Vertex-Minimal Paper Tori

Richard Evan Schwartz *

August 7, 2025

Abstract

A *paper torus* is an embedded polyhedral torus that is isometric to a flat torus in the intrinsic sense. We prove that there does not exist a paper torus with 7 vertices, and that there does exist a paper torus with 8 vertices. This settles the question of the minimum number of vertices needed for a paper torus.

1 Introduction

1.1 Context

A *flat torus* is a quotient of the form $\boldsymbol{R}^2/\Lambda$, where Λ is a lattice of translations of $\boldsymbol{R}^2$. More concretely, a flat torus is what you get when you identify the pairs of parallel sides of a parallelogram by translations.

A *polyhedral torus* is a piecewise affine embedding $\phi : T \to \Omega \subset \boldsymbol{R}^3$. Here T is a flat torus that has been triangulated, and ϕ is a continuous embedding that is affine on each triangle of the triangulation. If, additionally, ϕ is an affine isometry on each triangle, we say that ϕ is *flat*. When ϕ is flat, ϕ gives an isometry from T to the surface $\phi(T) \subset \boldsymbol{R}^3$ equipped with its intrinsic path metric. However, the image $\phi(T)$ could be crinkled up in space in a complicated way. Put another way, in the flat case, ϕ tells how to build a torus in $\boldsymbol{R}^3$ out of finitely many triangles so that the cone angle around each vertex is 2π.

Definition: A *paper torus* is a flat embedded polyhedral torus.

* Supported by N.S.F. Research Grant DMS-2505281

Surprisingly, paper tori exist. The 1960 paper of Y. Burago and V. Zalgaller [**BZ1**] gives the first constructions. The 1985 paper [**BZ2**] proves that one can realize every isometry class of flat torus as a paper torus. Their construction produces examples with thousands of faces. The 2024 preprint of F. Lazarus and F. Tallerie [**LT**] gives a universal combinatorial type of triangulation which does the job simultaneously for all isometry types.

The work (in progress) [**ALM**] of P. Arnoux, S. Lelievre, and A. Malaga gives a different construction which achieves every isometry class of flat torus as a paper torus, using far fewer faces than in the other cited works. Specifically, they prove that the embedded polyhedral tori described by U. Brehm [**Br**] in 1978 can be made flat. In 2025, Vincent Tugayé [**T**], discovered a 9-vertex paper torus. This is the record to date.

1.2 Seven Vertices

One needs at least 7 vertices to make a paper torus because there is no 6 vertex triangulation of a torus. There is a 7-vertex triangulation of a torus, and it is unique up to combinatorial isomorphism. This triangulation, sometimes called the *Moebius triangulation*, is shown in Figure 1.1.

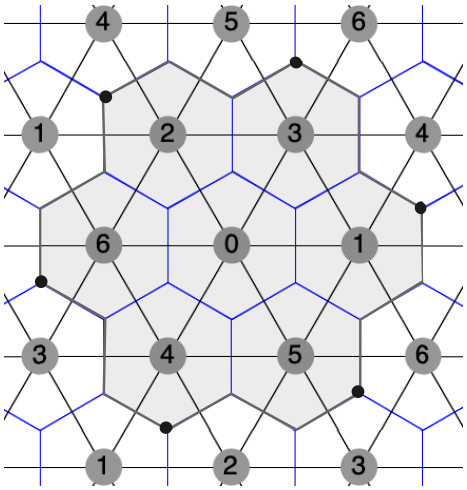

Figure 1.1: The universal cover of the Moebius triangulation

The Moebius triangulation gives the famous embedding of the complete graph K_7 in a torus. We are drawing part of the universal cover as well as a fundamental domain consisting of 7 hexagons.

In 1949, Á. Császár [**Cs**] showed that the Moebius torus has a polyhedral embedding. Császár's polyhedral embedding Ω is specified by giving coordinates for the vertices $P_0, ..., P_6$. With our labeling, the coordinates are

$$(0,0,15), (3,3,0), (-3,-3,1), (-1,-2,3), (1,2,3), (3,3,1), (3,-3,0).$$

This torus has 2-fold rotational symmetry and its convex hull contains 5 of the 7 vertices. See e.g. [**G**], [**HLZ**], and [**L**] for more detail and a discussion of related topics. The Császár torus is not flat. Could some other embedding of the Moebius torus be flat? The 2019 Ph.D. thesis of P. Quintanar Cortés, [**QC**] makes significant partial progress on this problem. Here is our first result.

Theorem 1.1 *There does not exist a 7-vertex paper torus.*

A key step in the proof is the following result.

