The Optimal Paper Moebius Band and Finite Tensegrity Calculations

Richard Evan Schwartz *

September 27, 2020

Abstract

There are two main conjectures about paper Moebius bands. First, a smooth embedded paper Moebius band must have aspect ratio at least $\sqrt{3}$. Second, any sequence of smooth embedded paper Moebius bands having aspect ratio converging to $\sqrt{3}$ converges, in the Hausdorff topology and up to isometries, to an equilateral triangle of semi-perimeter $\sqrt{3}$. We will reduce these conjectures to a finite number of statements that certain explicit piecewise-algebraic expressions are non-negative on the unit cubes of dimension 10 and 14. The inequalities involve the geometry of finite tensegrities. I have tested these inequalities extensively though not exhaustively.

1 Introduction

A *paper Moebius band* of aspect ratio λ is a smooth isometric embedding $I : M_\lambda \to \boldsymbol{R}^3$, where M_λ is the flat Mobius band

$$M_\lambda = ([0,1] \times [0,\lambda])/\sim, \qquad (x,0) \sim (1-x,\lambda) \tag{1}$$

An early work [**Sa**] establishes the existence of paper Mobius bands. The main question about these objects is the value of the smallest λ_0 such that a paper Moebius band of aspect ratio λ exists iff $\lambda > \lambda_0$. The paper [**HW**] shows that $\lambda_0 \in [\pi/2, \sqrt{3}]$. The book [**FT**, §14] gives an excellent exposition

*Supported by N.S.F. Grant DMS-1807320

of these bounds. The paper [**CF**] gives a general framework for understanding objects like smooth paper Moebius bands. See [**AHLM**], [**CKS**], [**HF**], [**MK**], [**S1**] for work that is more or less related.

The paper [**HW**] makes the following conjecture about λ_0.

Conjecture 1.1 (Optimality) *A smooth embedded paper Moebius band has aspect ratio greater than $\sqrt{3}$. Hence $\lambda_0 = \sqrt{3}$.*

The upper bound $\lambda_0 \leq \sqrt{3}$ is explained by an example. Figure 1.1 shows an immersed polygonal paper Mobius band of aspect ratio exactly $\sqrt{3}$.

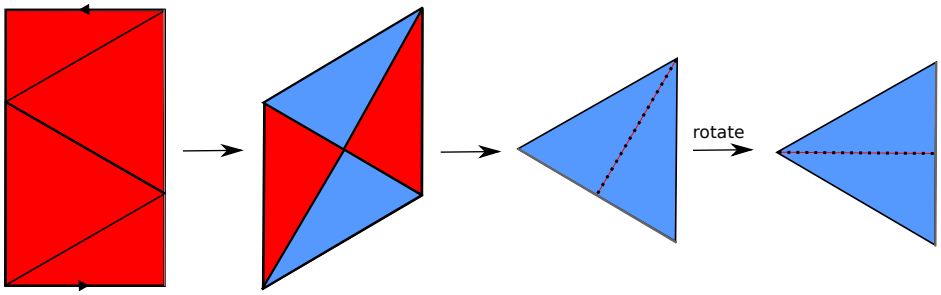

Figure 1.1: The conjectured optimizer

The final image is not embedded. It is an equilateral triangle of semi-perimeter $\sqrt{3}$. However, for any $\epsilon > 0$, one can approximate this map by smooth embeddings of $M_{\lambda+\epsilon}$. See [**FT**] for a discussion about this.

Here is an elaboration on the Optimality Conjecture:

Conjecture 1.2 (Rigidity) *A sequence of smooth embedded paper Mobius bands having aspect ratio converging to $\sqrt{3}$ converges, in the Hausdorff metric and up to isometries, to an equilateral triangle of semi-perimeter $\sqrt{3}$.*

In [**S3**] we show that $\lambda_0 > \sqrt{3} - (1/26)$. The actual bound involves a complicated algebraic number that is just slightly larger than this. For comparison, $\pi/2 < \sqrt{3} - (4/26)$, so our bound gets more than 3/4 of the way to the conjectured optimal bound. We also proved a result [1] which contains Theorem 3.1 as a subset.

In these notes, which are a sequel to [**S3**], we will reduce the Optimality and Rigidity Conjectures to a finite number of statements that certain explicit

[1]The result in Theorem 3.1 has a constant $\sqrt{3} - (1/24)$, which differs slightly from the bound $\sqrt{3} - (1/26)$. This is not a typo. The new constant refers to different objects.

piecewise-algebraic expressions are non-negative on the unit cube $[0,1]^N$, and also statements about where the zeros can occur. Sometimes $N = 10$ and sometimes $N = 14$.

The functions are not all that complicated; mostly they express distances between various pairs of points in space. It seems plausible that a computer-assisted method like the one I used in [**S2**] to solve Thomson's 5-electron problem could work here. However, I am somewhat daunted by the high-dimensionality of the calculations. I will present clear and (in my opinion) strong numerical evidence for all the calculations, but the method I use here (a stochastic hill-climbing algorithm) does not furnish a proof. It merely shows that it is not so easy to find any counterexamples.

In [**FT**] it is pointed out that the main difficulty in proving the conjecture is figuring out how to use the topological hypothesis that the paper Moebius band is embedded. Indeed, in [**FT**], the authors give an example of a sequence of immersed examples with aspect ratio converging to $\pi/2$. So, something more is needed. The advantage of our approach here and in [**S3**] is that it reduces all the topological concerns to finite dimensional geometric inequalities which work equally well for immersed examples.

These notes are organized as follows.

In §2 I prove the Optimality Theorem and the Rigidity Theorem modulo two main results, the Geometry Lemma and the Topology Lemma. The Geometry Lemma also works in the immersed case and the Topology Lemma requires an embedding. The rest of the paper is devoted to reducing these two results to tensegrity calculations.

In §3 I recall some geometric estimates from [**S3**] and also prove a new one along similar lines. The precise constants in these results are not important. They mainly serve to give us some *a priori bounds* which help us set up the limits and targets of our calculations.

In §4 I introduce the tensegrities of interest. I then explain the 3 main tensegrity calculations needed for the Geometry Lemma, and also 2 more tensegrity calculations of a similar nature which I use for the Topology Lemma. These calculations involve showing that certain piecewise algebraic functions are non-negative on the unit cube $[0,1]^{14}$. I call these calculations the Main Calculations.

In §5 I reduce the Topology Lemma to the last 2 calculations from §4 from some additional calculations, which I call Confiner Calculations. These calculations involve showing that certain piecewise algebraic functions are

non-negative on the unit cube $[0,1]^{10}$. The Topology Lemma (Lemma 2.3) a very natural statement which ought to have a traditional proof. A traditional proof of the Topology Lemma would remove the need for the Main Calculations 4,5 and all the Confiner Calculations. On the other hand, Main Calculations 1,2,3 seem essential to the whole approach.

In §6 I will describe my experiments so far with the calculations. In brief, I parametrize the relevant tensegrity families by unit cubes of dimension 10 and 14 and then use a hill climbing algorithm – a sort of random Newton's method – to optimize the relevant quantity. For the Main Calculations I have currently only worked in a simplified 10-dimensional slice of the 14-cube, the slice corresponding to planar configurations. As I will explain, my experience with the Confiner Calculations tells me that if the calculation works on the slice it also works in the whole space. Nonetheless, I will eventually code up the full 14-dimensional calculation. I hope to do all the calculations rigorously eventually.

I view these notes as a kind of time capsule. If the main conjectures are not proved in many years, it might be possible that some future computer will make short work of these tensegrity calculations, and then the conjectures will be proved. I think of it like this. Someone at the I.A.S. in 1950 might have a calculation that takes one second on a Macbook Pro, but while they were picking the dead crickets out of the vacuum tubes of the primitive computer there, they might think the calculation impossible. The calculation would not have been impossible but just premature.

I would like to thank Dan Cristofaro-Gardiner, Dmitry Fuchs, Steve Miller, and Sergei Tabachnikov for helpful discussions about this problem. I would especially like to thank Sergei for telling me about the problem and pointing me to his book with Dmitry. I would also like to acknowledge the support of the Simons Foundation, in the form of a 2020-21 Simons Sabbatical Fellowship, and also the support of the Institute for Advanced Study, in the form of a 2020-21 membership.

2 The Proof modulo Two Lemmas

2.1 Polygonal Moebius Bands

Basic Definition: Say that a *polygonal Moebius band* is a pair $\mathcal{M} = (\lambda, I)$ where $I : M_\lambda \to \boldsymbol{R}^3$ is an isometric embedding that is affine on each triangle of a triangulation of M_λ. We insist that the vertices of these triangles all lie on ∂M_λ, as in Figure 1.1. Any smooth isometric embedding $I' : M_\lambda \to \boldsymbol{R}^3$ can be approximated arbitrarily closely by this kind of map, so it suffices to work entirely with polygonal Moebius bands.

Associated Objects: Let $\delta_1, ..., \delta_n$ be the successive triangles of $\mathcal{M}$.

- The *ridge* of δ_i is edge of δ_i that is contained in ∂M_λ.
- The *apex* of δ_i to be the vertex of δ_i opposite the ridge.
- A *bend* is a line segment of δ_i connecting the apex to a ridge point.
- A *bend image* is the image of a bend under I.
- A *facet* is the image of some δ_i under I.

We always represent M_λ as a parallelogram with top and bottom sides identified. We do this by cutting M_λ open at a bend.

2.2 T Patterns

Say that a *T-pattern* is a pair of perpendicular coplanar disjoint bend images. Here is one of the central results from [**S3**].

Lemma 2.1 *A polygonal Moebius band of aspect ratio less than* $7\pi/12$ *has a T-pattern.*

Sketch of the Proof: A pair of perpendicular bend images is contained in a pair of parallel planes. The small aspect ratio allows us to rotate the image so that the bend images all make an angle of less than $\pi/4$ with the XY-plane. This property guarantees that the parallel planes just mentioned never contain vertical line segments. We consider the space $\mathcal{P}$ of perpendicular pairs of bend images. Generically this space is a topological 1-manifold.

We show that $\mathcal{P}$ contains a connected component $\mathcal{K}$ that is invariant under the involution which swaps the pair of bend images. Starting with a pair (α, β) of bend images in $\mathcal{K}$ we consider a path to (β, α). The corresponding pairs of perpendicular planes exchange their position and never contain vertical line segments. Hence, at some instant along the path, they coincide. ♠

The T-pattern in our polygonal Moebius band may not be unique, but we fix a T-pattern once and for all. Let β_1 and β_2 be two bends whose corresponding images $\beta_1^* = I(\beta_1)$ and $\beta_2^* = I(\beta_2)$ form a T-pattern. Since these segments do not intersect, we can label so that the line extending β_2^* does not intersect β_1^*. We cut M_λ open along β_1 and treat β_1 as the bottom edge. We now set $\beta_b = \beta_1$ and $\beta_t = \beta_2$ and (re)normalize as in §2.1. So, β_b^* connects $(-B, 0, 0)$ to $(0, 0, 0)$, and β_t^* is a translate of the segment connecting $(0, 0, 0)$ to $(0, T, 0)$. This translate still lies in the XY-plane. Here B and T are the lengths of these segments.

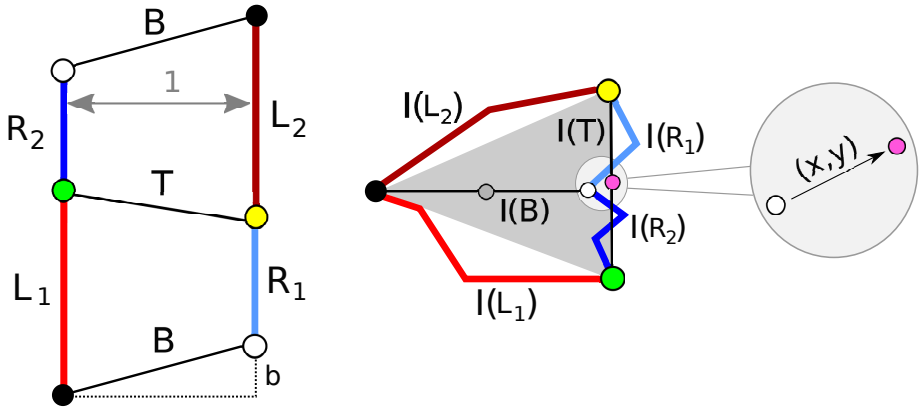

Figure 2.1: The standard normalization

The left side of Figure 2.1 shows M_λ. Reflecting in a vertical line, we normalize so that $L_1 \geq R_1$. This means that $L_2 \geq R_2$. We set

$$S_j = L_j + R_j. \tag{2}$$

We call this the *standard normalization*. Henceforth, we refer to (τ_1, τ_2) as a *standard pair*, and we let $\lambda(\tau_1, \tau_2)$ be the aspect ratio of the polygonal Moebius band from whence these trapezoids came.

