifying, yields the following inequalities (each equivalent to the desired inequality):

$$
\begin{aligned}
2\sin\phi - 1 & \geq 2\sin(\phi/2) + 2\cos(\phi/2) - 2, \\
4\sin(\phi/2)\cos(\phi/2) - 1 & \geq 2\sin(\phi/2) + 2\cos(\phi/2) - 2, \\
(2\sin(\phi/2) - 1)(2\cos(\phi/2) - 1) & \geq 0.
\end{aligned}
$$

The last inequality holds for $\pi/3 \leq \phi \leq \pi/2$. $\qquad \square$

Lemma 5.4.21. *Let $\mathcal{B}_1$ and $\mathcal{B}_2$ be the boundaries of two bands with distinct homologies, where $\mathcal{B}_1$ and $\mathcal{B}_2$ may overlap each other nontrivially (see Figure 25). If the length of $\mathcal{B}_1 \cup \mathcal{B}_2$ does not exceed 3, then the length of $\mathcal{B}_1 \cap \mathcal{B}_2$ is greater than or equal to 2.*
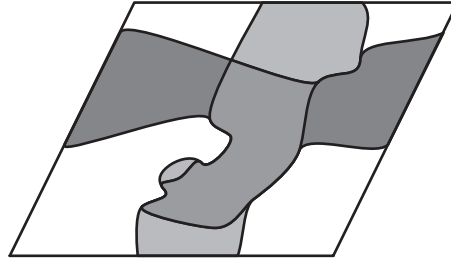

FIGURE 25. Given two bands with overlapping boundary, we can project the boundaries onto various directions to show that the length of the overlap is large. (Above, the bands are filled with dark and light shades, with a medium shade along their intersection.)

Proof. Assume that the length of $\mathcal{B}_1 \cup \mathcal{B}_2$ is as at most 3, and let $\gamma_1, \ldots, \gamma_k$ be the components of $\mathcal{B}_1 \cap \mathcal{B}_2$. For $i = 1, 2$, let L_i be the length of the closed geodesic g_i with the same homology as the components of $\mathcal{B}_i$. Observe that $L_i \leq 3/2$, because otherwise $\mathcal{B}_i$ would have length greater than 3. Let g_1 and g_2 meet each other at an angle $\phi \in (0, \pi/2]$.

Below, we will take each closed curve in $\mathcal{B}_1$ and project it onto two directions — onto the direction of g_1 (Figure 26(a)) and onto the direction orthogonal to g_2 (Figure 26(b)). We will also take similar projections of the components of $\mathcal{B}_2$. By analyzing these projections' effects on various pieces of $\mathcal{B}_1 \cup \mathcal{B}_2$, we will obtain our desired inequality.

For each γ_i, let $T(\gamma_i)$ be the sum of the projection of γ_i onto four directions: the two directions of the closed geodesics g_1 and g_2, and the two directions orthogonal to these closed geodesics.

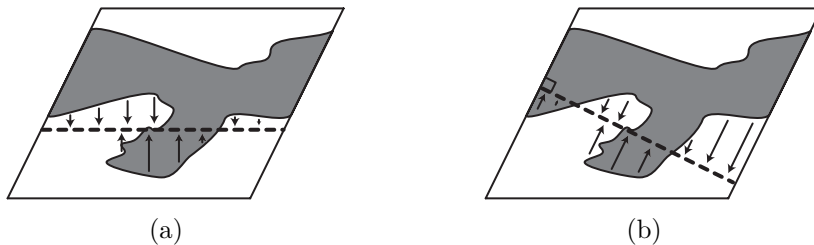(a)   (b)

FIGURE 26. By projecting each component of the boundary of the bands in two directions, we bound the length of the intersection of the two boundaries.

Take a component of $\mathcal{B}_1$, and project it onto the direction of g_1 and onto the direction orthogonal to g_2. The first projection has length at least L_1, and the second projection has length at least $L_1 \sin \phi$. Project the other component of $\mathcal{B}_1$ in a similar fashion.

Likewise, project each component of $\mathcal{B}_2$ onto the direction of g_2 and onto the direction orthogonal to g_1. In all, the projections of the components of $\mathcal{B}_1$ and $\mathcal{B}_2$ have total length greater than or equal to

$$2(L_1 + L_2)(1 + \sin \phi).$$

Each component of $\mathcal{B}_1 \cup \mathcal{B}_2 - \mathcal{B}_1 \cap \mathcal{B}_2$ is projected exactly twice, so they contribute at most

$$2 \cdot \text{Length}(\mathcal{B}_1 \cup \mathcal{B}_2 - \mathcal{B}_1 \cap \mathcal{B}_2)$$

to the total length of the projections. Each γ_i, on the other hand, contributes at most $T(\gamma_i)$. Therefore,

$$
\begin{aligned}
2(L_1 + L_2)(1 + \sin \phi) \quad &\leq \quad 2 \cdot \text{Length}(\mathcal{B}_1 \cup \mathcal{B}_2 - \mathcal{B}_1 \cap \mathcal{B}_2) + \sum_{i=1}^{k} T(\gamma_i) \\
&\leq \quad 2\left(3 - \sum_{i=1}^{k} \text{Length}(\gamma_i) \right) + \sum_{i=1}^{k} T(\gamma_i).
\end{aligned}
$$

We estimate the two sides of this inequality as follows. First, $L_1 + L_2 \geq 2$, and by Lemma 5.4.17, $L_1 + L_2 \geq 1 + 2\cos \phi$. Hence,

$$2 \max\{2, 1 + 2\cos \phi\}(1 + \sin \phi) \leq 2(L_1 + L_2)(1 + \sin \phi).$$

Second, from the two pairs of orthogonal directions we project onto, we can choose one direction from each pair so that they meet at an angle of $\max\{\phi, \pi/2 - \phi\} \in [\pi/4, \pi/2]$. Applying Proposition 5.4.19 then shows that

$$2\left(3 - \sum_{i=1}^{k} \text{Length}(\gamma_i) \right) + \sum_{i=1}^{k} T(\gamma_i)$$

is at most

$$6 + \left(2\sin\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2} \right) + 2\cos\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2} \right) - 2 \right) \sum_{i=1}^{k} \text{Length}(\gamma_i).$$

Therefore, to prove that $\sum_{i=1}^{k} \text{Length}(\gamma_i) \geq 2$, it suffices to prove that

$$\frac{\max\{2, 1 + 2\cos \phi\}(1 + \sin \phi) - 3}{\sin\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2} \right) + \cos\left(\frac{\max\{\phi, \pi/2 - \phi\}}{2} \right) - 1} \geq 2.$$

By Lemma 5.4.17, $\cos^{-1}(7/9) \leq \phi$; also, by definition, $\phi \leq \pi/2$. By Lemma 5.4.20, the above inequality holds for such ϕ. $\qquad \square$

Proposition 5.4.22. *No octagon-square tiling is perimeter minimizing.*

Proof. Let $\mathcal{B}_1$ be the boundary of the band enclosing one eight-sided region and one four-sided component (Figure 27(a)). Let $\mathcal{B}_2$ be the boundary of the band enclosing the same eight-sided region and the other four-sided component (Figure 27(b)). If the length of $\mathcal{B}_1 \cup \mathcal{B}_2$ is greater than 3, then by the perimeter bound (Proposition 3.8), the octagon square is not perimeter minimizing.