Lemma 1.2 (Hull) *Suppose Ω is an embedded 7-vertex polyhedral torus, not necessarily flat. Let H be the convex hull of Ω. Suppose Ω has all 7 vertices in ∂H. Then Ω has at least one vertex P such that all 6 triangles of Ω incident to P are in ∂H.*

I will deduce from the Hull Lemma that a flat polyhedral torus must have a vertex P contained in the interior of the convex hull. I will then use Crofton's formula to show that the cone angle at P exceeds 2π. This contradiction finishes the proof of Theorem 1.1.

The Hull Lemma is a consequence of a stronger result contained in the 1991 paper [**BE**] of J. Bokowski and A. Eggert. Bokowski and Eggert use oriented matroids to give classify the combinatorial types of 7-vertex polyhedral tori. One consequence of their classification is the *Hull Theorem*: there are no embedded 7-vertex polyhedral tori at all having all 7 vertices on their convex hull boundary. (See the remarks after [**BE**, Theorem 3.7]. This is my name for their result.) In my opinion, had these authors asked the flatness question for 7-vertex polyhedral tori, they most likely would have been able to answer it.

In order to keep this paper self-contained, I will give a light and entirely combinatorial computer-assisted proof of the Hull Lemma. For the interested reader, I will also take the extra step and deduce the Hull Theorem from our combinatorial analysis and a bit of projective geometry. Only the Hull Lemma is needed for our proof of Theorem 1.1, though the proof of Theorem 1.1 has one fewer step if we do use it.

1.3 Eight Vertices

Here is the main result.

Theorem 1.3 *There exists an* 8*-vertex paper torus.*

The example has 2-fold rotational symmetry and all 8 vertices in its convex hull boundary. Let me discuss how I found it.

Motivated by our proof of Theorem 1.1, I first tried to rule out the existence of 8-vertex paper tori having all 8 vertices on their convex hull boundary. Such a result would not be as decisive in the 7-vertex case, but it seemed like a good step in a potential proof of non-existence. I performed a combinatorial analysis similar to what I did in the 7-vertex case. This analysis ruled out most but not all examples. Afterwards, I could see experimentally that indeed this kind of 8-vertex embedded polyhedral tori exist. At that point, my attention turned to studying these kinds of examples in detail.

There are 7 combinatorial types of triangulation of the torus having 8 vertices. Of these 7 there is (in my opinion) a *worst triangulation* and a *best triangulation*. The worst triangulation is obtained from the Moebius triangulation by subdividing one of the triangles. The best triangulation has degree sequence $6, ..., 6$ and is vertex-transitive. This triangulation is very much like the Moebius triangulation. Figure 1.2 shows part of the universal cover of this triangulation, as well as a fundamental domain. I will explain the meaning of the blue triangles below.

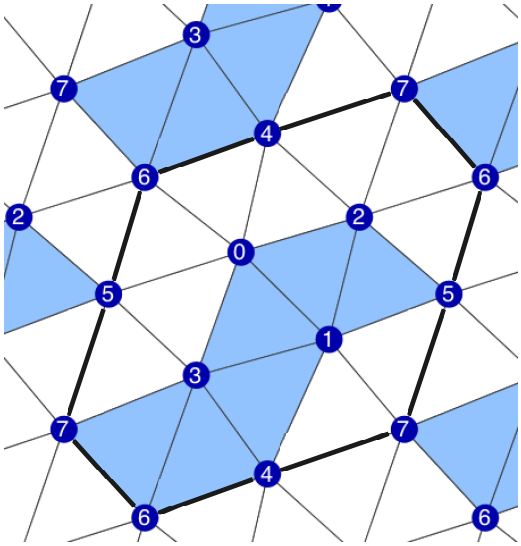

Figure 1.2: The universal cover of the best 8 vertex triangulation

For all but the worst triangulation, I noticed that one can place 8 points on the unit sphere and then fill in the triangles to get an embedded 8-vertex polyhedral torus. My computer program repeatedly places points at random on (the surface of) an ellipsoid and then checks the embedding condition until an embedded example turns up.

Next, I ran a hill-climbing algorithm which started with random collections of points on the unit sphere and then repeatedly perturbed them – not necessarily keeping them on the unit sphere – with the goal of minimizing an objective *flatness function*. The flatness function computes the maximum deviation of a cone angle from being 2π. The short story is that I got lucky and this worked. The longer story is that I made my good luck by incorporating some tricks into the algorithm, tricks that I gradually noticed after a lot of experimentation. Here are 4 key tricks.

1. It was useful to first perturb the points within an ellipsoid surface that contains them, and then to perturb the ellipsoid. This worked much better than just perturbing the points at random, probably because the method lets you separately control the normal and tangential variations.