The right side of Figure 2.1 shows the T pattern, and the corresponding images of the sets on the left under the isometry I. The wiggly curves we have drawn do not necessarily lie in the XY-plane but their endpoints do.

We normalize so that $(0,0,0)$ is the right endpoint of the B-bend image. In the figure (x,y) denotes the vector which points from the white to the pink vertex on the right side. We have blown this part of the figure up to make it more visible.

Remark: Actually, there are two such standard normalizations. We can make the replacements $M_\lambda \to \rho_1(M_\lambda)$ and $I \to \rho_2 \circ I \circ \rho_1$, where ρ_1 is reflection in the midpoint of M_λ and ρ_2 is reflection in the X-axis. This change preserves all our normalizations, and gives us the pair (τ_2, τ_1). Aside from swapping the names of the variables, the only thing that happens to Figure 2.1 is that the vector (x,y) changes to $(x,-y)$.

We call (τ_1, τ_2) *tame* if $\lambda(\tau_1,\tau_2) < \sqrt{3}+10^{-100}$. We add this tiny constant because, for the Rigidity Theorem, we will need to consider examples having aspect ratio slightly larger than $\sqrt{3}$.

2.3 Special Bends

We did a lot of numerical experiments and these led to the definitions we give here. Without these experiments, the definitions would seem very unmotivated.

Pitch: Given a polygonal Moebius band with a T-pattern, we use the standard normalization. For each bend β we let $\beta^* = \pi \circ I(\beta)$, the projection of the bend image $I(\beta)$ into the XY-plane. Each bend β of τ_1 has associated to it an angle $\theta \in [0,\pi]$ such that when we rotate the positive X-axis counterclockwise by θ we arrive at a ray parallel to β^*. We call θ the *bend pitch*.

Four Special Bends: Let $(k)_1$ stand for a bend β of τ_1 whose bend pitch is $+k\pi/12$. Let $(k)_2$ stand for a bend β of τ_2 whose bend pitch is $-k\pi/12$. We insist that $(0)_1 = (0)_2 = (0)$ is the bottom bend and $(6)_1 = (6)_2 = (6)$ is the top. We are interested in bends $(0), (1)_j, (4)_j, (6)$ for $j = 1,2$. We pick $(1)_j$ and $(4)_j$ to be the bends nearest (0) and (6), respectively, which satisfy the conditions. Let π be projection into the XY-plane. We illustrate our notation with an example:

$$(4)_1^* = \pi \circ I((4)_1) \tag{3}$$

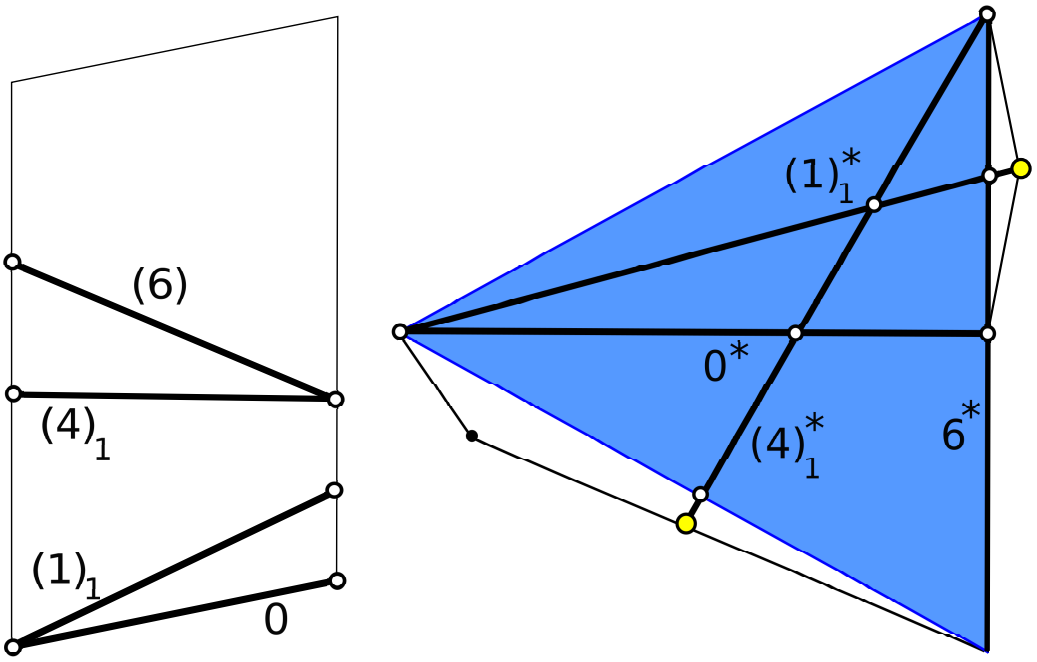

Figure 2.2: The 4 special bends and the yellow tips.

Tips and Bumps: We define the 1-*tip* of τ_j^* to be the right endpoint of $(1)_j^*$. We define the 4-tip of τ_j to be the left endpoint of $(4)_j^*$. These are the yellow vertices in Figure 2.2. Let Δ be the convex hull of the T-pattern. We say τ_j^* has a k-*bump* if the k-tip lies outside Δ. We make 2 more definitions:

- When $i \neq j$ we write $\tau_i \to_k \tau_j$ if some bend X of τ_i is such that X^* contains the k-tip of τ_j^*.
- We write $\tau_k \to_1 \tau_k$ if some bend $X \in [4,6]_k$ is such that X^* contains the 1-tip of τ_k.

2.4 The Two Main Lemmas

Suppose we have a polygonal Mobius band of aspect ratio less than

$$\sqrt{3} + 10^{-100}.$$

In the first result we assume that $\mathcal{M}$ has a T pattern but we do not assume that $\mathcal{M}$ is embedded. In the second result we require that $\mathcal{M}$ is embedded.

In either case, $\mathcal{M}$ has a T patterm. Let (τ_1, τ_2) be the associated tame pair. For the sake of working with closed subsets of objects, we allow the

limiting case when the horizontal and vertical segments $(0)^*$ and $(6)^*$ touch. That is, we allow $(0,0,0)$, the right endpoint of $(0)^*$, to lie in $(6)^*$.

Lemma 2.2 (Geometry) *Suppose that (τ_1, τ_2) is a tame pair with at least one of the following properties:*

- $\tau_1 \to_1 \tau_1$.
- $\tau_2 \to_1 \tau_1$.
- $\tau_2 \to_4 \tau_1$.

Then $\lambda(\tau_1, \tau_2) \geq 2\sqrt{3}$. Moreover, for any $\epsilon > 0$ there is a $\delta > 0$ such that if $\lambda(\tau_1, \tau_2) < 2\sqrt{3} + \delta$ then $(0)^* \cup (6)^*$ is within ϵ of an equilateral triangle of perimeter $2\sqrt{3}$ in the Hausdorff metric.

There is a symmetric result. Switching the roles of the indices, we get the same result when $\tau_2 \to_1 \tau_2$, etc. Thus, any of 6 hypotheses lead to the conclusion in the Geometry Lemma.

Our next result has the main topological component of the proof. In this result we insist that $\mathcal{M}$ is embedded.

Lemma 2.3 (Topology) *Suppose that (τ_1, τ_2) is a tame pair corresponding to an embedded polygonal Moebius band. Then at least one of 4 things is true.*

$$\tau_1 \to_1 \tau_1, \qquad \tau_2 \to_1 \tau_1, \qquad \tau_1 \to_4 \tau_2, \qquad \tau_2 \to_4 \tau_1.$$

2.5 Proofs of the Main Results

Given an immersed polygonal Moebius band with a T-pattern, we get a tame pair (τ_1, τ_2) provided that the aspect ratio satisfies $\lambda < \sqrt{3} + 10^{-100}$. Most of this chapter is devoted to proving the following result.

Proof of the Optimality Theorem: If we have an embedded polygonal Moebius band of aspect ratio less than $\sqrt{3}$ then we cut it open along the bends corresponding to a T-pattern. This gives rise to a tame pair (τ_1, τ_2) with $\lambda(\tau_1, \tau_2) < 2\sqrt{3}$. But then both the Topology Lemma and the Topology Lemma hold, and these results contradict each other. ♠

Proof of the Rigidity Theorem: Suppose we have a sequence of Paper Moebius bands whose aspect ratios converge to $\sqrt{3}$. The Topology Lemma applies for all examples sufficiently far along the sequence and thus we get a sequence $\{(\tau_{n,1}, \tau_{n,2}\})$ of tame pairs. By the Geometry Lemma, the convex hull Δ of the T pattern converges (modulo global isometries of $\boldsymbol{R}^3$) to the equilateral triangle Δ_0 of perimeter $2\sqrt{3}$. Since the length of the boundary $I(\partial M_\lambda)$ is at least as long as the perimeter of Δ, we see that this boundary must in fact converge in the Hausdorff Topology (modulo global isometries) to this same triangle Δ_0. This proves the Rigidity Theorem. ♠

3 A Priori Bounds

Theorem 3.1 and Lemma 3.2 are the main results in this chapter. The reader anxious to get to the main points should skip the proofs in this chapter and just use these results as black boxes. The precise statements of these results are not that important. We just need some *a priori* bounds to frame our calculations.

3.1 Geometric Bounds

Let Δ be the convex hull of the T-pattern $I(T) \cup I(B)$. Let b denote the slope of the B-bend and let t denote the slope of the T-bend. Figure 2.1 shows b but not t. Let Ω denote the set of pairs (b, t) which can arise in a standard pair with $\lambda < \sqrt{3}$. Figure 4.1 shows a plot of Ω in blue, as well as a quadrilateral $\widehat{\Omega}$ which we showed contains Ω.

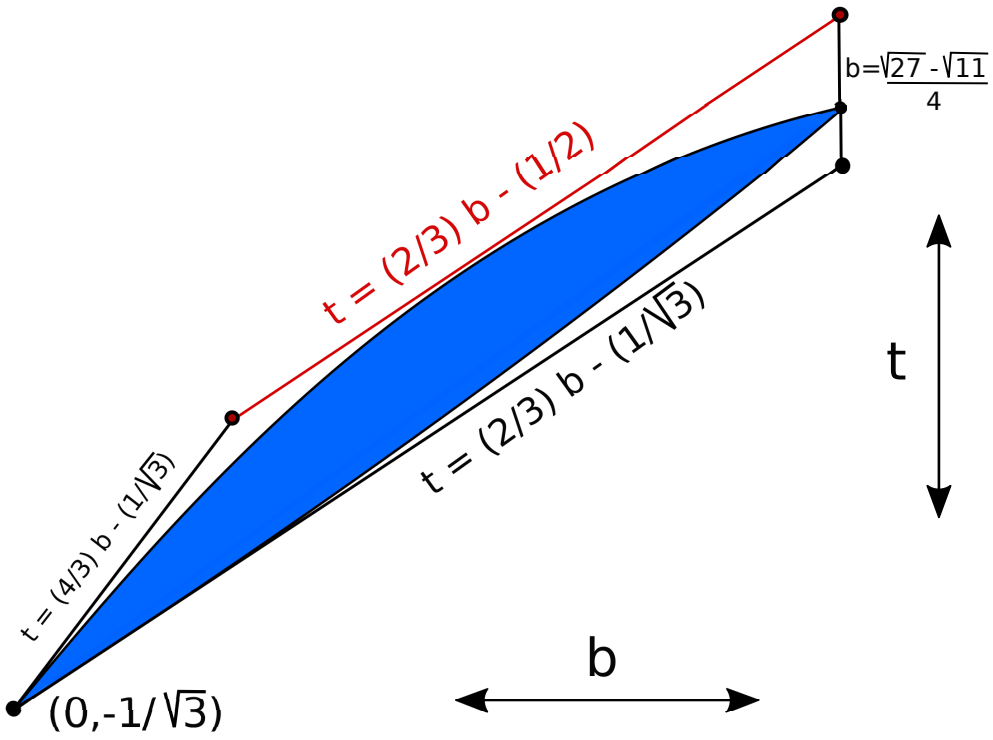

Figure 3.1: The range of slopes.

The following results are a subset of the main result in [**S3**, §4].