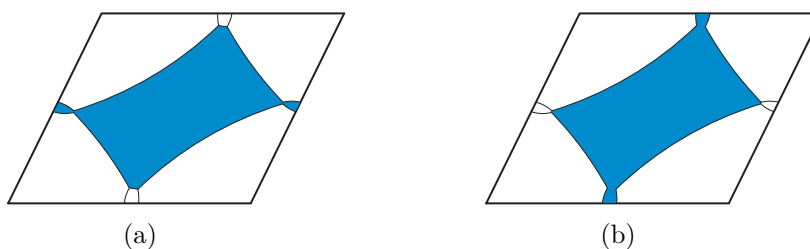(a)   (b)

FIGURE 27. Given the above octagon-square tiling, the boundaries of the two colored bands either have a union with length greater than three, or an intersection with length greater than or equal to two. In both cases, it can be shown that the tiling is not perimeter minimizing.

Otherwise, by Lemma 5.4.21, the length of $\mathcal{B}_1 \cap \mathcal{B}_2$ is at least 2. Let A_1 be the area of the region containing the four-sided components. By the isoperimetric inequality for the plane, the perimeter of this region is at least $2\sqrt{\pi A_1}$. Therefore, the perimeter of the octagon-square tiling is at least $2 + 2\sqrt{\pi A_1}$. By Corollary 5.4.15, the tiling is not a minimizer. □

Proposition 5.4.23. *The only tiling that can be perimeter minimizing is the standard hexagon tiling.*

Proof. By Corollary 5.2.9, the tiling must be a hexagon tiling or an octagon-square tiling. By Proposition 5.4.22, no octagon-square tiling is perimeter minimizing. By Proposition 5.4.13, the only minimizing hexagon tiling is the standard hexagon tiling. □

6. PERIMETER-MINIMIZING DOUBLE BUBBLES ON THE FLAT TWO-TORUS

The Main Theorem characterizes the perimeter-minimizing double bubbles on all flat two-tori, as depicted in Figure 1. Corollaries 6.3 and 6.5 characterize the minimizing double bubbles on the flat infinite cylinder and the flat infinite strip with free boundary. Remark 6.7 makes a conjecture about the double bubble problem on the closely related Klein bottle.

Main Theorem 6.1. *A perimeter-minimizing double bubble on a flat two-torus is of one of the following five types (perhaps after relabelling the three regions), depending on the areas to be enclosed:*

   (i) *the standard double bubble,*
   (ii) *a standard chain,*
  (iii) *the band lens,*
  (iv) *the double band, or*
   (v) *the standard hexagon tiling (only for the hexagonal torus).*

Proof. Proposition 4.6 shows that a minimizer must be a contractible double bubble, a tiling, a swath, a band adjacent to a set of components whose union is contractible, or the double band (after perhaps relabelling the two regions and the exterior).

By Proposition 5.1.2, the standard double bubble is the only possible contractible minimizer. By Proposition 5.2.15, a standard chain is the only possible minimizer

among swaths. By Proposition 5.3.5, the band lens is the only possible minimizer among bands adjacent to a set of components whose union is contractible. By Proposition 5.4.23, the standard hexagon tiling (on the hexagonal torus) is the only possible minimizing tiling.   □

Remark 6.2. Using the phase portrait computations of Section 8, it can be confirmed that all five types of minimizers do occur.

Corollary 6.3. *A perimeter-minimizing double bubble on the flat infinite cylinder is of one of the following four types, depending on the areas to be enclosed:*

   (i) *the standard double bubble,*
  (ii) *a standard chain,*
 (iii) *the band lens, or*
 (iv) *the double band.*

Proof. Without loss of generality, assume that the cylinder has circumference 1. Let Σ be a perimeter-minimizing double bubble on the cylinder that is not of one of the four types listed above. Because Σ has finite perimeter, it can be placed on some long rectangular torus (for example, one with length greater than 3). Then by Theorem 6.1, there exists a double bubble Σ' on the torus enclosing the same areas as Σ but with less perimeter. Because the torus has length greater than 3, and because Σ' has perimeter less than or equal to 3, it is contractible or wraps around a short direction of the torus. Therefore, Σ' can be placed on the cylinder, and Σ cannot be perimeter minimizing.   □

Corollary 6.4. *Every minimizing double bubble on the infinite cylinder has a line of reflective symmetry that is perpendicular to a shortest closed geodesic, or is congruent to a double bubble with this property.*

Proof. Given a minimizing standard double bubble, it can be rotated so that its axis of symmetry is perpendicular to a shortest closed geodesic. Each of the other minimizers in Corollary 6.3 has a line of reflective symmetry with the required property. (A standard chain on the infinite cylinder has a line of reflective symmetry that is parallel to a shortest closed geodesic, but it also has two lines of reflective symmetry perpendicular to a shortest closed geodesic.)   □

Corollary 6.5. *Any minimizing double bubble on the flat infinite strip with free boundary is the quotient of a minimizing double bubble on an infinite cylinder (with twice the width of the strip) by a reflection across one of its lines of symmetry, using the free boundary of the strip to complete the double bubble, as in Figure 28.*
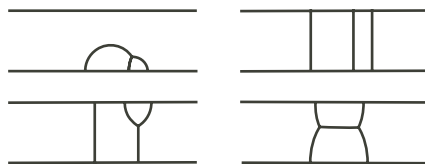

FIGURE 28. The four perimeter-minimizing double bubbles on the flat infinite strip with free boundary are each half of some perimeter-minimizing double bubble on the infinite cylinder.

Proof. The infinite strip is the quotient space of the infinite cylinder by some reflection. Hence we can think of the infinite cylinder as a double cover of the infinite strip. Lifting a minimizer Σ on the strip yields a double bubble Σ' on the cylinder with twice the perimeter and enclosed areas as Σ. If Σ' is a minimizing double bubble, then we are done. Otherwise, by Corollary 6.4, there exists a minimizer Γ' that encloses the same areas as Σ' and that has at least one line of reflective symmetry perpendicular to a shortest closed geodesic. Quotienting by the corresponding reflection yields the original infinite strip with free boundary, along with a new double bubble Γ enclosing the same areas as Σ with less perimeter, a contradiction.   □

Remark 6.6. For finite rectangular cylinders with free boundary (i.e., the quotient space of a rectangular torus by some reflection), a result analogous to Corollary 6.5 holds if the distance between the two components of the boundary is sufficiently small with respect to the length of each component of the boundary.