2. It was extremely useful to control the combinatorics on the convex hull boundary. The convex hull boundary has 12 faces, some of which – say K of them – are also faces in the triangulation of the torus. Call K the *face number*. I noticed that the examples with large face number turned up much more often, but they never could lead to flat examples. (There always appeared a star in the torus that was part of a convex fan.) So I eventually constrained my search to favor examples with low face number. The examples with low face number are more crinkly.

3. It helped to bound some auxiliary quantities, like the minimum angles and dihedral angles, to prevent degenerations. Bounds of about 0.00001 on the dihedral angles worked well.

4. It helped to force the examples to have 2-fold rotational symmetry. This simplification came later. After I found some asymmetric examples, Samuel Lelievre asked me about symmetric examples and I tweaked the program. The program works more reliably after incorporating the symmetry.

With all these constraints in place, and then a lot of playing around, I found some embedded and 2-fold symmetric examples of face number 6 which are flat up to 10^{-16}. I took the nicest one and simplified the coordinates, feeding it back into the program to make it near-flat again. Then I implemented a high precision version of Newton's method and improved my example to one that was flat up to 10^{-400}. Here are the coordinates of this extremely near flat example, truncated after 32 digits.

$$\begin{array}{lll} +0.755 & +0.650 & z_0 \\ -0.455 & +0.345 & z_1 \\ -0.170 & +1.140 & z_2 \\ +0.455 & -0.345 & z_1 \\ -0.755 & -0.650 & z_0 \\ -0.090 & +0.665 & 0 \\ +0.170 & -1.140 & z_2 \\ +0.090 & -0.665 & 0 \end{array} \qquad \begin{aligned} z_0 &= 0.9805\ 0571\ 5859\ 7793\ 5561\ 6538\ 2008\ 5693 \\ z_1 &= 0.9902\ 8162\ 4334\ 3054\ 2934\ 3176\ 1585\ 8328 \\ z_2 &= 0.9765\ 3883\ 4703\ 1231\ 7624\ 1842\ 4567\ 2434 \end{aligned} \tag{1}$$

The 0th row of this array gives coordinates for vertex 0, and so on. This example has 2-fold rotational symmetry. The symmetry is the permutation (04)(13)(26)(57). When we just use the numbers in Equation 1 we get an embedded example that is flat up to an error of about 10^{-32}. (See §4.) Here is some more precision.

$$\begin{aligned} z_0 &= 0.9805\ 0571\ 5859\ 7793\ 5561\ 6538\ 2008\ 5693\ 0055\ 3244\ 8886\ 6182\ 3764\ 0925\ 1031\ 3998 \\ &\quad\ 5289\ 6900\ 0219\ 1976\ 3052\ 9441\ 6054\ 1397\ 2979\ 6612\ 3914\ 4736\ 3149\ 9461\ 8320\ 0299 \\ z_1 &= 0.9902\ 8162\ 4334\ 3054\ 2934\ 3176\ 1585\ 8328\ 0396\ 3159\ 0105\ 1356\ 1699\ 8363\ 9853\ 0130 \\ &\quad\ 1662\ 2960\ 4417\ 0122\ 2479\ 7351\ 1673\ 2074\ 0218\ 8931\ 2952\ 9639\ 3940\ 0849\ 9822\ 4101 \\ z_2 &= 0.9765\ 3883\ 4703\ 1231\ 7624\ 1842\ 4567\ 2434\ 3419\ 2723\ 3847\ 8043\ 1363\ 4317\ 6666\ 5944 \\ &\quad\ 1656\ 5345\ 8266\ 6325\ 9911\ 2468\ 3851\ 5637\ 1435\ 4186\ 4964\ 4745\ 4121\ 8471\ 8234\ 9025 \end{aligned} \tag{2}$$

Again, this example has face number 6. The blue triangles in Figure 1.2 correspond to the 6 triangles of the torus that lie in the convex hull boundary. Figure 1.3 shows a numerical computation of the intrinsic geometric structure of the triangles. We are showing part of the universal cover of the triangulation.