Theorem 3.1 *For any immersed paper Moebius with a T-pattern and aspect ratio less than $\sqrt{3}$, the following is true.*

1. $S_j \geq \sqrt{3} - (1/24)$ *for* $j = 1, 2$.
2. $x < 1/18$ *and* $|y| < 1/30$.
3. $\Omega \subset \widehat{\Omega}$.

Sketch of the Proof: The distances relations in Figure 2.1 give rise to 3 constraints:

1. $R_1 \geq \sqrt{x^2 + (T/2 - y)^2}$ and $R_2 \geq \sqrt{x^2 + (T/2 + y)^2}$.
2. $L_1 \geq \sqrt{(B + x)^2 + (T/2 + y)^2}$ and $L_2 \geq \sqrt{(B + x)^2 + (T/2 - y)^2}$.
3. $B^2 - L_j^2 + (T - R_j)^2 \leq 0$ for $j = 1, 2$.

Constraints 1 and 2 follow directly from the Pythagorean Theorem. Constraint 3 comes from a more subtle argument involving the Ridge Curve, which we define later in this chapter. We give the argument in [**S3**].

We also have the following relations:

$$L = \frac{S + b - t}{2}, \quad R = \frac{S - b + t}{2}, \quad B = \sqrt{1 + b^2}, \quad T = \sqrt{1 + t^2}. \tag{4}$$

Plugging these relations into our various constaints and doing some calculus, we arrive at the statements in Theorem 3.1. ♠

3.2 Angle Range Bounds

Now we prove a new result, similar in spirit to Theorem 3.1. This new result also helps frame our calculations by giving *a priori* bounds.

Let the interval $[0, 1]_1$ denote all the bends in the trapezoid τ_1 that interpolate between 0 and 1_1. Let $\Theta([0, 1]_1)$ denote the range of bend pitches for bend images in $[0, 1]_1$. (We make similar definitions for the other intervals.) The initial bend image $(0)^*$ has bend pitch 0. The final bend image $(1_1)^*$ has bend pitch $\pi/12$. It is tempting to guess that $\Theta([0, 1]_1) = [0, \pi/12]$ but, since the bend pitches need not vary monotonically, we cannot conclude this. In this section we establish the following bounds.

Lemma 3.2 *Let $\theta_0 = \pi/30$. For any tame pair (τ_1, τ_2) of trapezoids,*

1. $\Theta([0,1]_1) \subset [-\theta_0, \pi/12]$.
2. $\Theta([0,4]_1) \subset [-\theta_0, 4\pi/12 + \theta_0]$.
3. $\Theta([4,6]_1) \subset [4\pi/12, \pi/2 + \theta_0]$.

Remark: One consequence Lemma 3.2 is that any bend in τ_1 has pitch in $[-\theta_0, \pi/2 + \theta_0]$. This is a more precise version of the statement above that the pitches essentially lie in $[0, \pi/2]$ except for a bit of slope at either end.

The rest of the chapter is devoted to proving this result. The technique used here is not used anywhere else in this paper, though we used it extensively in [**S3**]. At this point, the reader can use Theorem 3.1 and Lemma 3.2 as black boxes. Again, they only serve to place *a priori* bounds on our tensegrity calculations.

3.3 The Ridge Curve

Here we recall some more notions from [**S3**].

The Sign Sequence: Let $\delta_1, ..., \delta_n$ be the triangles of the triangulation associated to $\mathcal{M}$, going from bottom to top in P_λ. We define $\mu_i = -1$ if δ_i has its ridge on the left edge of P_λ and $+1$ if the ridge is on the right. The sequence for the example in Figure 1.1 is $+1, -1, +1, -1$.

The Core Curve: There is a circle γ in M_λ which stays parallel to the boundary and exactly $1/2$ units away. In Equation 1, this circle is the image of $\{1/2\} \times [0, \lambda]$ under the quotient map. We call $I(\gamma)$ the *core curve*.

The left side of Figure 3.2 shows M_λ and the pattern of bends. The vertical white segment is the bottom half of γ. The right side of Figure 3.2 (which has been magnified to show it better) shows $I(\tau)$ where τ is the colored half of M_λ. All bend angles are π and the whole picture is planar. The colored curve on the right is the corresponding half of the core curve. Incidentally, for τ we have $L + R = 1.72121... < \sqrt{3}$.

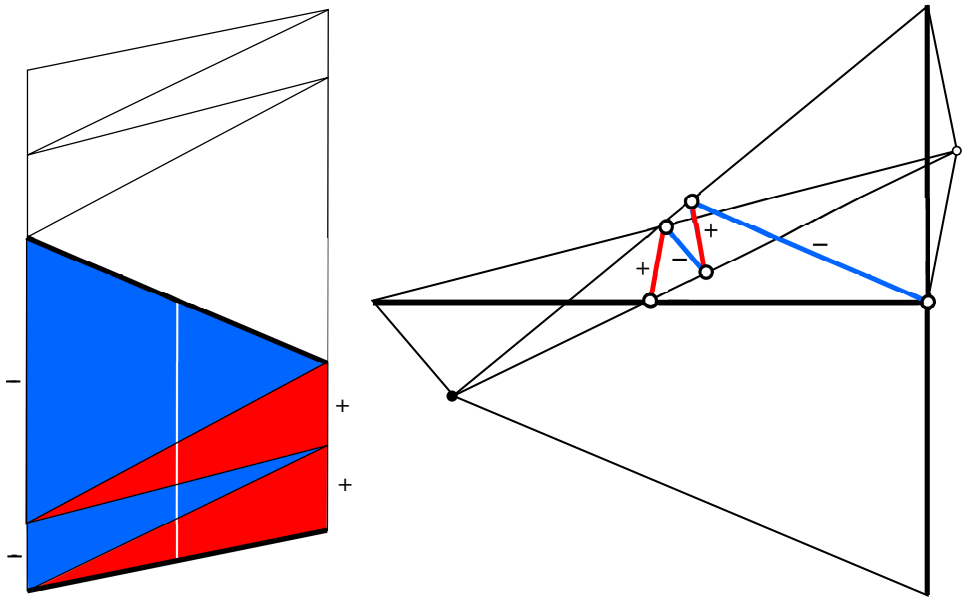

Figure 3.2: The bend pattern and the bottom half of the image

The Ridge Curve: We show the picture first, then explain.

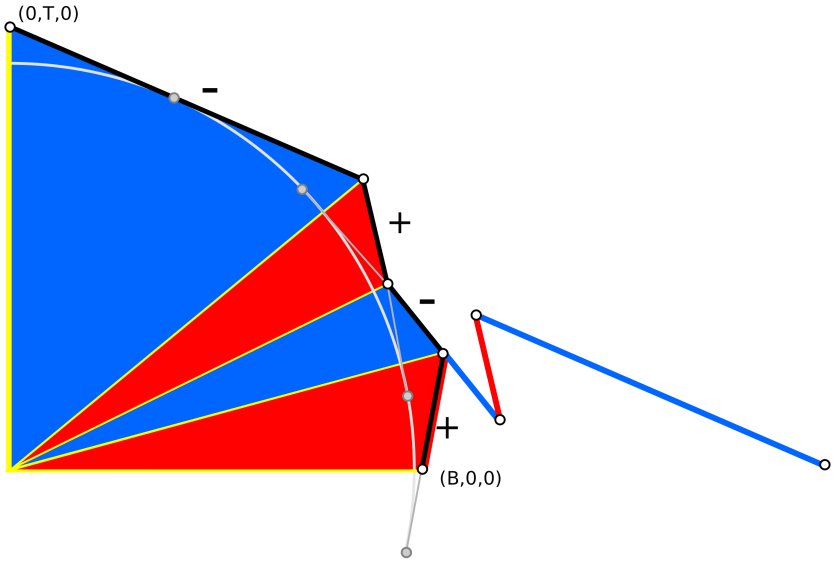

Figure 3.3: Half of the core curve (red/blue) and half of the ridge curve (black). The core curve is scaled up by a factor of 2.

Let β_b be the bottom edge of the parallelogram representing M_λ. We normalize so that I maps the right vertex of β_b to $(0,0,0)$ and the left vertex to $(-B,0,0)$, where B is the length of β_b. Let $E_1,...,E_n$ be the successive edges of the core curve, treated as vectors. Let

$$\Gamma_i' = 2\mu_i E_i, \qquad i = 1,...,n. \tag{5}$$

Let Γ be the curve whose initial vertex is $(B,0,0)$ and whose edges are $\Gamma_1',...,\Gamma_n'$. Here $\mu_1,...,\mu_n$ is the sign sequence.

Γ has length 2λ, connects $(B,0,0)$ to $(-B,0,0)$, and is disjoint from the open unit ball. The lines extending the sides of Γ are tangent to the unit sphere. We rotate so that Γ contains $(0,T,0)$ for some $T>1$. If we cone Γ to the origin, we get a collection $\Delta_1,...,\Delta_n$ of triangles, and Δ_i is the translate of $\mu_i I(\delta_i)$ whose apex is at the origin. In particular, the vectors pointing to the vertices of Γ are parallel to the corresponding bend images. Figure 3.3 shows the portion of the ridge curve (in black) associated to the example in Figure 2.1. We have also scaled the core curve by 2 and translated it to show the relationships between the two curves.

3.4 Proof of Lemma 3.2

Let us dispense with a representative case first. The reason why $\Theta([0,1]_1)$ ends at $\pi/12$ rather than $\pi/12+\theta$ for some positive θ is that we take $(1)_1$ be the first bend after $(0)_1$ with pitch $\pi/12$. Similar remarks apply to the other cases where we have no slop over the endpoint in the bound.

Now we consider the other cases. Let Γ_1 be the portion of the ridge curve associated to τ_1. The curve Γ_1 connects $(B,0,0)$ to $(0,T,0)$. This curve perhaps does not stay entirely in the positive sector consisting of points (x,y,z) with $x,y \geq 0$, but it certainly does not "go around the back". For instance, the projection $\pi(\Gamma_1)$ avoids the ray of slope 1 which starts at $(0,0)$ and goes through $(-1,-1)$. Here π is projection into the XY-plane.

Let Π_1 and Π_2 be two planes in $\boldsymbol{R}^2$ which contain the Z-axis. Suppose that the dihedral angle between Π_1 and Π_2 is θ. Suppose these planes are ordered so that Γ_1 must hit Π_1 before hitting Π_2. We give a lower bound on the length if Γ_1 goes from $(B,0,0)$ to a point $p \in \Pi_2$, to a point $q \in \Pi_1$ to $(0,T,0)$. Figure 3.4 shows what $\Pi(\Gamma_1)$ would look like in several cases.

It is worth pointing out that Γ_1 lies outside the open unit ball, but $\pi(\Gamma_1)$ does not necessarily lie outside the open unit disk. The shaded region in Figure 3.4 is one quarter of the unit disk.

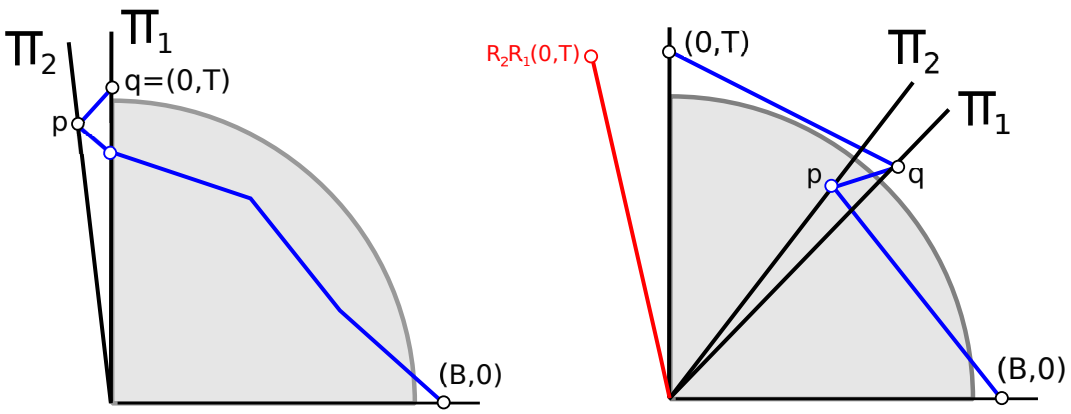

Figure 3.4: Projection of the ridge curve.

Let ℓ_1 denote the length of Γ_1. Let R_j denote reflection in Π_j. We write $\Gamma_1 = \Gamma_{11} \cup \Gamma_{12} \cup \Gamma_{13}$, where Γ_{11} goes from $(B, 0, 0)$ to p and Γ_{12} goes from p to q and Γ_{13} goes from q to $(0, T, 0)$.