Remark 6.7 (Klein bottles). We conjecture that perimeter-minimizing double bubbles on the flat Klein bottle are of virtually the same form as on the flat torus. Like the flat torus, we can represent the flat Klein bottle as a planar parallelogram with opposite sides identified and no angle less than 60 degrees, except that we identify the top and bottom sides with a flip. Because of this flip, on a Klein bottle with reasonably long top and bottom, any minimizing double band is bounded by two closed geodesics each of length 2, and has total perimeter 4 (instead of 3). This is an important difference, and for proofs on the torus where we required a perimeter bound of 3, it will be necessary to modify our approach to accommodate the higher bound for the Klein bottle. Note also that the placement of perimeter-minimizing double bubbles on a wide Klein bottle is less flexible than on a torus. The standard chain, band lens, and double band must all be oriented symmetrically about one of the two unique shortest closed geodesics of a wide Klein bottle (which are always vertical). However, all of the conjectured minimizers (including the standard hexagon tiling) do remain well defined, so long as they are placed carefully.

We also conjecture that the perimeter-minimizing double bubbles on the infinite cylinder (which can be thought of as an infinitely long rectangular torus) carry over in the analogous space corresponding to the Klein bottle, that is, a Möbius strip of infinite length — the surface between two horizontal lines in the plane, where the lines are identified with a flip. Here the same comments apply, but the absence of tilings may make the full solution more accessible.

7. Formulas for perimeter and area

This section gives formulas for the areas of the enclosed regions and the perimeters of the four potential minimizers different from the double band: the standard double bubble (Proposition 7.1.1), the standard chain (Propositions 7.2.1, 7.2.2), the band lens (Proposition 7.3.1), and the standard hexagon tiling (Proposition 7.4.1). In addition, Lemmas 7.2.4 and 7.2.5 prove that when a standard chain with axis-length one or a standard hexagon tiling *seems* to be minimizing, based on formal perimeter calculations that we can perform even if the double bubble actually does not exist on a given torus, then the double bubble actually exists and is minimizing. As described in Section 8, we use these formulas and results to produce the phase portraits in Figure 2.

Remark 7.1. We begin by stating a few simple geometric formulas that describe circular arcs meeting chords. Given a circular arc subtended by a chord of length C that it meets at an angle θ, the area between the arc and the chord, the length of the arc, and the radius of the circle are given by

$$(7.1) \qquad A(\theta, C) \quad = \quad \frac{C^2(\theta - \sin\theta\cos\theta)}{4\sin^2\theta},$$

$$(7.2) \qquad L(\theta, C) \quad = \quad \frac{C\theta}{\sin\theta},$$

$$(7.3) \qquad R(\theta, C) \quad = \quad \frac{C}{2\sin\theta}.$$

7.1. Formulas for the standard double bubble.

Proposition 7.1.1 gives formulas for perimeter and area of the standard double bubble. These formulas first appeared in [F, Section 2].

Proposition 7.1.1. *Let R_1 and R_2 be the regions of higher and lower pressure enclosed in a standard double bubble. The area and perimeter formulas of a standard double bubble in terms of the separating chord length C and the angle θ at which it meets the interior arc are*

$$A_{R_1}(\theta, C) \quad = \quad A(2\pi/3 - \theta, C) + A(\theta, C),$$

$$A_{R_2}(\theta, C) \quad = \quad A(2\pi/3 + \theta, C) - A(\theta, C),$$

$$P(\theta, C) \quad = \quad L(2\pi/3 + \theta, C) + L(2\pi/3 - \theta, C) + L(\theta, C).$$

Proof. The arc between R_1 and the exterior meets the separating chord C at $\frac{2\pi}{3} - \theta$, and the arc between R_2 and the exterior meets the separating chord C at $\frac{2\pi}{3} + \theta$. The given formulas follow immediately. □

7.2. Formulas for the standard chain.

To parameterize the standard chain, we distinguish between those in which the enclosed regions have unequal pressure (Proposition 7.2.1) and those in which the enclosed regions have equal pressure (Proposition 7.2.2).

Proposition 7.2.1. *Consider a standard chain with axis-length L_0 and with two components of unequal pressure. Let R_1 and R_2 be the enclosed regions, and suppose that R_1 has greater pressure than R_2. Let C_1 be the length of either chord subtending an arc separating R_1 from the exterior, where the chord and arc meet at (interior) angles θ_1. (See Figure 29.) The area and perimeter formulas are*

$$A_{R_1}(\theta_1, C_1) \quad = \quad 2A(\theta_1, C_1) + 2A(\pi/6 - \theta_1, C_3) + C_1 C_3,$$

$$A_{R_2}(\theta_1, C_1) \quad = \quad 2A(\pi/3 - \theta_1, C_2) - 2A(\pi/6 - \theta_1, C_3) + C_2 C_3,$$

$$P(\theta_1, C_1) \quad = \quad 2L(\theta_1, C_1) + 2L(\pi/3 - \theta_1, C_2) + 2L(\pi/6 - \theta_1, C_3),$$

where C_2 and C_3 are given by

$$C_2 \quad = \quad L_0 - C_1,$$

$$C_3 \quad = \quad \frac{C_1(L_0 - C_1)\sin(\frac{\pi}{6} - \theta_1)}{(L_0 - C_1)\sin\theta_1 - C_1\sin(\frac{\pi}{3} - \theta_1)}.$$
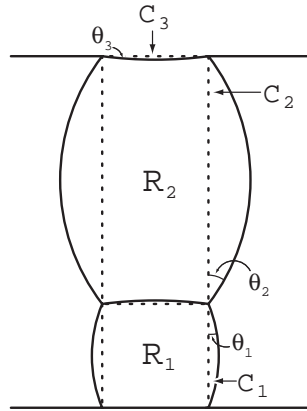
FIGURE 29. This picture illustrates the basic parameters involved in computing perimeter and area for a standard chain.

A standard chain with these parameters exists on an infinite cylinder of circumference L_0 if and only if $C_1 \in (0, \frac{1}{2}L_0)$ and $\theta_1 \in (\sin^{-1}(C_1/L_0), \pi/6)$.

Proof. Let R_0 be the low-pressure exterior region. By the results of Section 5.2, the vertices of R_1 and R_2 form a rectangle. For $i = 1, 2$, let either arc separating R_i from R_0 have curvature κ_i and radius of curvature $r_i = 1/\kappa_i$, let C_i denote the length of the chord with the same endpoints as the arc, and let θ_i be the (interior) angle the chord and the arc make. Define κ_3, r_3, C_3, and θ_3 similarly with respect to the interface between R_1 and R_2.

Because the arcs in the boundary of the chain meet at $2\pi/3$, we have

$$\theta_2 = \pi/3 - \theta_1 \qquad \text{and} \qquad \theta_3 = \pi/6 - \theta_1.$$

Of course,

$$C_2 = L_0 - C_1.$$

Also, by the Extended Law of Sines,

$$C_i = 2r_i \sin \theta_i \qquad \text{for } i = 1, 2, 3,$$

allowing us to solve for r_1 and r_2 in terms of θ_1 and C_1. By the cocycle condition (as in Proposition 3.1),

$$-\frac{1}{r_1} + \frac{1}{r_2} + \frac{1}{r_3} = -\kappa_1 + \kappa_2 + \kappa_3 = 0,$$

and we may solve for r_3 in terms of r_1 and r_2. These equations allow us to find C_3 in terms of θ_1 and C_1:

$$\begin{aligned} C_3 &= 2\frac{1}{1/r_1 - 1/r_2} \sin \theta_3 = \frac{2r_1 r_2}{r_2 - r_1} \sin \theta_3 \\ &= \frac{C_1 C_2}{C_2 \sin \theta_1 - C_1 \sin \theta_2} \sin \theta_3 = \frac{C_1(L_0 - C_1)\sin(\pi/6 - \theta_1)}{(L_0 - C_1)\sin \theta_1 - C_1 \sin(\pi/3 - \theta_1)}. \end{aligned}$$

It remains to show that $C_1 \in (0, L_0/2)$, $\theta_1 \in (\sin^{-1}(C_1/L_0), \frac{\pi}{6})$, and that a corresponding standard chain exists (on some infinite cylinder) if these bounds hold.