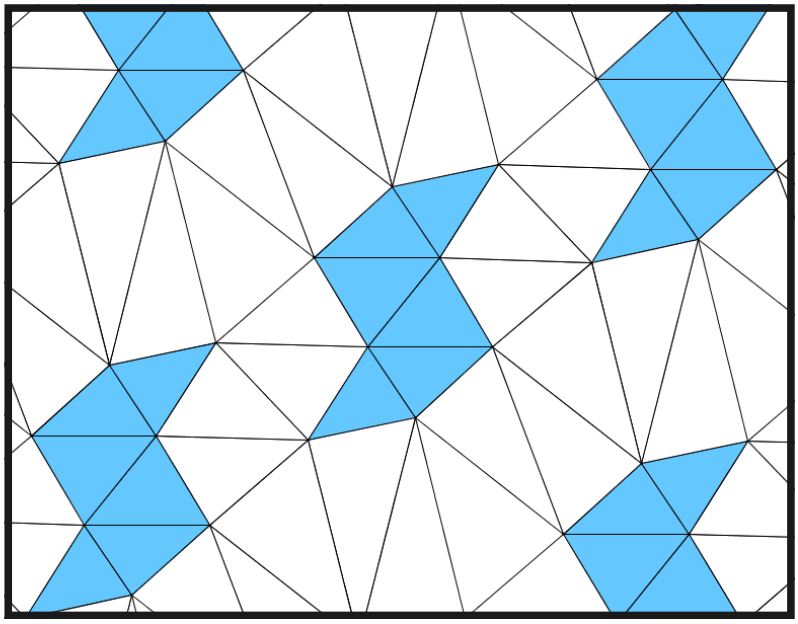

Figure 1.3: The intrinsic structure.

Figure 1.3 looks quite tame, but the extrinsic geometry is quite beautiful and intricate. Figure 1.4 shows one of the slices of the object by the XY plane.

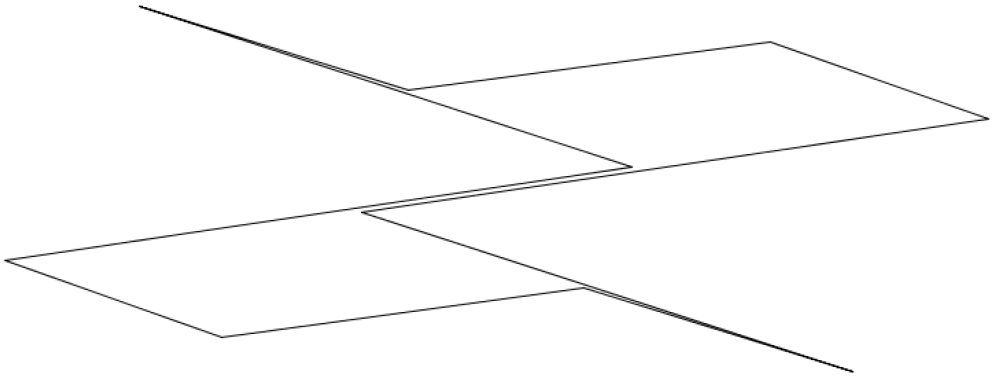

Figure 1.4: The XY-slice of the example

This slice is very spiky and sort of "barely embedded", but the embeddedness is quite robust compared to the plot precision of 10^{-16}.

Figure 1.5 shows the slice by the XZ plane. You are looking at two nested embedded polygons. The annulus between these two polygons is a slice of the

solid handlebody bounded by the torus. The two "parallel vertical segments" on the right are neither parallel nor segments. Each is a bigon made from two smaller and symmetrically placed segments which meet at an angle that deviates from π by an amount on the order of 0.01.