The continuous path

$$\Gamma_1^* = \Gamma_{11} \cup R_2(\Gamma_{12}) \cup R_2R_1(\Gamma_{13})$$

has the same length as Γ_1 and connects $(B, 0)$ to $R_2R_1(0, T, 0)$. This latter point lies in the XY plane and makes an angle 2θ with the Y-axis, as shown (projected into the XY-plane) on the right side of Figure 7.2. But then

$$\ell_1 = |\Gamma_1^*| \geq (\pi/2) + 2\theta. \tag{6}$$

Let ℓ_2 denote the length of Γ_2, the portion of the ridge curve associated to τ_2. From Theorem 3.1 we have $\ell_1 < \sqrt{3} + (1/24)$. Combining this with Equation 6, we have

$$2\theta < \sqrt{3} + (1/24) - (\pi/2) < \pi/15. \tag{7}$$

Hence $\theta < \pi/30$. If one of our estimates failed, Γ_1 would make exactly the kind of path just studied, where the angle θ between Π_1 and Π_2 would be $\theta = \pi/30$. This is a contradiction. ♠

4 Proof of the Geometry Lemma

4.1 Trapezoids and Capacities

We first mention a general property of our notation. For any object X in the plane X' will be the image of X under some kind of piecewise linear map to $\boldsymbol{R}^3$. Suppose that

$$Q = (Q_0, Q_1, Q_2, Q_3), \qquad Q' = (Q'_0, Q'_1, Q'_2, Q'_3)$$

are as follows: Q is a trapezoid in the plane having vertical sides and width 1 and Q' is a quadrilateral in $\boldsymbol{R}^3$. Figure 4.1 shows Q. The numbers indicate the labeling of the vertices of Q.

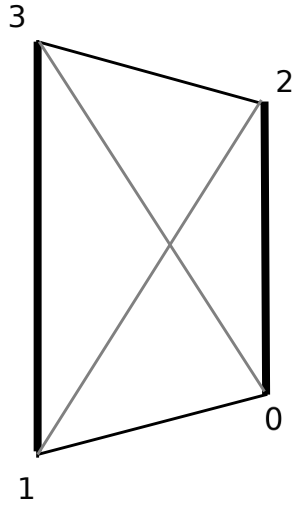

Figure 4.1: The quadrilateral tensegrity.

Let $d_{ij} = \|Q_i - Q_j\|$. Let $d'_{ij} = \|Q'_i - Q'_j\|$. We define the *height* of Q to be the sum of the vertical edge lengths of Q, namely $d_{02} + d_{13}$. We call the pairs (Q_0, Q_1) and (Q_2, Q_3) the *bends* and we call the pairs (Q'_0, Q'_1) and (Q'_2, Q'_3) the *bend images*. This lines up our discussion here with Figure 2.3.

We say that (Q, Q') is a *plain tensegrity pair* if $d_{ij} \geq d'_{ij}$ for all the edges of Q, with equality for the pairs $(0, 1)$ and $(2, 3)$. So, we have 2 equalities and 6 inequalities. We define the *capacity* of Q' to be the minimum height of Q where (Q, Q') is a tensegrity pair. We denote this by $\kappa(Q')$.

4.2 Compound Tensegrities

This time let $\Psi' = (Q'_0, ..., Q'_{k-1})$ be a collection of k quadrilaterals in space. We call k the *complexity*. We have

$$Q'_j = (Q'_{j,0}, Q'_{j,1}, Q'_{j,2}, Q'_{j,3}), \qquad j = 0, ..., (k-1).$$

We insist that these quads abut, in the sense that

$$Q_{j,2} = Q_{j+1,0}, \qquad Q_{j,3} = Q_{j+1,1}, \qquad j = 1, ..., (k-2).$$

When these conditions are satisfied, we call Ψ' a *compound tensegrity*. We will always take $k = 2$ here.

We call Ψ' *cyclic* if, additionally,

$$Q_{k-1,0} = Q_{0,1}, \qquad Q_{k-1,1} = Q_{0,0}. \tag{8}$$

The mismatch of the indices is deliberate, and is designed to reflect the structure of a Moebius band. When Ψ' we will take $k = 4$.

In either case, we define

$$\kappa(\Psi') = \sum_{i=0}^{k-1} \kappa(Q'_i). \tag{9}$$

Suppose now that (τ_1, τ_2) is a tame pair (or indeed any pair) of trapezoids coming from the process of cutting a polygonal Moebius band open along the bends corresponding to a T-pattern. Suppose we choose a finite number of additional bends $\beta_{1,1}, ..., \beta_{1,k_1-1}$ in the interior of τ_1 and $\beta_{2,1}, ..., \beta_{1,k_2-1}$ in the interior of τ_2. Then we can get a compound tensegrities in 3 ways:

- We set $\beta_{1,0} = b$ and $\beta_{1,k_1} = t$ and then let Q_j be the quadrilateral whose vertices are the endpoints of $I(\beta_1, j)$ and $I(\beta_1, j+1)$ with the vertices labeled as in Figure 4.1. Then $Q_0, ..., Q_{k_1-1}$ is a compound tensegity. Here b and t are the bottom and top bend of τ_1. Call this Ψ'_1.

- We set $\beta_{2,0} = t$ and $\beta_{2,k_2} = b'$ and then repeat the same construction. Here b' is the top of τ_2 which is parallel to b. Call this Ψ'_2.

- We can take the concatenation $\Psi' = \Psi'_1, \Psi'_2$ of the previous two tensegrities. This gives us a cyclic compound tensegrity.

Here is the key inequality:

$$S_1 \geq \kappa(\Psi'_1), \qquad S_2 \geq \kappa(\Psi'_2), \qquad \lambda(\tau_1, \tau_2) \geq \kappa(\Psi'). \tag{10}$$

4.3 Main Calculations

Now we describe 2 families of cyclic tensegrities of complexity 4. The input for both families is a pair Θ_1, Θ_2 of angle intervals. We have 2 middle bend images

$$\beta'_j = \overline{Q'_{j,0}Q'_{j,1}}, \qquad j = 1, 2.$$

We call the bend images β'_j for $j = 0, 1, 2, 3$. Let π be projection into the XY-plane. Recall that the *pitch* of β' is the angle that $\pi(\beta')$ makes with the X-axis.

The family $\mathcal{F}(h, \Theta_1, \Theta_2)$ is the set of tensegrities with these properties.

- $\beta'_0 \cup \beta'_3$ is a T-pattern, with $\pi(\beta'_0)$ horizontal and $\pi(\beta'_3)$ vertical.
- β'_j has pitch in the interval Θ_j for $j = 1, 2$.
- $\pi(\beta'_2)$ contains the point $\pi(Q_{1,j})$.

The family $\mathcal{G}(j, \Theta_1, \Theta_3)$ is the set of tensegrities with these properties.

- $\beta'_0 \cup \beta'_2$ is a T-pattern, with $\pi(\beta'_0)$ horizontal and $\pi(\beta'_2)$ vertical.
- β'_j has pitch in the interval Θ_j for $j = 1, 3$.
- $\pi(\beta'_3)$ contains the point $\pi(Q_{1,j})$.

Referring to polygonal Moebius bands, the family $\mathcal{F}$ encodes the situation where we focus on two bends contained in τ_1 and the family $\mathcal{G}$ encodes the situation where we focus on two bends, one contained in τ_1 and the other contained in τ_2.

We define the *capacity* of one of these families to be the infimal capacity of members of the family. Let Δ_0 denote the equilateral triangle of perimeter $2\sqrt{3}$ coming from the equilateral triangle shown on the right side of Figure 1.1. Here are our main calculations.

1. $\mathcal{F}(0, [\pi/12, \pi/12], [4\pi/12, 6\pi/12+\theta_0]$ has capacity $2\sqrt{3}$. The only member of this family achieving $2\sqrt{3}$ has the property that θ_0 is the convex hull of $\beta'_0 \cup \beta'_3$.
2. $\mathcal{G}(0, [\pi/12, \pi/12], [6\pi/12-\theta_0, 8\pi/12]$ has capacity $2\sqrt{3}$. The only member of this family achieving $2\sqrt{3}$ has the property that θ_0 is the convex hull of $\beta'_0 \cup \beta'_2$.

3. $\mathcal{G}(1, [4\pi/12, 4\pi/12], [-\pi/2, 3\pi/2]$ has capacity $2\sqrt{3}$. The only member of this family achieving $2\sqrt{3}$ has the property that θ_0 is the convex hull of $\beta_0' \cup \beta_2'$.

4. $\mathcal{F}(0, [4\pi/12, 4\pi/12], [-\theta_0, \pi/12])$ has capacity greater than $2\sqrt{3}$.

5. $\mathcal{G}(0, [4\pi/12, 4\pi/12], [-\theta_0, \pi/12])$ has capacity greater than $2\sqrt{3}$.

Here $\theta_0 = \pi/30$.

Calculations 1,2,3 immediately imply the truth of the Geometry Lemma. Calculations 4,5 will be used in the proof of the Topology Lemma. Figures 4.2-4.6 show one member from each of the 5 families. The figure on the left emphasizes the bends and the figure on the right emphasizes the complementary edges. These complementary edges correspond to the boundary of our polygonal Moebius band.