Because $r_2 > r_1$ and $\theta_2 > \theta_1$, we have $L_0 - C_1 = C_2 = 2r_2 \sin \theta_2 > 2r_1 \sin \theta_1 = C_1$, proving the bound on C_1. The upper bound on θ_1 holds because $\pi/6 - \theta_1 = \theta_3 > 0$.
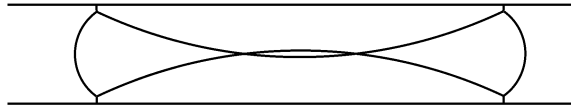

FIGURE 30. If the angles and lengths in the parameterization of a standard chain do not meet certain bounds, then the boundary arcs of the chain intersect each other.
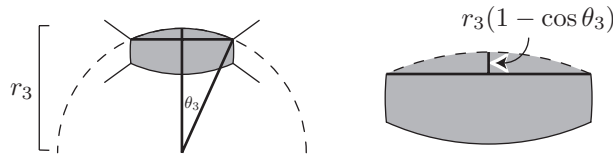

FIGURE 31. The arcs separating the components enclosed in a standard chain cannot intersect each other, yielding bounds on the parameters used to describe standard chains.

Given $C_1 \in (0, L_0/2)$ and $0 < \theta_1 < \pi/6$, a corresponding standard chain exists (on some infinite cylinder) if and only if it can be drawn so that the arcs separating R_1 and R_2 do not intersect. Consider one of these arcs. As shown in Figure 31, the distance between the midpoint of the arc and the chord subtending the arc is $r_3 - r_3 \cos \theta_3$. The two arcs separating R_1 and R_2 do not intersect each other if and only if this distance is less than $\frac{1}{2}C_2$ — that is, if and only if

$$r_3(1 - \cos \theta_3) < \frac{1}{2}C_2.$$

(This inequality fails if and only if the arcs intersect each other as in Figure 30.) Writing r_3, θ_3, and C_2 in terms of C_1 and θ_1, this inequality becomes (after much simplification)

$$C_1 < L_0 \sin \theta_1$$

or $\theta_1 > \sin^{-1}(C_1/L_0)$. □

Proposition 7.2.2. *Consider a standard chain with axis-length L_0 and with two regions R_1 and R_2 of equal pressure. The area and perimeter formulas are*

$$A_{R_1}(C_3) = A_{R_2}(C_3) \quad = \quad \frac{2\pi - 3\sqrt{3}}{24}L_0^2 + \frac{1}{2}L_0 C_3$$

$$P(C_3) \quad = \quad \frac{2\pi}{3}L_0 + 2C_3,$$

where C_3 is the length of the chord separating R_1 from R_2.

Proof. For $i = 1, 2, 3$, define κ_i, r_i, C_i, and θ_i as in the first paragraph of the proof of Proposition 7.2.1. (Figure 29 shows many of these variables, although the standard chain in the figure does not contain two regions of equal pressure.)

Observe that $\kappa_1 = \kappa_2$ and hence $r_1 = r_2$. Also, $\theta_3 = 0$ and $\theta_1 = \theta_2 = \pi/6 - \theta_3 = \pi/6$. Therefore, $C_1 = 2r_1 \sin\theta_1 = 2r_2 \sin\theta_2 = C_2$, implying that $C_1 = C_2 = \frac{1}{2}L_0$ and $r_1 = r_2 = \frac{1}{2}L_0/(2\sin(\pi/6)) = \frac{1}{2}L_0$.

Although C_1 and C_2 are uniquely determined, C_3 is not. Still, because all other arc lengths and angles in the chain are fixed, we have simple formulas expressing the areas and perimeters of such chains in terms of C_3. Given the fixed angles and chord lengths found earlier, we have $A_{R_1}(C_3) = A_{R_2}(C_3) = 2A(\pi/6, \frac{L_0}{2}) + (\frac{1}{2}L_0)C_3$ and $P(C_3) = 4L(\pi/6, \frac{L_0}{2}) + 2C_3$. Using (7.1)–(7.2) to evaluate these expressions yields the desired formulas. $\square$

Remark 7.2.3. The plots based on this parameterization prove the existence of a family of standard chains with fixed axis-length. In particular, we know that there exists a family of standard chains that wrap around a short direction of the torus.

Lemma 7.2.4. *Consider any standard chain on the unit-circumference cylinder enclosing areas A_1, A_2, and any torus with area greater than $A_1 + A_2$. The perimeter of the chain is greater than or equal to that of the minimizing double bubble(s) enclosing areas A_1, A_2 on the torus. If equality holds, the standard chain fits on the torus and is itself a minimizing double bubble on the torus.*

Proof. Consider a fundamental domain of the torus, a parallelogram with side lengths 1 and L and with interior angle $\alpha \in [\pi/3, \pi/2]$. We orient the parallelogram so that the side with length 1 is vertical. Also, let R_1 be a component of higher (or equal) pressure enclosed in the chain, and let R_2 be the other component enclosed in the chain. For $i = 1, 2, 3$ define C_i and θ_i as in the first paragraph of the proof of Proposition 7.2.1.

We immerse the chain in the torus, such that the chain and the side of length 1 have the same homology. It suffices to prove that if the chain is not embedded, then it has more perimeter than some (embedded) minimizer enclosing areas A_1, A_2 on the torus. Because the chain is not embedded in the torus, it does not only wrap around the torus in the direction of the side with length 1; it also wraps around the torus in the direction of the side with length L. We consider two cases: first, where the chain wraps around to intersect itself because the two components intersect each other; second, where instead at least one component intersects itself but not the other component.