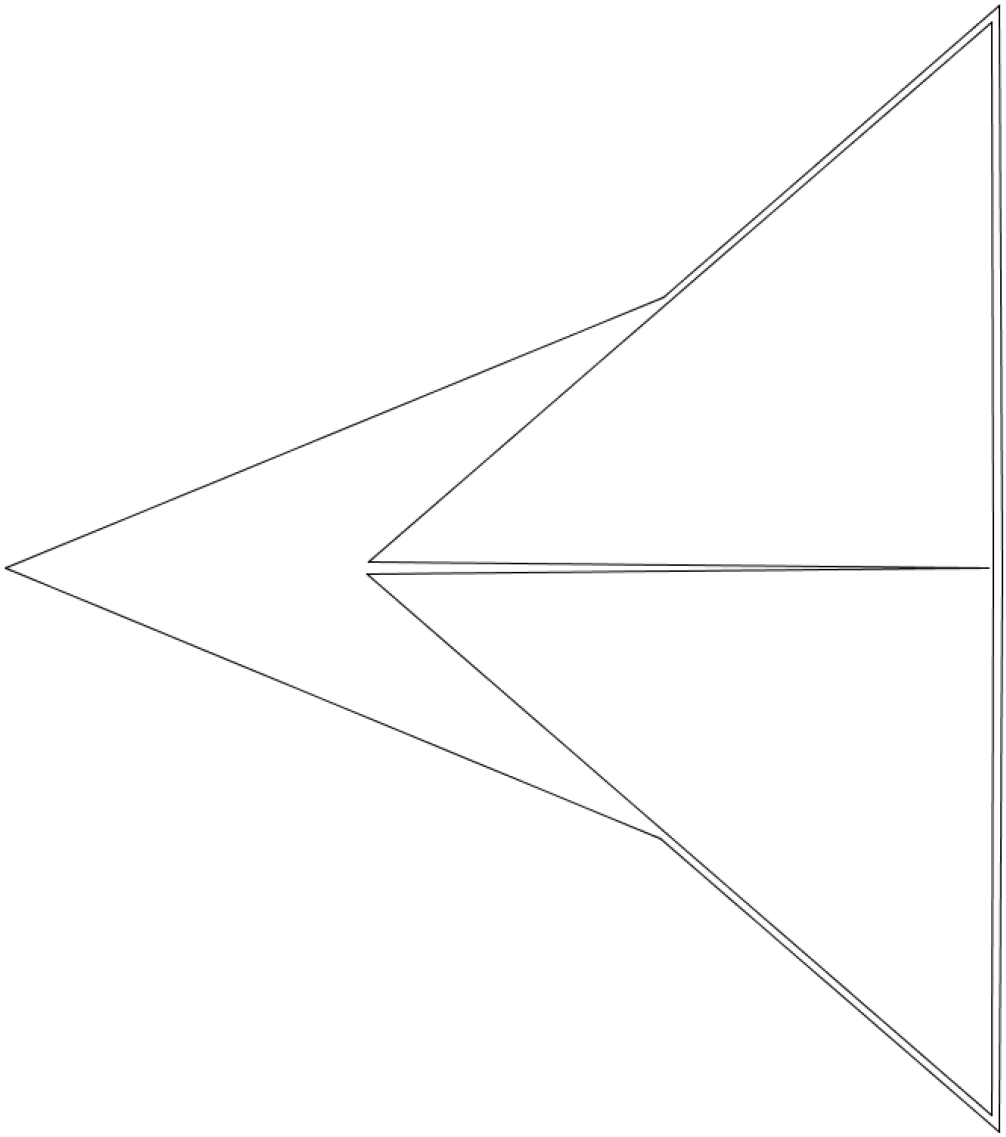

Figure 1.5: The XZ slice of the example

Figure 1.6 shows the slice by the YZ plane. We have shown some close-ups to indicate that the two polygons are disjoint and embedded, even though without magnification it looks close.

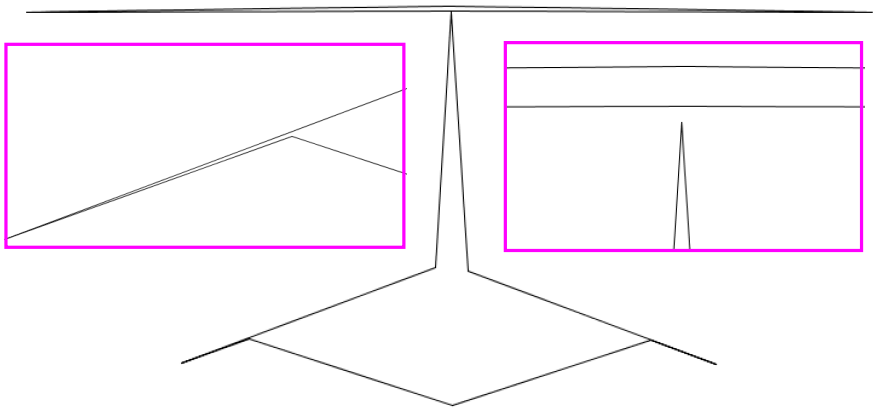

Figure 1.6: The YZ slice of the example.

Figure 1.7 shows a slice in a random direction. Once again, we have an embedded loop.