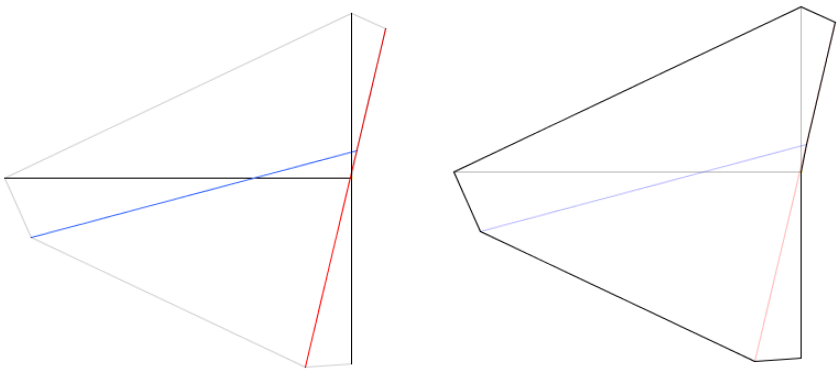

Figure 4.2: The constraints for Main Calculation 1

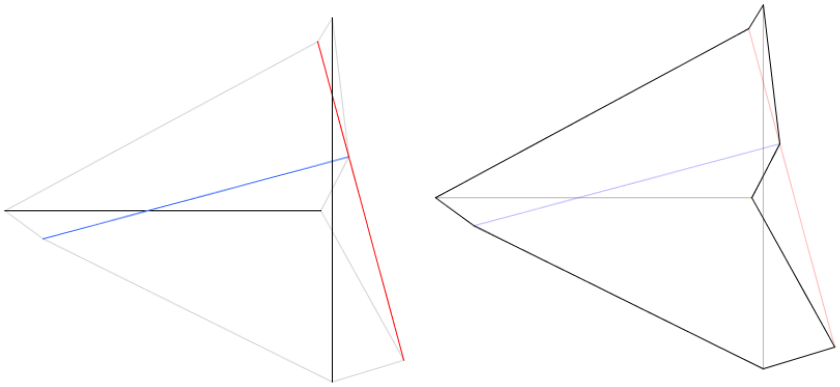

Figure 4.3: The Constraints for Main Calculation 2

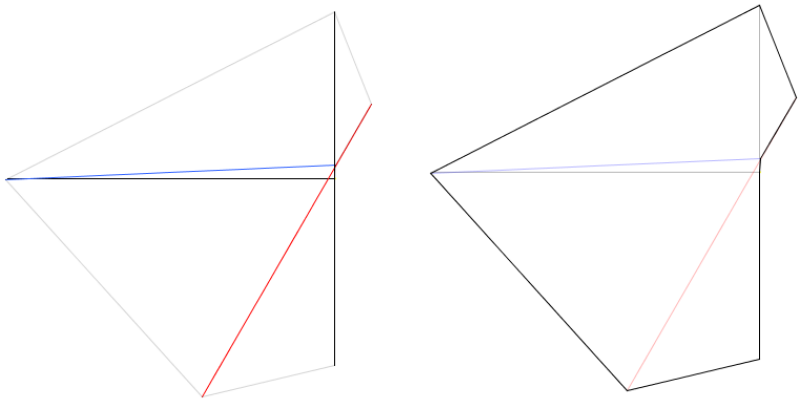

Figure 4.4: The Constraints for Main Calculation 3

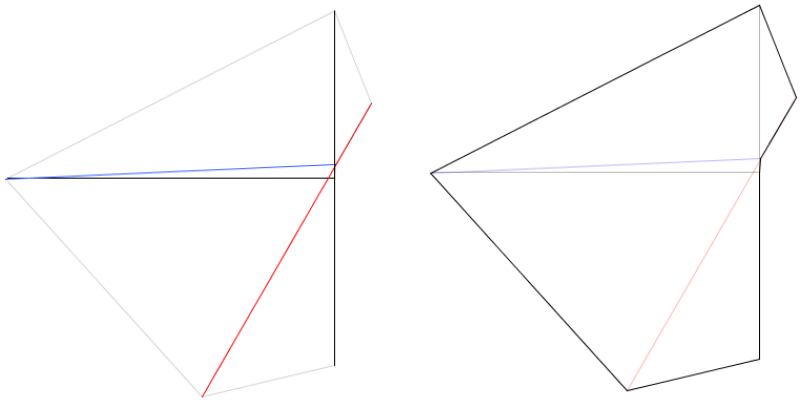

Figure 4.5: The Constraints for Main Calculation 4

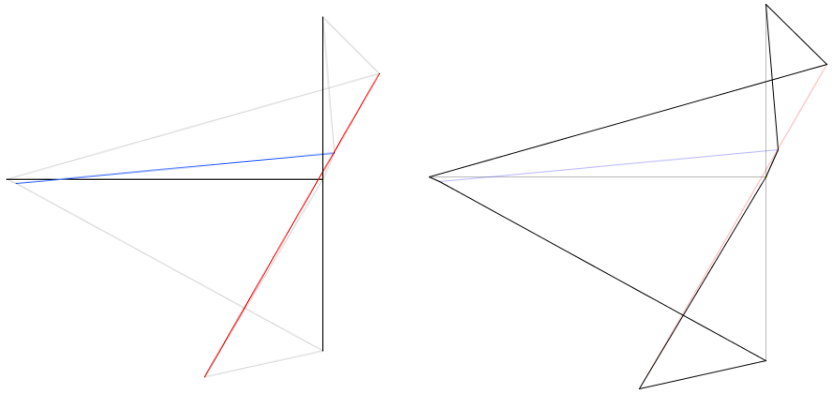

Figure 4.6: The Constraints for Main Calculation 5

5 Proof of the Topology Lemma

The rest of the chapter is devoted to proving the Topology Lemma. We first prove the Topology Lemma under a condition we call *topological goodness*. Following this, we reduce topological goodness to various geometric properties that can be established by tensegrity calculations.

5.1 Topological Goodness

Given 2 bends β_1 and β_2 we write $\beta_1|\beta_2$ if β_1^* and β_2^* intersect in a point that is interior to both. Assuming that $\beta_1|\beta_2$, we write $\beta_1 \uparrow \beta_2$ if the vertical line through $\beta_1^* \cap \beta_2^*$ intersects $I(\beta_1)$ above where it intersects $I(\beta_2)$. Otherwise we write $\beta_1 \downarrow \beta_2$. We must have one or the other when we have an embedded Mobius band. We will encode the information below in a picture reminiscent of a knot diagram.

We call $(\beta_1, \beta_2, \beta_3)$ a *topologically bad triple* triple if the crossings are inconsistent in one of two ways:

- $\beta_1 \uparrow \beta_2$ and $\beta_1 \downarrow \beta_3$.
- $\beta_1 \downarrow \beta_2$ and $\beta_1 \uparrow \beta_3$.

Otherwise we call the triple *topologically good*.

Good Pairs: We say that the pair (τ_1, τ_2) is *topologically good* if:

1. $(0)|(4)_j$ and $(1)_1|(4)_j$ and $(4)_1|(4)_2$.
2. $(4_j, 1_1, 0)$ is topologically good.
3. $(1_1, 4_j, 6)$ is topologically good unless $\tau_j \to_1 \tau_i$.
4. $(4_{3-j}, 4_j, 0)$ is topologically good unless $\tau_j \to_4 \tau_{3-j}$.

Here $j = 1, 2$ in all statements.

Lemma 5.1 *The Topology Lemma holds for topologically good pairs.*

Proof: Suppose first that $(1)_1^*$ and 6^* do not cross. Since $(1)_1|(4)$, there is by continuity some $\beta \in [4, 6]_1$ such that β^* contains the 1-tip of $(1)_1^*$. This means that $\tau_1 \to \tau_1$.

Henceforth we assume $(1)_1|(6)$. We normalize by an ambient isometry so that $(1)_1 \downarrow (6)$. Figure 5.1 shows the 4 possibilities.

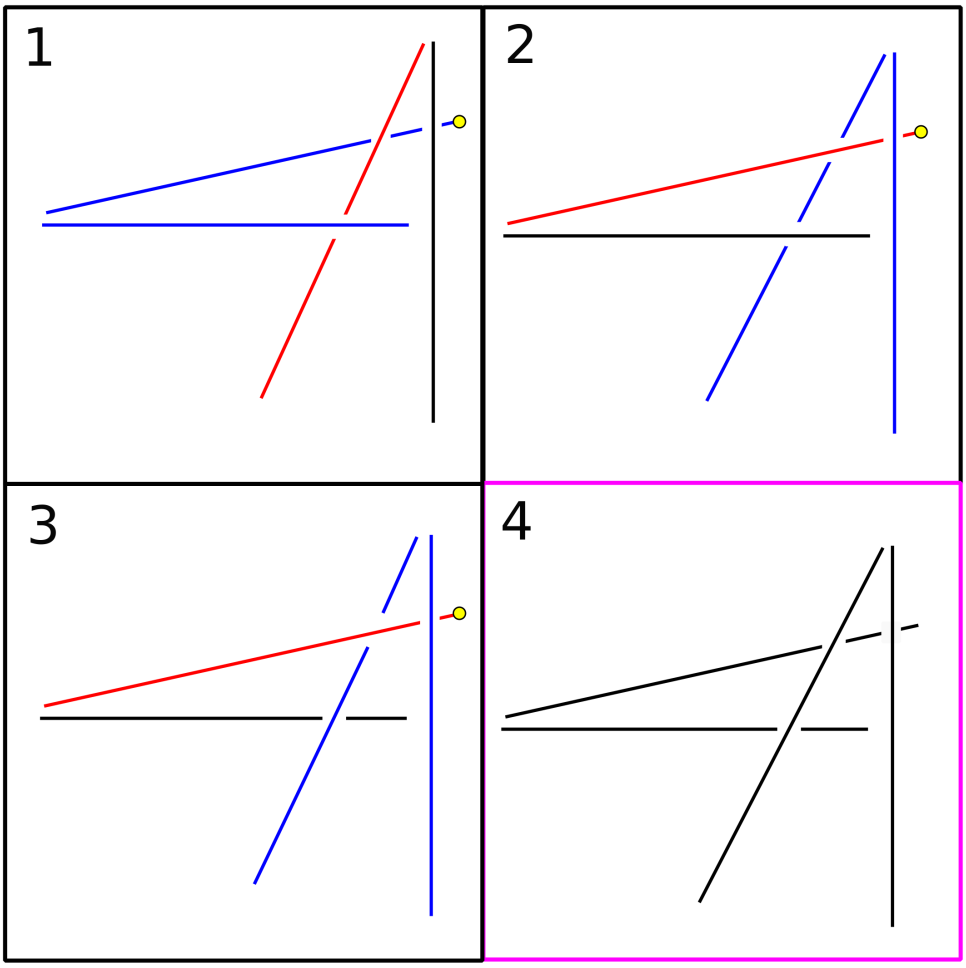

Figure 5.1: Four possibilities

Case 1 has the topologically bad triple $(4_1, 1_1, 0)$, so this case cannot happen. Cases 2 and 3 have the topologically bad triple $(1_1, 4_1, 6)$. Hence $\tau_1 \to_1 \tau_1$ in these cases. This leaves Case 4.

Let us explore Case 4 in more detail. Figure 5.2 shows the situation.

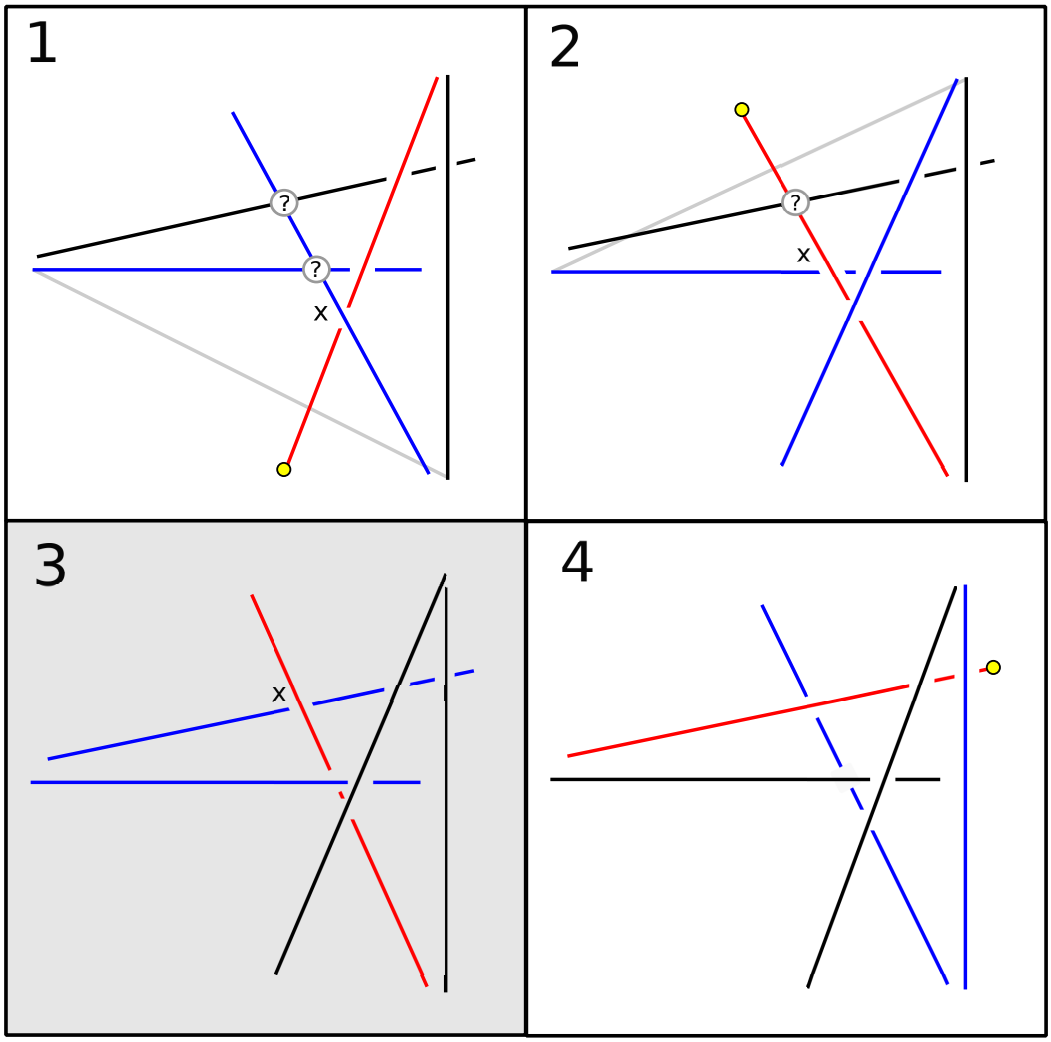

Figure 5.2: Four possibilities

In Case 1 we have left 2 crossings undetermined and we place an x by the crossing we study. If the crossing is as appears, then $(4_1, 4_2, 0)$ is a topologically bad triple. But then $\tau_2 \to_4 \tau_1$.

In Case 2 we are working on the next crossing, and we still leave one undetermined. This has the same structure as Case 1 but with the indices 1 and 2 reversed. Hence $\tau_1 \to_4 \tau_2$ here.

In Case 3, the situation is impossible because $(4_2, 1_1, 0)$ is a topologically bad triple.

In Case 4, the triple $(1_1, 4_2, 6)$ is topologically bad. Hence $\tau_2 \to_1 \tau_1$. ♠

5.2 Geometrical Goodness

Here we give conditions under which the pair (τ_1, τ_2) is topologically good. We call the triple $(\beta_1, \beta_2, \beta_3)$ of bends *geometrically good* if $\beta_1 | \beta$ for all bends β that interpolate between β_2 and β_3. Since the space of bends is topologically a circle, we we have to be careful about what we mean here. We never allow our bends in an interpolating family to cross the bend labeled T on the left side of Figure 2.1. In the next result, we have our standard normalization. We call (τ_1, τ_2) *geometrically good* if the various conditions on the triples in the definition of topological goodness hold with the word *geometrically* replacing the word *topologically* in every instance.