First suppose that both components of the chain intersect each other. For $i = 1, 2$, consider either boundary arc of R_i that lies along the boundary of the chain. The horizontal distance between the midpoint of this arc and the chord subtending it is $d_i = (C_i/2)\tan(\theta_i/2)$. Also, the vertices of the standard chain lie on two vertical closed geodesics, and the horizontal distance between these closed geodesics is C_3. Now, because the chain wraps around the torus in a nonvertical direction, the sum of d_1, d_2, and C_3 must exceed the horizontal distance $L\sin\alpha$ between the two vertical sides of the parallelogram. Hence,

$$\frac{C_1}{2}\tan\left(\frac{\theta_1}{2}\right) + \frac{C_2}{2}\tan\left(\frac{\theta_2}{2}\right) + C_3 > L\sin\alpha \geq \sin(\pi/3) = \frac{\sqrt{3}}{2}.$$
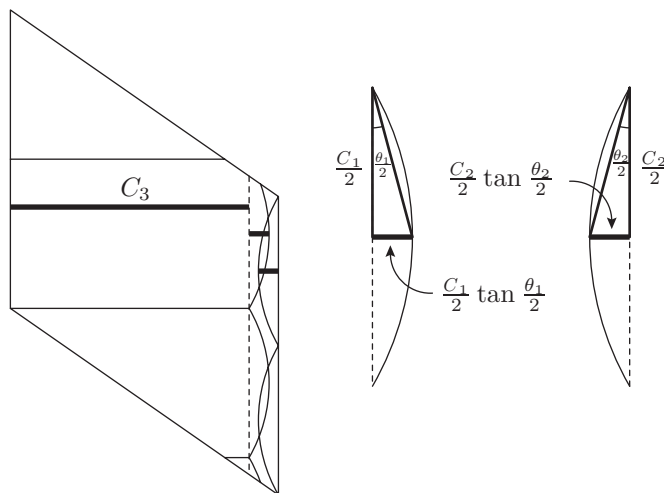
FIGURE 32. One way a standard chain might be nonembedded is if the two components intersect each other. In this case, by considering three segments perpendicular to the short direction with total length at least $\sqrt{3}/2$, it can be shown that the arcs separating the two components in the chain are long enough (greater than .5) to prove that the chain does not have minimizing perimeter.

From our parameterization, $C_1 + C_2 = 1$ and $\theta_2 < \theta_1 < \pi/3$. Therefore,

$$\frac{C_1}{2} \tan\left(\frac{\theta_1}{2}\right) + \frac{C_2}{2} \tan\left(\frac{\theta_2}{2}\right) < \frac{C_1}{2} \tan\left(\frac{\pi}{6}\right) + \frac{C_2}{2} \tan\left(\frac{\pi}{6}\right) = \frac{1}{2} \tan\left(\frac{\pi}{6}\right) = \frac{\sqrt{3}}{6}.$$

Combining these two inequalities, we find that

$$C_3 > \frac{\sqrt{3}}{2} - \frac{\sqrt{3}}{6} = \frac{\sqrt{3}}{3} > \frac{1}{2}.$$

Hence, each of the arcs separating the two components enclosed in the chain has length greater than $\frac{1}{2}$. It easily follows that the perimeter of the chain is greater than 3. Therefore, by the perimeter bound (Proposition 3.8), the standard chain has perimeter greater than that of the (embedded) minimizer(s) enclosing the same areas.

Next suppose that one of the components — say, R_i — intersects itself but not the other component. The area enclosed by R_i on the torus thus forms a band wrapping around the torus in a nonvertical direction; let $\mathcal{B}_1$ be the boundary of this band. Also, let $\mathcal{B}_2$ be the boundary of a second band containing the two components of the chain (and excluding any area that lies in the overlap of the immersed chain). By Lemma 5.4.21, either the length of $\mathcal{B}_1 \cup \mathcal{B}_2$ is greater than 3, or the length of $\mathcal{B}_1 \cap \mathcal{B}_2$ is greater than or equal to 2.

If the length of $\mathcal{B}_1 \cup \mathcal{B}_2$ is greater than 3, then by the perimeter bound (Proposition 3.8), the perimeter of the standard chain has perimeter greater than that of the (embedded) minimizer(s) enclosing the same areas.

If instead the length of $\mathcal{B}_1 \cap \mathcal{B}_2$ is greater than or equal to 2, then let R_j be the region in the chain different from R_i, and let A_j be the area of this region. Then
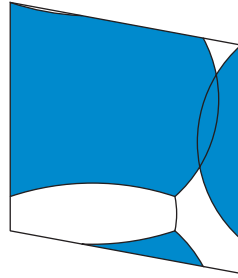
FIGURE 33. A second way a standard chain might be nonembed-
ded is if at least one component intersects itself. The component
then forms a band wrapping around the torus, implying (as was the
case for octagon-square tilings) that the chain has a large perime-
ter.

$\mathcal{B}_1 \cap \mathcal{B}_2$ does not intersect the boundary of R_j. By the isoperimetric inequality for
the plane, the perimeter of the region is at least $2\sqrt{\pi A_j}$. Therefore, the perimeter
of the standard chain is at least $2 + 2\sqrt{\pi A_j}$. By Corollary 5.4.15, this exceeds the
perimeter of the minimizer(s) enclosing the same areas. □

7.3. **Formulas for the band lens.** By Lemma 5.3.6, any minimizing band lens
has the same homology as a shortest closed geodesic. Proposition 7.3.1 gives area
and perimeter formulas for these band lenses.

Proposition 7.3.1. *Consider a band lens with the same homology as a shortest
closed geodesic. Let $r = (1/\text{curvature})$ be the radius of curvature of the lens, and
let d be the width of the band. The area and perimeter formulas for the band lens
in terms of r and d are*

$$A_{\text{Lens}}(r) = 2r^2 \left(\frac{\pi}{3} - \frac{\sqrt{3}}{4} \right),$$

$$A_{\text{Band}}(r, d) = d - r^2 \left(\frac{\pi}{3} - \frac{\sqrt{3}}{4} \right),$$

$$P(r) = \left(\frac{4\pi}{3} - \sqrt{3} \right) r + 2.$$

Furthermore, $r < \frac{1}{\sqrt{3}}$ and $d > \frac{r}{2}$.

Proof. The chord between the vertices of the lens, with length $C = 2r \sin(\pi/3) =
\sqrt{3}r$, must fit along a short direction of the torus. Hence, $r < \frac{1}{\sqrt{3}}$. Also, because
the band and lens do not overlap, the width of the band is greater than half the
width of the lens. Hence, $d > r/2$.

Using (7.1)–(7.2), we can compute the area and perimeter formulas:

$$
\begin{aligned}
A_{\text{Lens}}(r) &= 2A\left(\frac{\pi}{3}, \sqrt{3}r\right) = 2r^2\left(\frac{\pi}{3} - \frac{\sqrt{3}}{4}\right), \\
A_{\text{Band}}(r, d) &= d - \frac{1}{2}A_{\text{Lens}}(r) = d - r^2\left(\frac{\pi}{3} - \frac{\sqrt{3}}{4}\right), \\
P(r) &= 2L\left(\frac{\pi}{3}, \sqrt{3}r\right) + (2 - \sqrt{3}r) = \left(\frac{4\pi}{3} - \sqrt{3}\right)r + 2.
\end{aligned}
$$

$\square$

7.4. **Formulas for the standard hexagon tiling.** Corollary 5.4.12 showed that the perimeter of every standard hexagon tiling is equal to three. This section gives area formulas for the standard hexagon tiling (Proposition 7.4.1), proves that at most one standard hexagon tiling encloses two given areas (Lemma 7.4.2), and describes exactly which pairs of areas can be enclosed by a standard hexagon tiling (Lemma 7.4.4).