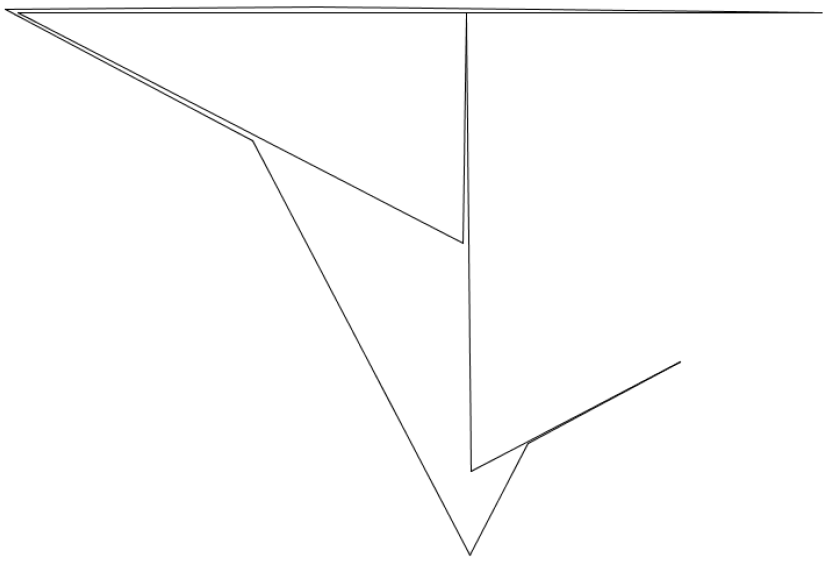

Figure 1.7: A random slice of the example

Once we fix this direction we get a *level set foliation* of the torus by curves that lie in planes perpendicular to this direction. In figure 1.8 we have drawn the level set foliation in the intrinsic flat structure and then lifted the picture to the universal cover. The highlighted loop in Figure 1.8 corresponds to the loop shown in Figure 1.7. This Figures 1.7 and 1.8 show two views of the same polygon, one in the intrinsic geometry and one as it appears in space. Comparing these two pictures gives you a sense of the really weird and twisted way our example sits in space.

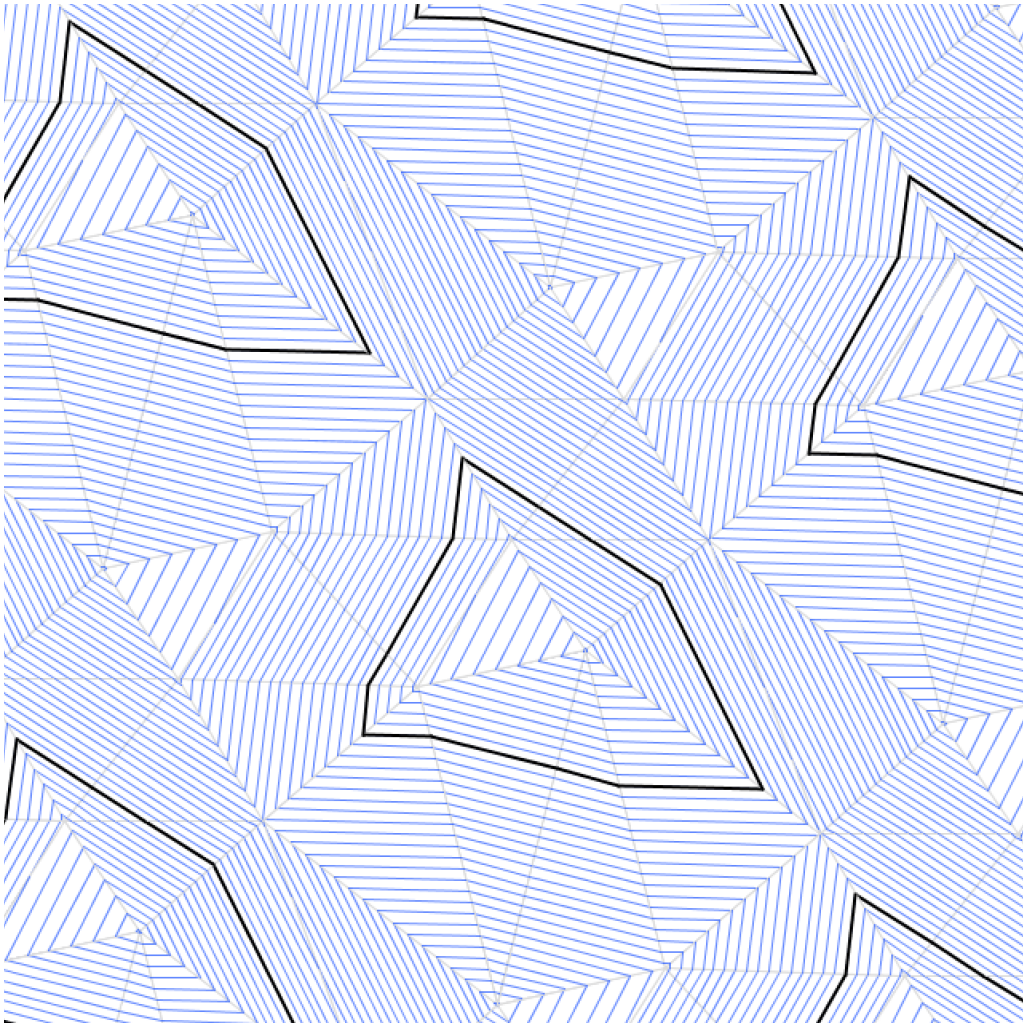

Figure 1.8: Intrinsic view if the level set foliation

Now let me explain how I convert the example above into a proof of

Theorem 1.3. Call the example above $T(z_0, z_1, z_2)$, where we think of z_0, z_1, z_2 as variables rather than as the fixed numbers given above. Let $\theta_k(z_0, z_1, z_2)$ be the cone angle of $T(z_0, z_1, z_2)$ at vertex k. Let