Lemma 5.2 *(τ_1, τ_2) is topologically good if (τ_1, τ_2) is geometrically good.*

Proof: For a given triple, the implication "geometric implies topological" works the same way in all cases. We just consider two representative cases.

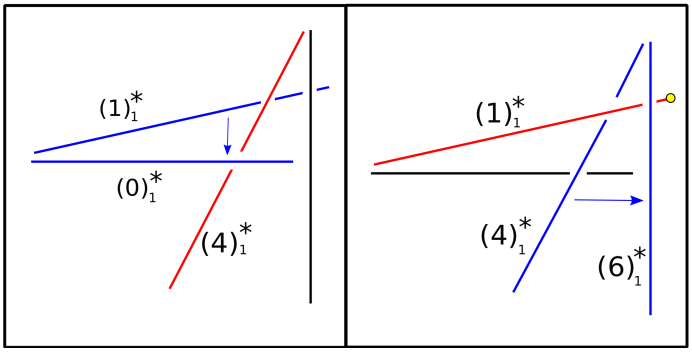

Figure 5.3: A topological contradiction.

First look at the left half of Figure 5.3. Consider the triple $(4_1, 1_1, 0)$. In $\boldsymbol{R}^3$, as we sweep the blue $I(I_1)$ over to the blue $I(0)$ we end up on the wrong side of the red $I(4_1)$. The geometrical goodness prevents the crossing type from changing.

Now look at the right half of Figure 5.3. As we sweep the blue $I(4_1)$ over to the blue $I(6)$, we get the same contradiction as in the previous case, except that now the crossing can change if the projection of one of the blue segments in our sweepout contains the right endpoint of $(1)_1^*$. In other words, $(1_1, 4_1, 6)$ is topologically good unless $\tau_1 \to_1 \tau_1$. ♠

5.3 Confiner Pairs

We now set up the tensegrity calculations needed to establish Geometrical Goodness for (τ_1, τ_2).

A *prism* is a set of the form $P \times \boldsymbol{R}$, where $P \subset \boldsymbol{R}^2$ is an open convex polygon. The prism is the intersection of finitely many halfspaces whose boundaries project to lines in the XY-plane. We call these halfspaces the *defining halfspaces*. In general, we say that a halfspace whose boundary projects to a line in the XY-plane is a *prism halfspace*.

Our calculations all involve compound tensegrities with 2 quadrilaterals. Given such a tensegrity Υ', we let $\theta(\Upsilon')$ denote the pitch of the middle bend image $\overline{Q'_{10}Q'_{11}}$. Here are the constraints.

1. The horizontal and vertical bends make a T-pattern.
2. $\theta(\Upsilon') \in \Theta \subset [-\theta_0, \pi/2 + \theta_0]$, for some interval Θ.
3. A selected endpoint of the middle bend lies in the prism halfspace H.

Again, $\theta_0 = \pi/30$ as in Lemma 3.2.

For various choices of (Θ, H) the calculations, if valid, show

$$\kappa(\Upsilon') > \sqrt{3} + \frac{1}{24} + 10^{-100}.$$

We call (j, Θ, H) an *excluder*. Suppose that (j, Θ, H) is an excluder. Suppose (τ_1, τ_2) is a tame pair and $I(\beta)$ is a bend image associated to τ_1 having pitch in Θ. Then the relevant endpoint of $I(\beta)$ cannot lie in H. Here is the argument. The calculation above, applied to $\Upsilon' = \Psi'_1$, shows that if this fails then $S_1 > \sqrt{3} + \frac{1}{24} + 10^{-100}$. On the other hand, Theorem 3.1 gives us that $S_2 > \sqrt{3} - \frac{1}{24}$. But then $\lambda(\tau_1, \tau_2) > \sqrt{3} + 10^{-100}$, contradicting tameness.

Confining Pairs: We say that a *prism pair* is a pair (P_1, P_2) of disjoint prisms. We say that this pair *confines* a line segment if one endpoint of the line segment lies in P_1 and the other lies in P_2. Let $H_{j,1}, ..., H_{j,k_j}$ be the closed complements of the defining halfspaces for P_j. For fixed Θ we set up calculations which would show that $(j, \Theta, H_{j,i})$ is an excluder for all relevant indices. This shows that (P_1, P_2) confines the middle bend image $I(\beta)$ associated to (τ_1, τ_2) provided that and pitch θ of the middle bend lines in Θ. We abbreviate this by saying that (P_1, P_2) is a Θ *confiner*.

5.4 The Confiner Calculations

Now we describe the Confiner Calculations we need in order to establish Geometrical Goodness.

Figure 5.4 shows 6 regions in the plane. These regions are projections of prism pairs. The red and green regions are interchanged by reflection in the X-axis. As always, we show the projection to the XY-plane. The circle in the picture denotes a place where the pieces are actually disjoint but look tangent. The green piece is entirely below the line extending the bottom edges of the blue piece.

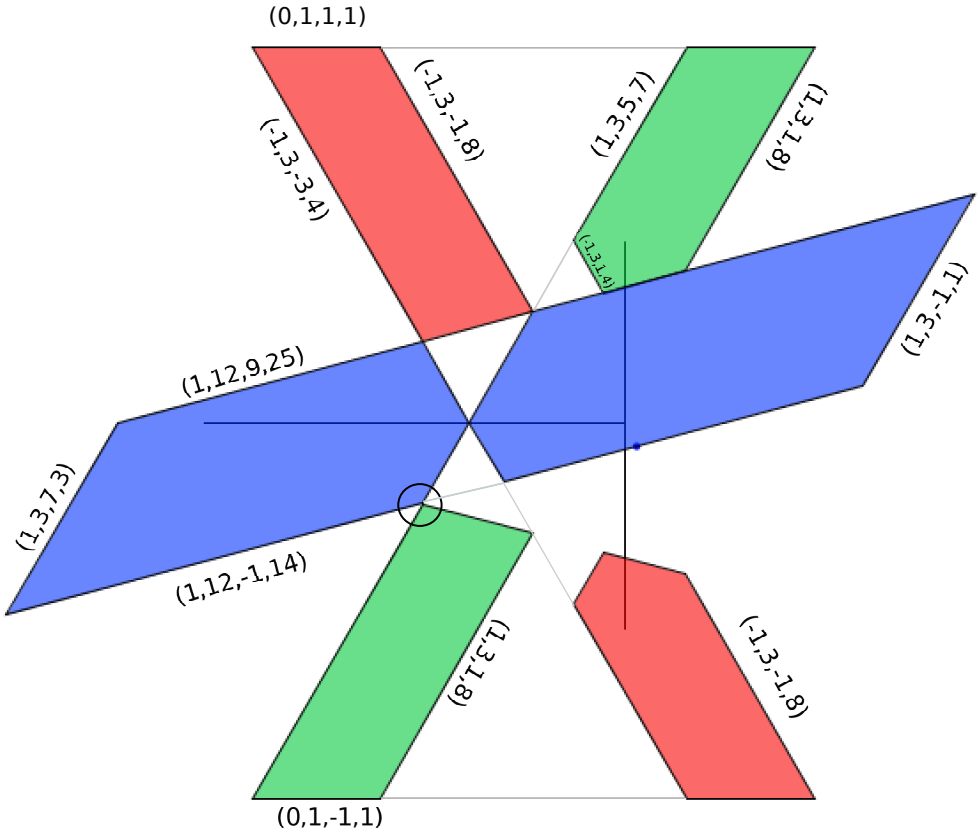

Figure 5.4: Confining prisms

Figure 5.5 shows some new regions. The blue regions in Figure 5.5 are the same as in Figure 5.4. The magenta regions are reflections of the blue regions in the X-axis.

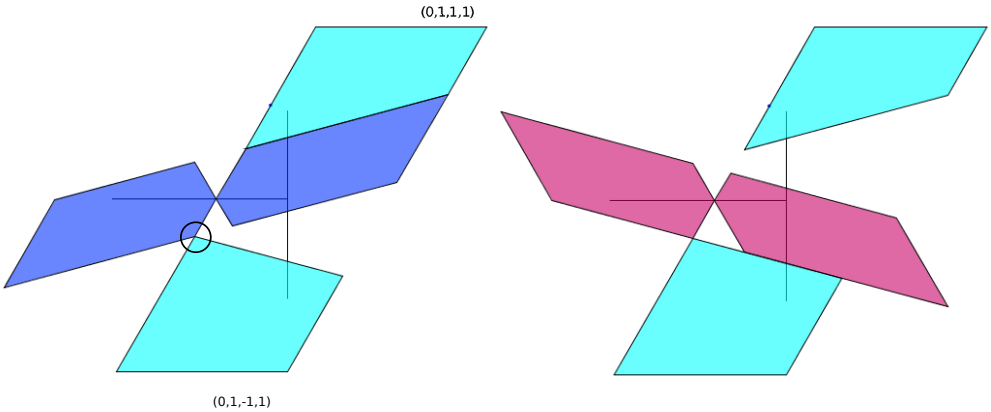

Figure 5.5: Confining prisms

The red regions in Figure 5.6 are the same as in Figure 5.4.

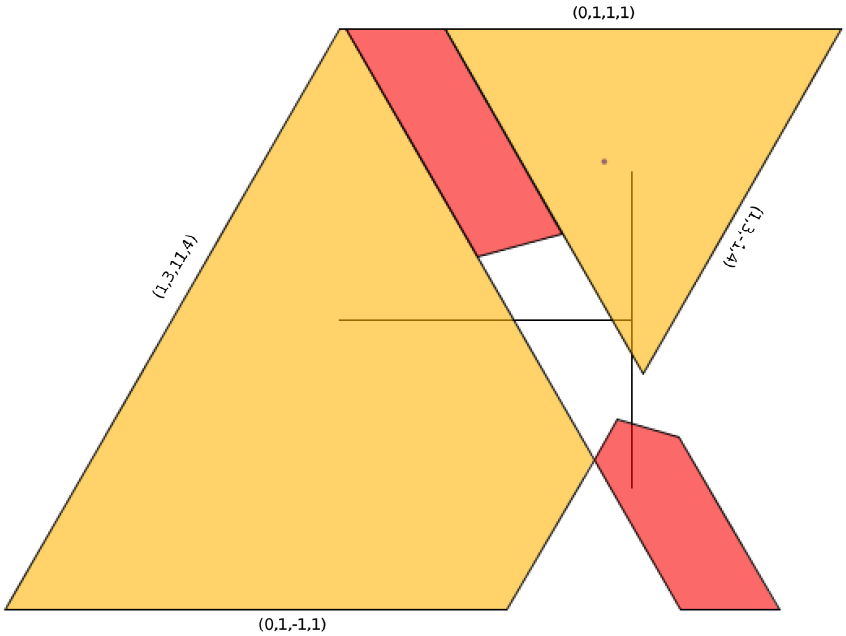

Figure 5.6: Confining prisms

Motivated by Lemma 3.2, we define

$$\Theta_{01} = [-\theta_0, \pi/12], \quad \Theta_{44} = [4\pi/12, 4\pi/12],$$
$$\Theta_{46} = [4\pi/12, 6\pi/12 + \theta_0], \quad \Theta_{04} = [-\theta_0, 4\pi/12 + \theta_0]. \tag{11}$$

Here are the calculations.

1. The green prisms are a Θ_{44} confining pair.

2. The blue prisms are a Θ_{01} confining pair.

3. The cyan prisms are a Θ_{46} confining pair.

4. The orange prisms are a Θ_{04} confining pair.

5.5 Establishing Geometrical Goodness

Now we deduce Geometrical Goodness from the confiner tensegrity calculations we discussed in the previous chapter.

Lemma 5.3 *The triples* $(4_i, 1_j, 0)$ *are geometrically good for all* $, i, j \in \{1, 2\}$.

Proof: If symmetry it suffices to take $j = 1$. Our argument refers to Figure 5.7. Call a segment *red* if it has one endpoint in the interior of one red region and the other endpoint in the interior of the other. Likewise define *green segments* and *blue segments*. Looking at Figure 5.7, we see that a green segment fails to cross a blue segment only if the line extending the green segment lies to the right of the entire blue segment.