Proposition 7.4.1. *In any standard hexagon tiling, there exist a, b, c with sum 1 such that one region R_1 is a hexagon with side lengths alternating between a and b, another region R_2 is a hexagon with side lengths alternating between b and c, and the exterior region R_0 is a hexagon with side lengths alternating between c and a. Conversely, given positive a, b, c with sum 1, there exists a corresponding standard hexagon tiling.*

The areas of the three regions are

$$
\begin{aligned}
A_{R_1} &= \frac{\sqrt{3}}{4}(a^2 + 4ab + b^2), \\
A_{R_2} &= \frac{\sqrt{3}}{4}(b^2 + 4bc + c^2), \\
A_{R_0} &= \frac{\sqrt{3}}{4}(c^2 + 4ca + a^2).
\end{aligned}
$$

Proof. Consider the closed curve formed by the four segments p, q, r, s depicted in Figure 34(a). Segments p and r are parallel to a short direction of the torus, and the segment σ connecting the endpoints of the curve is parallel to the same short direction of the torus. Therefore, the segments q and s must veer the same distance away from σ. Because they both meet σ at $2\pi/3$, they have equal length.

Just as q and s must be congruent, *any* two sides of a hexagon must be congruent if they are separated by one other side. It follows that the sides of the hexagons have one of three lengths a, b, c, as depicted in Figure 34(b). By Corollary 5.4.12, the tiling has perimeter three, implying that $a + b + c = 1$. Conversely, given any positive a, b, c such that $a + b + c = 1$, one can construct the corresponding standard hexagon tiling as in Figure 34(b).

To find the area of R_1, extend the sides of R_1 with length a, as in Figure 35, in order to form an equilateral triangle with sides of length $a + 2b$. The new equilateral triangle consists of R_1 and three small equilateral triangles with sides of length b. Therefore, the area of R_1 equals

$$
\frac{\sqrt{3}}{4}(a + 2b)^2 - 3 \cdot \frac{\sqrt{3}}{4}b^2 = \frac{\sqrt{3}}{4}(a^2 + 4ab + b^2).
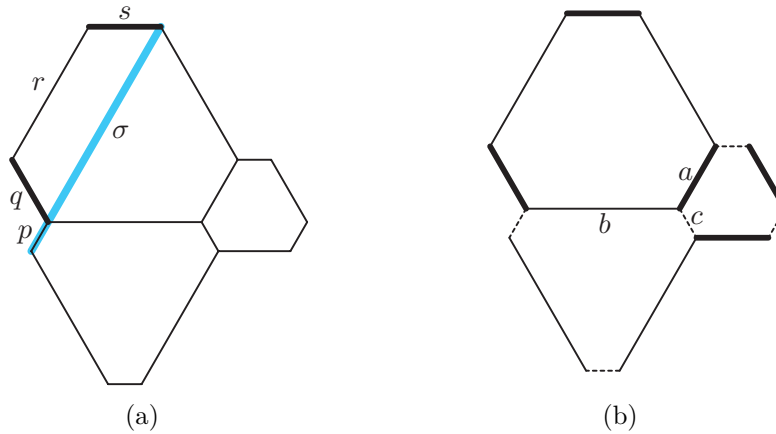$$

FIGURE 34. By examining how far various edges of a standard hexagon tiling veer from shortest closed geodesics, one can show that the sides of each hexagon alternate between two different lengths.
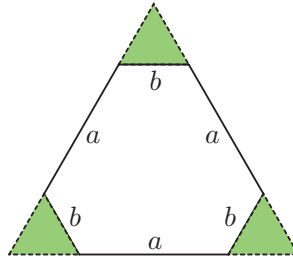

FIGURE 35. If the side lengths of a hexagon with interior angles of $2\pi/3$ alternate between two lengths, one can calculate the area of the hexagon by extending three sides to form an equilateral triangle.

We find the analogous formulas for the areas of R_2 and R_0 in a similar manner. $\quad\square$

Lemma 7.4.2. *Given positive real numbers* A_0, A_1, A_2, *there exists at most one triple* (a, b, c) *of nonnegative numbers with sum* 1 *such that*

$$A_1 = \frac{\sqrt{3}}{4}(a^2 + 4ab + b^2),$$

$$A_2 = \frac{\sqrt{3}}{4}(b^2 + 4bc + c^2),$$

$$A_0 = \frac{\sqrt{3}}{4}(c^2 + 4ca + a^2).$$

Proof. Suppose, for sake of contradiction, that there existed two such triples (a_1, b_1, c_1) and (a_2, b_2, c_2). Without loss of generality, assume that $a_1 < a_2$. Because $a_1^2 + 4a_1b_1 + b_1^2 = a_2^2 + 4a_2b_2 + b_2^2$, it follows that $b_1 > b_2$. Similarly, because

$b_1^2 + 4b_1c_1 + c_1^2 = b_2^2 + 4b_2c_2 + c_2^2$, it follows that $c_1 < c_2$. Likewise, $c_1 < c_2$ implies that $a_1 > a_2$, a contradiction.   □

Corollary 7.4.3. *Given prescribed areas A_1, A_2, there exists at most one standard hexagon tiling enclosing those areas.*

Proof. The result follows from our parameterization (Proposition 7.4.1) and Lemma 7.4.2.   □

Lemma 7.4.4. *Given positive areas A_0, A_1, A_2 with sum $\sqrt{3}/2$ (the area of the hexagonal torus), there exists a standard hexagon tiling enclosing any two of those areas if and only if $\frac{2}{\sqrt[4]{3}}\left(\sqrt{A_j} + \sqrt{A_k}\right) > 1$ for $0 \le j < k \le 2$. That is, if and only if the side length of an equilateral triangle with area A_j and the side length of an equilateral triangle with area A_k have a sum greater than 1. (When equality holds for some pair, say A_1, A_2, a degenerate hexagon tiling exists, consisting of two equilateral triangles enclosing areas A_1, A_2 and with an exterior hexagonal region of area A_0.)*

Proof. Consider the following map from $\{(a,b) \mid 0 \le a, b \text{ and } a+b \le 1\}$ to $\mathbb{R}^2$:

$$(a,b) \mapsto \left(\frac{\sqrt{3}}{4}(a^2 + 4ab + b^2), \frac{\sqrt{3}}{4}(b^2 + 4b(1-a-b) + (1-a-b)^2)\right).$$

By our parameterization (Proposition 7.4.1), this map sends a pair of side lengths of a hexagon tiling to the areas of the region with side lengths a, b and of the region with side lengths $b, 1-a-b$. Thus, as we vary a and b, we obtain all ordered pairs of areas that can be enclosed by a standard hexagon tiling.