$$F(z_0, z_1, z_2) = (\theta_0, \theta_1, \theta_2).$$

If $F(z_0, z_1, z_2) = (2\pi, 2\pi, 2\pi)$ and $T(z_0, z_1, z_2)$ is embedded then we have our proof. We use the Inverse Function Theorem. Let dF be the Jacobian of F, computed at our example. We display dF and its inverse up to 2 digits:

$$\begin{bmatrix} -0.91\ldots & +0.74\ldots & +0.39\ldots \\ +0.74\ldots & -1.92\ldots & +1.14\ldots \\ +0.39\ldots & +1.14\ldots & -0.06\ldots \end{bmatrix} \qquad \begin{bmatrix} -0.56\ldots & +0.23\ldots & +0.76\ldots \\ +0.23\ldots & -0.04\ldots & +0.64\ldots \\ +0.76\ldots & +0.64\ldots & +0.57\ldots \end{bmatrix} \tag{3}$$

Notice that dF is very well conditioned because dF^{-1} is small. We will need just two digits of precision, and some crude bounds on the variation of these matrices, to promote our example to a genuinely flat embedded example *via* an effective Inverse Function Theorem. (Of course, I computed these matrices to about 400 decimal places.) This method is a tried-and-true one. For instance, it is similar in spirit (thought not in the details) to the method in Matthew Ellison's 2023 preprint [**E**].

Here is a quick corollary of our method.

Corollary 1.4 *There exists a* 6*-dimensional open manifold of non-similar* 8*-vertex paper tori.*

Proof: We can normalize our example so that the line of symmetry is the Z-axis. We can then perturb our points symmetrically by changing the X and Y coordinates but not the Z-coordinates. As long as the perturbation is small enough, the Inverse Function Theorem selects Z-coordinates for us which give a paper torus. We have 8 free parameters to play with and then these are grouped into a 2-dimensional foliation under the action of similarities of the XY-plane which fix the origin. A cross-section of this foliation gives us our 6-dimensional open manifold of inequivalent 8-vertex paper tori. ♠

1.4 Paper Organization

- In §2, we will prove Theorem 1.1.
- In §3, we will scale our example from Equation 1 by a factor of 10^{32} to get an integer example Ω. We prove that Ω is 10^{28}-*robustly embedded*, meaning that if we change any of z_0, z_1, z_2 by less than 10^{28} then the result is still embedded.
- In §4, we prove that Ω is 10^{-30}-*flat*, meaning that the deviation of each cone angle from 2π is less than 10^{-30}
- In §5, we prove an effective version of the Inverse Function Theorem.
- In §6, we use the results from §2-5 to reduce Theorem 1.3 to a final estimate on the derivatives of the Jacobian matrix dF discussed above.
- In §7, we establish these estimates and thereby finish the proof of Theorem 1.3.
- In §8, an appendix, I show how to promote the proof of the Hull Lemma to the proof of the Hull Theorem. Again, this is not needed for the proofs of our results.

1.5 Computer Assistance

My proof is heavily computer assisted, though all the important calculations are done with exact integer arithmetic. The reader can download the code for Theorems 1.1 and 1.3 respectively at
www.math.brown.edu/∼res/Java/SevenTorus.tar.
www.math.brown.edu/∼res/Java/EightTorus.tar.

I used ChatGPT-4o in various ways to help me with this project. Since this is my first time doing a project with the assistance of ChatGPT, I think that this warrants some discussion. Here are the ways I used the program.

- I used it to discuss high level mathematical ideas. The act of explaining math to a tireless and enthusiastic listener was quite fun.
- I used it to find typos and other glitches in the paper.
- I used it to format data. If the computer printed out some numerical data, I would paste it into ChatGPT and have it format the data in a way that was usable either in the paper or in my computer programs.

- I used it to evaluate my code checking it, at least in a semantic way, for errors. This mostly served as a sanity check for my programming.

One final thing I would like to say is that I did not allow ChatGPT to compose prose for me or even to give me writing advice. I want to retain 100 percent of my own voice when I write something. ChatGPT was perfectly happy to stick to fixing typos and inconsistencies in the writing.

Overall, I found ChatGPT to be an enthusiastic and friendly companion as I worked.