At the same time, the conclusions from Main Calculation 4 say that the right endpoint of a blue segment cannot lie in the line extending a green segment. The green segment is always $(4)_1^*$ and the blue segment can be any β^* with $\beta \in [0, 1]_1$. But β^* varies continuously. Hence, from what we have already said, either β^* crosses $(4)_1^*$ for all choices of β or for none of them. But $(0)^*$, which is just the horizontal bend, definitely crosses the green segment. Hence all the blue segments cross the green segment. This shows that $(4_1, 1_1, 0)$ is geometrically good.

The same argument works with *red* in place of *green*, and shows that $(4_1, 1_1, 0)$ is geometrically good. ♠

Lemma 5.4 $(1_i, 4_j, 6)$ *is geometrically good unless* $\tau_j \to_1 \tau_i$. *Again,* $i, j \in \{1, 2\}$.

Proof: By symmetry, it suffices to take $j = 1$. This time our argument refers to Figure 5.4. Looking at Figure 5.4, we see that a blue segment crosses a green segment unless the line extending the green segment lies to the right of

the blue segment. We have already shown that $(1)_1^*$ and $(4)_1^*$ cross. Consider the continuous family of segments of the form β^*, where β sweeps through $[4,6]_1$. The first segment crosses $(1)_1^*$. Therefore, from what we have observed about blue and green segments, either every β^* crosses $(1)_1^*$, or one of them contains the right endpoint of $(1)_1^*$. In this latter case, we have $\tau_1 \to_1 \tau_1$ by definition. Hence $(1_1, 4_1, 6)$ is geometrically good unless $\tau_1 \to_1 \tau_1$.

By symmetry, any segment of the form β^* with $\beta \in [0,1]_2$ is magenta. The same argument as above, with *magenta* replacing *blue*, impliage that $(1_2, 4_1, 6_1)$ is good unless $\tau_1 \to \tau_2$. ♠

Lemma 5.5 *The triple* $(4_2, 4_1, 0)$ *is geometrically good unless* $\tau_1 \to_4 \tau_2$. *Likewise,* $(4_1, 4_2, 0)$ *is geometrically good unless* $\tau_2 \to_4 \tau_1$.

Proof: We prove the first result. The second one follows from interchanging the roles played by the indices. Our argument refers to Figure 5.5. Note that an orange segment always crosses a red segment unless that line extending the orange segment lies above the red segment. We have already seen that $4_1|4_2$. The same argument as in Round 2 now shows that $(4_2, 4_1, 0)$ is geometrically good unless $\tau_1 \to_4 \tau_2$. ♠

This covers all triples we need to consider in order to establish Geoemtrical Goodness for a tame pair (τ_1, τ_2) of trapezoids. This completes our reduction of the Topology Lemma to tensegrity calculations.

6 Numerical Computations

6.1 Parametrizing Tensegrities

For all our calculations we have a compound tensegrity satisfying certain constraints. The trensegrity consists of either 2 or 4 quadrilaterals and it satisfies certain constraints. We want to parametrize all the tensegrities which satisfy the given constraints. There is a tradeoff between the simplicity of the parametrization and the efficiency of the calculation. We first describe a very simple but inefficient parametrization, and then we describe the parametrization we actually use.

The Simple Approach: All the compound tensegrities of interest to us certainly have all their points in the interval $[-2,2]^3$. Thus we can parametrize a compound tensegrity having k quadrilaterals as a point in cube $X = [-2,2]^{6k+6}$. We then have some constraint set $Y \subset X$ and a function $f : X \to \boldsymbol{R}_+$. We want to bound $f|_Y$ from below. This approach is not very efficient because it is hard to directly sample points in Y. It is better to parametrize just Y.

The Interpolation Function: We define

$$\iota(a,b,r) = (1-r)a + rb. \tag{12}$$

We can think of $\iota(a,b,*)$ as being a map from $[0,1]$ to the interval $[a,b]$. We will use this map repeatedly.

The Bends: Now we start on a more subtle approach that works better. The tensegrities we consider have either 3 or 4 bend images Let us say that there are $\ell+1$ bend images, $\beta_0,...,\beta_\ell$. So, we have $\ell = 2$ or $\ell = 3$. We arrange the indices so that $j = 0$ and $j = \ell$ correspond to the horizontal and vertical bend images.

We make one remark about notation. We let k stand for the complexity. When the complexity equals 2 we have $\ell = 2$. When the complexity equals 4 have $\ell = 3$. The reason is that in the complexity 4 case we are considering cyclic tensegrities. This drops the bend image count by 1. This is why we are using ℓ here in place of k.

Given a point $(r_0,...,r_\ell) \in [0,1]^{\ell+1}$ we set

$$s_0 = \iota(0,1/2,r_0), \qquad s_j = \iota(-1,1,r_j), \;\; j = 1,...,\ell,$$

$$s_\ell = \iota(a(s_0), b(s_0), r_\ell). \tag{13}$$

Here $a(s_0)$ and $b(s_0)$ are obtained as follows. We intersect the vertical line $x = s_0$ with the quadrilateral $\widehat{\Omega}$ from Theorem 3.1 and take the upper and lower limits. The precise formulas are

$$a(s_0) = (2/3)s_0 - (1/\sqrt{3}), \quad b(s_0) = \min(b_1(s_0), b_2(s_0)),$$

$$b_1(s_0) = \min((2/3)s_0 - (1/2), \quad b_2(s_0) = (4/3)s_0 - (1/\sqrt{3})). \tag{14}$$

In our definition of s_0 we could take $s_0 \in [0, \frac{1}{4}(\sqrt{27} - \sqrt{11})]$, but it is simpler just to take $s_0 \in [0, 1/2]$. In terms of Figure 2.3, we are setting $b = s_0$ and $t = s_k$.

This approach allows (s_0, s_ℓ) to sample all of $\widehat{\Omega}$. The other bend slopes $s_1, ..., s_{k-1}$ are each allowed to lie in $[-1, 1]$, which certainly covers all the possibilities.

The Offsets: Given $(r_1, r_2) \in [0, 1]^2$ we define

$$x = \iota(0, 1/18, r_1), \qquad y = \iota(-1/30, 1/30, r_2). \tag{15}$$

These are the coordinates of the vector shown in Figure 2.3.

The coordinates (b, t, x, y) determine the placement of the horizontal and vertical bend images. The horizontal bend image is the segment with endpoints $(0, 0)$ and $(-\sqrt{1 + b^2}, 0, 0)$. The vertical bend image is centered at (x, y) and has length $\sqrt{1 + t^2}$.

Complexity Two Case: For the prototypical calculation and the confiner calculations we are working with a tensegrity having 2 quadrilaterals. In this case it only remains to describe how we parametrize the middle bend. The parameter s_1 determines the length of this bend. Given a point $(r_1, r_2, r_3, r_4, r_5) \in [0, 1]^5$ we set

$$m_j = \iota(-1, 1, r_j), \;\; j = 1, 2, 3, \quad \theta = \iota(\theta_1, \theta_2, r_4), \quad \phi = \iota(-\pi/4, \pi/4.r_5). \tag{16}$$

Here (m_1, m_2, m_3) is the center of the middle bend. Looking at the figures in the previous chapter (and looking at our computer program) the constraints we have placed cover all the reasonable cases.

The interval $[\theta_1, \theta_2]$ is the angle range for the middle bound. For instance, in the prototypical calculation we have $\theta_1 = \theta_2 = 2\pi/12$. The angle ϕ is the

angle that the middle bend makes with the XY-plane. We call ϕ the *spatial angle* Our coordinates specify the center, length, pitch, and spatial angle, and thus we can specify the middle bend exactly. (In a rigorous calculation we perhaps would want to avoid the trig functions needed to specify the middle bend.)

So a point in $[0,1]^3 \times [0,1]^2 \times [0,1]^5 = [0,1]^{10}$ completely specifies any tensegrity of this type we need to consider.

Complexity Four Case: In this case we have two middle bend images to worry about, β_1' and β_2'. We parametrize β_1' just as the middle bend image above. We specify β_2' using a point $(r_1, r_2, r_3, r_4) \in [0,1]^4$. We already know the length of β_2': It is $\sqrt{1+s_2^2}$.

The parameter r_1 determines a point $p \in \beta_2'$ as follows:

$$p = \iota(\partial_0 \beta_2', \partial_1 \beta_2', r_1). \tag{17}$$

Here $\partial_j \beta_2'$ for $j = 0, 1$ are the endpoints of β_2'. What we are doing here is interpolating between the two endpoints of β_2'. (We are also abusing the notation, because ι is a scalar valued function and we are taking the obvious vector-valued extension.)

We set $h = \iota(-1, 1, r_2)$. This determines the distance that p lies above the corresponding point $q \in \beta_1'$. What we mean here is that p and q project to the same point in the XY plane.

We then define the angles θ and ϕ as above. We use the following intervals $[\theta_1, \theta_2]$:

1. In Main Calculation 1 we take $\theta \in [5\pi/12, 7\pi/12]$.
2. In Main Calculation 2 we take $\theta \in [5\pi/12, 7\pi/12]$.
3. In Main Calculation 2 we take $\theta \in [1\pi/12, 3\pi/12]$.

We will not rigorously justify why these intervals cover all relevant cases, but we note that the midpoints of the intervals correspond to the pitches we get when we take the equilateral example from Figure 1.1. If one does not agree with these choices, one can take wider intervals.

We need a point in $[0,1]^4$ to specify β_2', so all in all a point in $[0,1]^{14}$ specifies a tensegrity of complexity 4 that we use for the Main Calculations.

The Distinguished Slice: We notice that our Confinng calculations converge much more quickly to the presumed minimum when we work with planar examples. For the middle bend in a complexity 2 case, the parameters (m_3, ϕ) determine the extent to which this segment does not lie in the XY-plane. To get a planar example, we set $m_3 = \phi = 0$. This drops us down to $[0,1]^8 \subset [0,1]^{10}$. We call this copy of $[0,1]^8$ the *distinguished slice*.

In the complexity 4 case, the parameters (h, ϕ) determine the extend to which β_2' sticks out of the XY-plane. So, we set these parameters to 0. This (when combined with the steps taken for β_1) drops us down to $[0,1]^{10} \subset [0,1]^{14}$. We call this copy of $[0,1]^{10}$ the *distinguished slice*.

So far I have only experimented with the Main Calculations on the distinguished slice. Guided by what happens with the Confining Calculations I think that very likely that if the Main Calculations work on the distinguished slice then they will also work on the total space. My impression of what is going on, in all cases, is that the minimum capacity configuration lies in the distinguished slice, so it is faster just to move within this slice. Of course, a rigorous proof will require a consideration of the whole space. I plan to code up the Main Calculations for the whole space soon.

6.2 Hill Climbing Algorithms

For all our calculations we have our capacity function $\kappa : [0,1]^N \to \boldsymbol{R}_+$. More specifically...

- For the Prototypical Calculation we have $N = 9$ rather than $N = 10$ because we have $\theta_1 = \theta_2 = 2\pi/12$.
- For the Confiner Calculations we have $N = 10$. In the cases involving Θ_{44} we can take $N = 9$.
- For the Main Calculations we have $N = 14$.

We have some lower bound λ and we are trying to show that

$$\min_{[0,1]^N} \kappa \geq \lambda. \tag{18}$$

More precisely, with our numerical calculations we are searching for possible counter-examples to this inequality and we want the most efficient way of finding them. If we don't find any counter-examples after running an efficient

optimizing algorithm for a reasonably long time, we guess that there are no counter-examples – i.e., the inequality is true.

The optimization algorithm we use is similar to gradient flow, except that we do not actually compute a gradient. It would not really be a good idea to compute the gradient here because the function κ is not differentiable. It is piecewise algebraic. One could perhaps do gradient flow on the pieces, so to speak, and then assemble the results, but this is not what we do.

We first describe the simplest possible thing. We evaluate κ on random points of $[0,1]^N$ and simply keep track of the minimum value attained. If this value always exceeds λ then we have some evidence for Equation 18. This approach does not work so well in high dimensions. Even if we sample 10^N points are are really only searching on a scale of about $1/10$.

One way to do much better is to use gradient flow (or some version of Newton's method) to locate minima. One problem is that our function κ is only piecewise smooth. Another problem is that we are lazy and do not want to work out the very complicated formulas for the gradient where it is defined. The optimization algorithm we use is similar to gradient flow, except that we do not actually compute a gradient.

We start with some initial point $r_0 \in [0,1]^N$. We then pick a small random step size s and a random point $r_1 \in [-1,1]^N$. We then replace r_0 by $r_0 + sr_1$. If some coordinate of r_1 is negative, we set that coordinate equal to 0. If some coordinate of r_j exceeds 1, we set it equal to 1. In other words, we take a small step in a random direction and then retract to $[0,1]^N$ if necessary. If $f(r_1) < f(r_0)$ we replace r_0 with r_1 and repeat. This produces a sequence of points $r_0, r_1, r_2, ...$ on which κ is decreasing. The hope is that this sequence converges to a global minimum.