By Lemma 7.4.2, the above map is injective. Also, it maps the boundary of its domain continuously to a simple closed curve $\mathcal{C}$ in $\mathbb{R}^2$, as depicted in Figure 36(a). (The figure is drawn with barycentric coordinates, not Cartesian coordinates.) Because the map is injective and continuous, it follows that its range is precisely the closed region bounded by $\mathcal{C}$. This boundary $\mathcal{C}$ consists of three pieces, corresponding to $a = 0, b = 0, c = 0$. When $a = 0$, we have $1 = b + c = \frac{2}{\sqrt[4]{3}}(\sqrt{A_1} + \sqrt{A_0})$; similarly, along the other two pieces of the boundary, we have $1 = \frac{2}{\sqrt[4]{3}}(\sqrt{A_j} + \sqrt{A_k})$ for $(j,k) = (1,2)$ and $(0,2)$. Hence, the interior of the region bounded by these three pieces consists of those area triples (A_0, A_1, A_2) such that $\frac{2}{\sqrt[4]{3}}(\sqrt{A_j} + \sqrt{A_k}) > 1$ for $0 \le j < k \le 2$. Therefore, for each of these triples and no others, some standard hexagon tiling encloses some two of the three areas.   □

Lemma 7.4.5. *If a minimizing double bubble enclosing areas A_1, A_2 on the hexagonal torus has perimeter equal to three, then there exists a (perimeter-minimizing) standard hexagon tiling enclosing those areas.*

Proof. Let the exterior region of the given double bubble have area A_0, and without loss of generality assume that $A_1 \le A_2 \le A_0$. By Lemma 5.4.14,

$$2 + \sqrt{A_1\left(\frac{8\pi}{3} - 2\sqrt{3}\right)}$$

is greater than or equal to the given minimal perimeter, three. The (nonsharp) bound $A_1 \ge \frac{4}{25}$ follows easily. Hence, $A_0 \ge A_2 \ge A_1 \ge \frac{4}{25}$. For $0 \le j < k \le 2$, we

$$(A_1, A_2, A_0) = (0, 0, \tfrac{\sqrt{3}}{2})$$

Area(R_1)

Area(R_2)

$$(A_1, A_2, A_0) = (0, \tfrac{\sqrt{3}}{2}, 0)$$ Area(R_0) $$(A_1, A_2, A_0) = (\tfrac{\sqrt{3}}{2}, 0, 0)$$
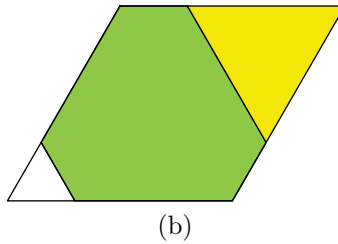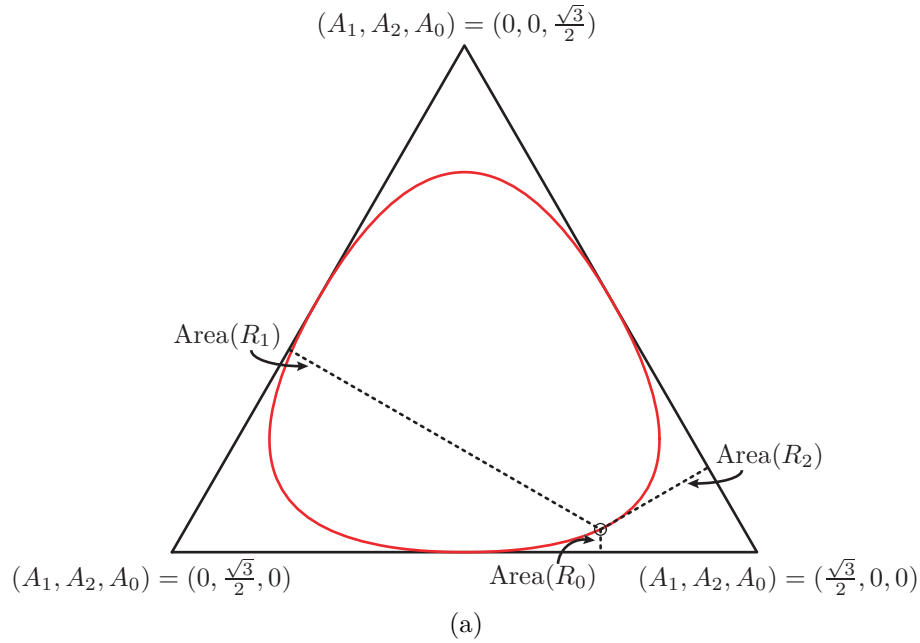
(a)

(b)

FIGURE 36. As in the phase portraits in Figure 2, the triangle in (a) represents the possible area pairs (A_1, A_2) on the hexagonal torus, where the areas of the three regions are given by the distances to the three sides. The region bounded by the curve represents exactly those area pairs enclosed by some standard hexagon tiling. The curve itself represents area pairs enclosed by degenerate hexagon tilings consisting of a hexagon and two triangles, such as the one depicted in (b).

thus have

$$\sqrt{A_j} + \sqrt{A_k} \geq \frac{4}{5} \geq \frac{\sqrt[4]{3}}{2}.$$

By Lemma 7.4.4, there exists a standard hexagon tiling enclosing areas A_1, A_2. By Corollary 5.4.12, the tiling has perimeter three, so it is perimeter minimizing.   □

8. Phase portraits

The phase portraits of Figure 2 show which of the minimizers of Theorem 6.1 win for which pairs of prescribed areas on four different tori — where we represent each torus as a parallelogram with side lengths $1, L$ and interior angle θ. To create these phase portraits, we plot perimeter as a function of area pairs for each candidate in Theorem 6.1 using the parameterizations from Section 7. For each candidate in Theorem 6.1 besides the standard chain, we know that at most one double bubble exists enclosing any given pair of areas, so we can solve for its perimeter numerically. Because we do not know that at most one standard chain (of a fixed homology) exists enclosing two given areas, we instead create the perimeter-areas plot for the standard chain parametrically. Also, using Lemma 4.2, it is possible to show that any minimizing standard chain has axis-length 1, L, or $\sqrt{1 + L^2 \pm 2\angle \cos \theta}$, so we do not plot perimeter as a function of area pairs for other types of standard chains.
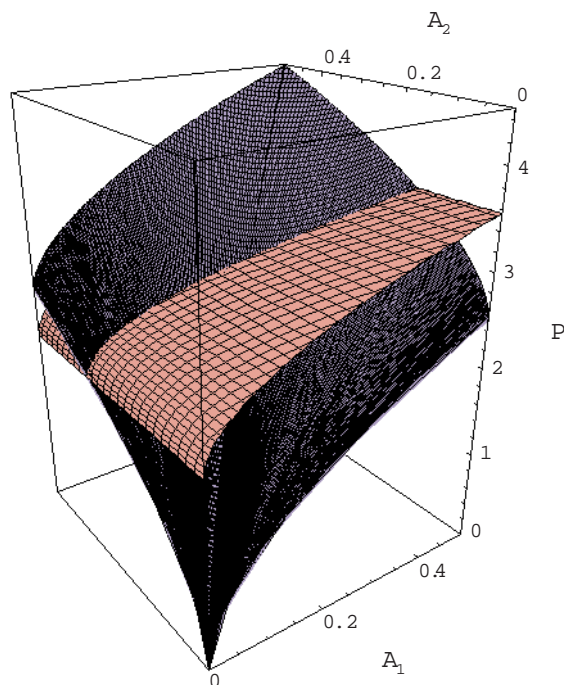
FIGURE 37. This *Mathematica* plot of the intersection of the perimeter plots for the standard double bubble and band lens illustrates the process used to create the phase portraits in Figure 2. The plots suggest that perimeters do not fluctuate wildly along certain sets of unusual area pairs, increasing our confidence that the phase portraits are accurate.

After plotting all these perimeter-areas plots simultaneously, we then view the combined plot from below so that at each area pair the plot(s) of the minimal calculated perimeter is visible. Figure 37 illustrates a simplified version of this
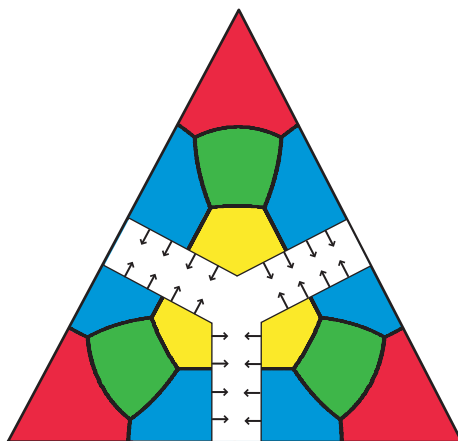
FIGURE 38. Removing pieces of this phase portrait for a rectangular torus with length 1.2 yields the phase portrait plot for some torus with less area.

process, comparing the perimeters of only two candidates, the standard double bubble and the band lens.

This process works if whenever the perimeter calculated for a given double bubble is minimizing for two given areas, the double bubble actually exists on the torus for those areas. Indeed, this is the case. When the standard double bubble does not fit on the torus, it has diameter greater than 1, so by Propositions 5.1.1 and 3.8 it cannot have minimizing perimeter. Lemma 5.4.14 shows that when the calculated perimeter for the band lens is minimizing, that band lens exists. By Lemma 7.2.4, when a standard chain with axis-length one has minimizing perimeter, it can be embedded on the torus; and according to the plots, no standard chain with axis-length greater than one has minimizing perimeter. As the simplest case, we know that the double band exists for all area pairs. Finally, by Lemma 7.4.5, when the minimal perimeter on the hexagonal torus is 3 for two given areas, there exists a standard hexagon tiling enclosing those two areas.

The following conjecture describes a simple connection *among* the various phase portraits. As usual, we assume that the shortest geodesics on the torus and the cylinder have unit length.

Conjecture 8.1. *Given areas $A_0 \geq A_2 \geq A_1$ and a torus with area $A_0 + A_1 + A_2$, the nontiling minimizing double bubble(s) enclosing A_1, A_2 on this torus are the same as the minimizing double bubble(s) enclosing A_1, A_2 on the infinite cylinder. Consequently the phase portrait of any flat two-torus can be obtained from the phase portrait of any flat two-torus of larger area (or from the phase portrait of the infinite cylinder) as in Figure 38. In particular, tori of equal areas have identical phase portraits.*

To prove Conjecture 8.1 it would suffice to show that every minimizing standard chain has axis-length one. The proof would also use the fact (which we can show using Theorem 6.1) that in a minimizer, a region of greater pressure has no more

area than a region of less pressure, as well as Proposition 5.1.1, Lemma 5.3.6, and Lemma 7.2.4.

Acknowledgements

References

[CCWB] Miguel Carrión Álvarez, Joseph Corneli, Geneveive Walsh, and Shabnam Beheshti, *Double bubbles in the three-torus*, Exp. Math., **12** (2003), 79-89.

[FT1]   László Fejes-Tóth, *Lagerungen in der Ebene auf der Kugel und im Raum*, in *Die Grundlehren der Math. Wiss.*, Vol. 65, Springer-Verlag, Berlin, 1953. MR **15:**248b

[FT2]   ———, *Regular Figures*, A Pergamon Press Book, Macmillan Co., New York, 1964. MR **29:**2705

[FT3]   ———, *Über das kürzeste Kurvennetz, das eine Kugeloberfläche in flächengleiche konvexe Teile zerlegt*, Math. Naturwiss. Anz. Ungar. Akad. Wiss., **62** (1943), 349-354. MR **9:**460d

[F]   Joel Foisy, Manual Alfaro, Jeffrey Brock, Nickelous Hodges, and Jason Zimba, *The standard double soap bubble in $\mathbf{R}^2$ uniquely minimizes perimeter*, Pac. J. Math., **159** (1993), 47-59. MR **94b:**53019

[H]   Thomas C. Hales, *The honeycomb conjecture*, Discr. Comput. Geom., **25** (2001), 1-22. MR **2002a:**52020

[HHM]   Hugh Howards, Michael Hutchings, and Frank Morgan, *The isoperimetric problem on surfaces*, Amer. Math. Monthly, **106** (1999), 430-439. MR **2000i:**52027

[HMRR] Michael Hutchings, Frank Morgan, Manuel Ritoré, and Antonio Ros, *Proof of the double bubble conjecture*, Ann. Math., **155** (2002), 459-489. Research announcement of same title, Electron. Res. Announc. Amer. Math. Soc., **6** (2000), 45-49. MR **2003c:**53013

[Ma1]   Joseph D. Masters, *The double bubble on the flat torus*, Unpublished notes on file with F. Morgan, Williams College, 1994.

[Ma2]   ———, *The perimeter-minimizing enclosure of two areas in $\mathbf{S}^2$*, Real Analysis Exchange, **22** (1996/1997), 645-654. MR **99a:**52010

[M1]   Frank Morgan, *Geometric Measure Theory: A Beginner's Guide*, 3rd ed., Academic Press, San Diego, 2000. MR **2001j:**49001

[M2]   ———, *Soap bubbles in $\mathbf{R}^2$ and in surfaces*, Pac. J. Math., **165** (2) (1994), 347-361. MR **96a:**58064

[MW]   Frank Morgan and Wacharin Wichiramala, *The standard double bubble is the unique stable double bubble in $\mathbf{R}^2$*, Proc. AMS, **130** (2002), 2745-2751. Available on the web at `http://www.ams.org/journal-getitem?pii=S0002-9939-02-06640-6`. MR **2003c:**53016

[RHLS] Ben W. Reichardt, Cory Heilmann, Yuan Y. Lai, and Anita Spielman, *Proof of the double bubble conjecture in* $\mathbf{R}^4$ *and certain higher dimensional cases*, Pac. J. Math., **208** (2003), 347-366.

[W]   Wacharin Wichiramala, *The planar triple bubble problem*, Ph.D. thesis, University of Illinois at Urbana-Champaign, 2002.

C/O Frank Morgan, Department of Mathematics and Statistics, Williams College, Williamstown, Massachusetts 01267 – and – Department of Mathematics, University of Texas, Austin, Texas 78712

E-mail address: Frank.Morgan@williams.edu

E-mail address: jcorneli@math.utexas.edu

Department of Mathematics and Statistics, Williams College, Williamstown, Massachusetts 01267

E-mail address: pholt@wso.williams.edu

Department of Mathematics, Harvard University, Cambridge, Massachusetts 02138

E-mail address: lee43@fas.harvard.edu

Department of Mathematics, University of Texas, Austin, Texas 78712

E-mail address: nickleger@mail.utexas.edu

Department of Mathematics and Statistics, Williams College, Williamstown, Massachusetts 01267

E-mail address: eschoenf@wso.williams.edu

Department of Mathematics and Statistics, Williams College, Williamstown, Massachusetts 01267

E-mail address: Benjamin.A.Steinhurst@williams.edu