This algorithm is a reasonable approximation to gradient flow. At least at points where κ is smooth, a small step in a random direction has a 50 percent change of moving in the direction of the gradient. and a somewhat lower but still decent chance of making a small angle with the gradient. The concentration of measure phenomenon – randomly chosen vectors in high dimensions are usually almost orthogonal – does not really kick in when N is reasonable small. So, the algorithm we use is not so different from gradient flow, at least when s is small.

The art to the algorithm comes from choosing s appropriately. If we choose s too large, then we are essentially bouncing around $[0,1]^N$ randomly, and we are not doing much better than random sampling. If we choose s too large, then we are decently approximating gradient flow but the process

takes a long time and also can get stuck in local minima. The best approach (I think) is to choose s randomly in $[0,1]$ but with a bias towards small values. This way the process is typically doing something like gradient flow and occasionally makes big jumps to possibly escape local minima.

Here is a concrete example: I find that setting

$$s = \frac{1}{10}(f_0 - 2\sqrt{3})\rho$$

works pretty well for the Main Calculations. Here κ_0 is the current value of κ and ρ is some random number in $[0,1]$.

6.3 The Confiner Calculations

There are two main troubles with the kind of calculations we run. The first trouble is that the calculation could get stuck in a local minimum. The second trouble is that the calculation could have a very shallow global minimum which is lower than the bound we seek. So, even if the process finds the global minimum, it might happen that it takes a long time to get very close.

There is one other simplification I often made when running the Confiner Calculations. I noticed that when I constrain the endpoint of a segment in each of the subspace, the minimum capacity configuration occurs at the boundary. The subspaces are all chosen so that the tensegrity has to stretch, so to speak, in order to meet the constraint. The object is the least stretched for boundary points. Anyway, the process seems to find the minimum much more quickly with the stronger constraint. Most of my experiments involved the stronger constraint, after I checked that the stretching phenomenon seemed to be in place.

The idea behind the Confiner Calculations is quite simple. In order for the given point to lie outside the given prism halfplane, the tensegrity has to be "stretched" a lot. The stretching bumps up the capacity beyond $\sqrt{3}+(1/24)$. Figure 6.1 illustrates this with two examples. The first example shows a tensegrity near the minimum for the given constraint. The capacity here is about $\sqrt{3} + 0.14$. The second example shows a much more radical stretch, where the capacity is about $\sqrt{3} + 1.5$. The red edges correspond to the vertical sides of the associated planar quadrilaterals.

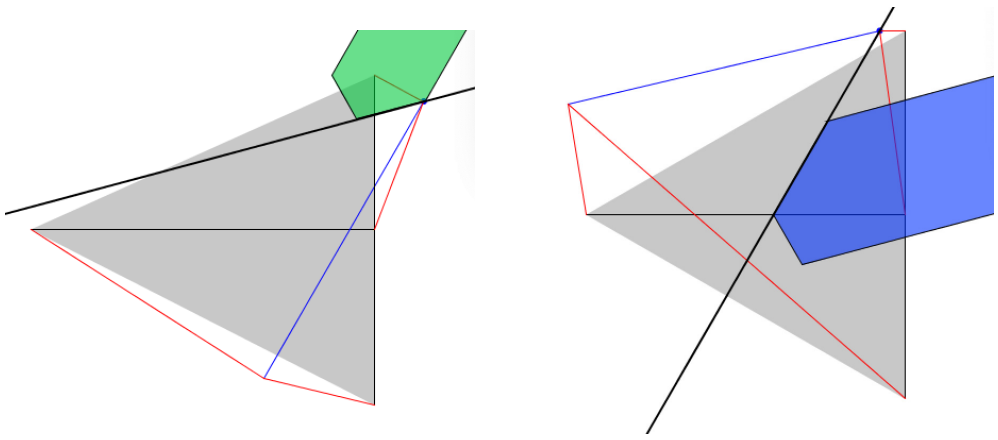

Figure 6.1: Calculation Records

Figures 6.2-6.3, which correspond to Figures 5.4-5.6 (with redundant pieces removed) give some idea of what the calculations are like. An edge labeled by, say, 0.1, indicates that after running the algorithm for a while with respect to the line extending the edge we conclude that the minimum of the capacity function is at least $\sqrt{3}+0.1$. Our target is $\sqrt{3}+(1/24) < \sqrt{3}+0.05$, so even in this case we have a healthy margin for success. The red lines indicate the most delicate calculations, and the orange lines indicate the next more delicate calculations. Everything else is pretty crudely true. I ran each of the calculations associated to the red edges at least 2 hours, and I ran the ones associated to the small red edge labeled 0.09 for about 8 hours.

There is something I want to say about the calculation where the bound is 0.09. When I run this calculation I notice that when κ tends towards 0.09 the corresponding b value (the slope of the bottom bend) tends to the upper bound $(\sqrt{27}-\sqrt{11})/4$. This accounts for the elongated appearance of the grey triangle on the left side of Figure 6.1. If we wanted, we could do additional calculations to strengthen Statement 3 of Theorem 3.1 to say that we must have $b < 1/3$ whenever we have a tame pair (τ_1, τ_2). This would give us a more robust calculation result in this case. What I am saying is that there are tricks available to make these calculations true by wider margins.

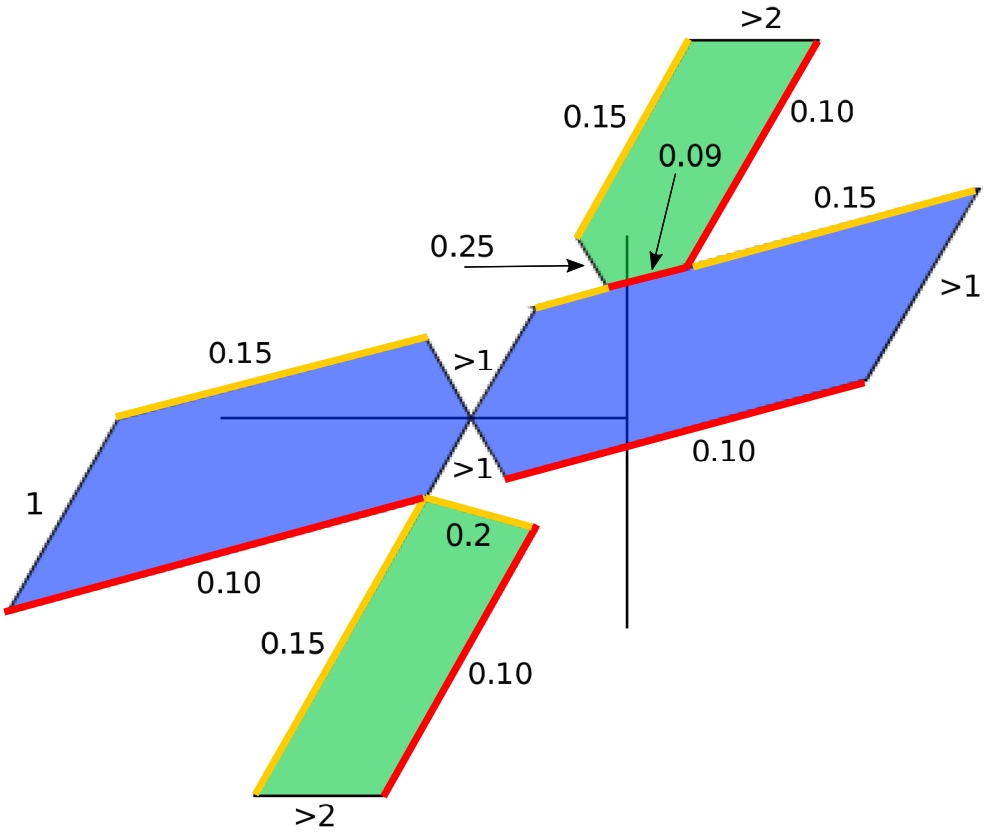

Figure 6.2: Calculation Records

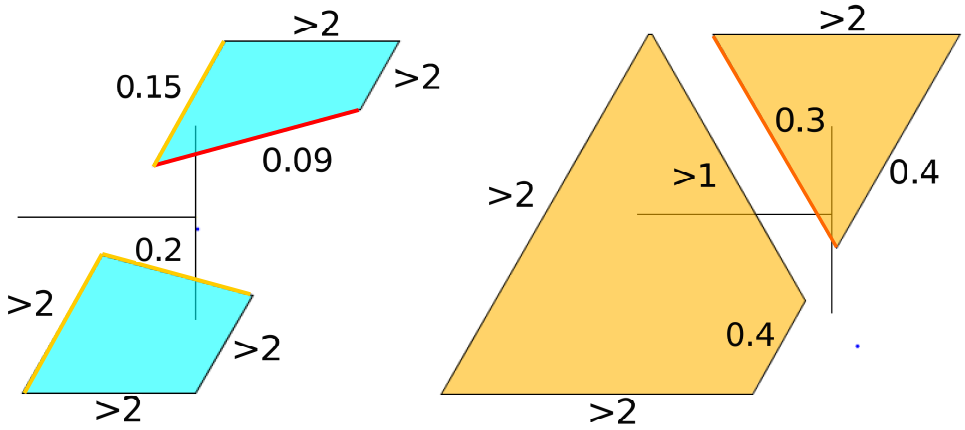

Figure 6.3: Calculation Records

6.4 The Main Calculations

Now we say something about the Main Calculations 1,2,3. These calculations also assert that the unique minimum occurs when the union of the horizontal and vertical bends has convex hull Δ_0. Here Δ_0 is the equilateral triangle of semi-perimeter $\sqrt{3}$, which appears on the right side of Figure 1.1. When I run Calculations 1 and 2 for a long time, the picture seems to converge to this situation. Thus strongly suggests that the calculation is really finding the global minimum.

Calculation 3 is a bit wonky. If I run it in an unrestricted way, the point falls into a local minimum having capacity about $2\sqrt{3} + 2$, which is vastly larger than the supposed true minimum. The easiest way to coerce the points toward the presumed true minimum is to constrain the slope of the bottom bend to be less than 1/3. In a more polished version of this draft I will say more about this.

7 References

[**AHLM**] A. Malaga, S. Lelievre, E. Harriss, P. Arnoux, *ICERM website*: https://im.icerm.brown.edu/portfolio/paper-flat-tori/ (2019)

[**CB**], R. Connolly, A Back, *Mathematics and Tensegrity*, in **American Scientist**, Vol 82 **2**, (March-April 1998) pp 142-151.

[**CF**] Y. Chen and E. Fried, *Mobius bands, unstretchable material sheets and developable surfaces*, Proceedings of the Royal Society A, (2016)

[**CKS**] J. Cantarella, R. Kusner, J. Sullivan, *On the minimum ropelength of knots and links*, Invent. Math. **150** (2) pp 257-286 (2003)

[**FT**], D. Fuchs, S. Tabachnikov, *Mathematical Omnibus: Thirty Lectures on Classic Mathematics*, AMS 2007

[**HF**, D.F. Hinz, E. Fried, *Translation of Michael Sadowsky's paper 'An elementary proof for the existence of a developable MÖBIUS band and the attribution of the geometric problem to a variational problem'*. J. Elast. 119, 3–6 (2015)

[**HW**], B. Halpern and C. Weaver, *Inverting a cylinder through isometric immersions and embeddings*, Trans. Am. Math. Soc **230**, pp 41.70 (1977)

[**MK**] L. Mahadevan and J. B. Keller, *The shape of a Mobius band*, Proceedings of the Royal Society A (1993)

[**Sa**], M. Sadowski, *Ein elementarer Beweis für die Existenz eines abwickelbaren MÖBIUSschen Bandes und die Zurückführung des geometrischen Problems auf einVariationsproblem*. Sitzungsberichte der Preussischen Akad. der Wissenschaften, physikalisch-mathematische Klasse 22, 412–415.2 (1930)

[**S1**] G. Schwarz, *A pretender to the title "canonical Moebius strip"*, Pacific J. of Math., **143** (1) pp. 195-200, (1990)

[**S2**] R. E. Schwartz, *The* 5 *Electron Case of Thomson's Problem*, Journal of Experimental Math, 2013.

[**S3**] R. E. Schwartz, *An improved bound on the optimal paper Moebius band*, preprint, 2020

[**W**] S. Wolfram, *The Mathematica Book, 4th Edition*, Wolfram Media and Cambridge University Press (1999).