1.6 Acknowledgements

I thank Samuel Lelievre and Alba Malaga-Sabogal for telling me all about flat tori (over a period of some years) and in particular telling me about the flatness problem and supplying me with some historical context. The main thing that inspired me to work on this project was the infectious enthusiasm of Samuel and Alba. I also thank Peter Doyle and Fabian Lander for helpful comments about this material. I thank Frank Lutz for making his manifolds webpage, which lets you download the combinatorial data for triangulations of surfaces and other manifolds. This page was very useful to me.

I thank the IHES, the Hamilton Institute, and the Isaac Newton Institute, where I started working on this paper. This work is also supported by a Simons Sabbatical Fellowship and the Mazur Chair at IHES.

2 Proof of Theorem 1.1

2.1 Crofton's Formula

Crofton's Formula is a classic result from integral geometry. See [**S**]. It applies to any rectifiable arc on the unit sphere S^2. We just need this result for finite unions of arcs of great circles, which we call *spherical polygons*.

The space of oriented great circles in S^2 is canonically bijective with S^2 itself, and inherits a canonical probability measure σ. Given a spherical polygon γ and a great circle C, we let $\#(C \cap \gamma)$ denote the number of intersection points. (We can ignore the finitely many great circles for which this is infinite.) Crofton's formula is as follows.

$$\text{length}(\gamma) = \pi \int_{\mathcal{S}} \#(C \cap \gamma)\ d\sigma. \tag{4}$$

At least for spherical polygons, this has a swift proof. Both sides of Equation 4 are additive, so it suffices to prove the equation for a spherical polygon that is just an arc of a single great circle. By continuity, it suffices to prove the result for great circle arcs whose length is a rational multiple of 2π. By additivity once again, it suffices to prove Crofton's formula for the great circle itself. But then the formula is obvious.

2.2 Proof of the Main Theorem

Here is a well-known consequence of Crofton's formula.

Lemma 2.1 *Suppose γ is a spherical polygon which is also a topological loop. If γ has length at most 2π then γ lies in some hemisphere of S^2.*

Proof: Crofton's formula says that some great circle C intersects γ at most once. But then γ lies in one of the two hemispheres defined by C. When γ has length exactly 2π we can shorten γ a bit by cutting a small corner off. The shortened curve then lies in a hemisphere. Taking a limit as the cut corner tends to 0 in length, we see that γ itself lies in a hemisphere. ♠

Now suppose that Ω is a flat embedded 7-vertex polyhedral torus. Let H be the convex hull of Ω. Call a vertex of Ω *interior* if it does not lie in ∂H. Given any vertex P of Ω let Ω_P denote the union of 6 triangles of Ω incident to P.

Lemma 2.2 *Ω has an interior vertex.*

Proof: Suppose that Ω does not have an interior vertex. By the Hull Lemma, Ω has a vertex P such that all 6 triangles of Ω incident to P lie in ∂H. Let Ω_P be the union of triangles incident to P as above.

A *support plane* for H is a plane which intersects the boundary ∂H and is disjoint from the interior of H. Since the cone angle of Ω at P is 2π, and since $\Omega_P \subset \partial H$, we see that H cannot be strictly convex at P. Hence some support plane through P intersects Ω_P in (at least) a line segment through P that contains P in its relative interior. But then the intersection of ∂H with a small sphere centered at P is a spherical polygon consisting of two arcs which connect the same pair of antipodal points.

Since the total length of γ is 2π, and since a great semicircle is the unique path on S^2 of length at most π connecting two antipodal points, we see that γ must be a union of two great semicircles. Hence Ω_P is contained in the union of two planes whose intersection contains P.

One of these two planes, Π, contains at least 3 consecutive triangles of Ω_P. But then Π contains at least 5 vertices of Ω. Since the 1-skeleton of the triangulation of Ω is the complete graph K_7, we see that Π contains an embedded copy of K_5. This contradicts the fact that K_5 is not planar. ♠

Now we know that our flat embedded 7-vertex torus Ω has an interior vertex P.

Lemma 2.3 *The cone angle θ at P exceeds 2π.*

Proof: We translate so that P is the origin. Let S^2 be the unit sphere. Let θ_P denote the cone angle of Ω at P. Let $\ell(\cdot)$ stand for length.

Let Ω_P be the union of 6 triangles of Ω incident to P, as above. Let $\widehat{\Omega}_P$ denote the union of rays emanating from P whose initial portions are contained in Ω_P. Let $L_P = \widehat{\Omega}_P \cap S^2$. We have $\theta_P = \ell(L_P)$.

Since P is in the interior of the convex hull of Ω and since $\widehat{\Omega}_P$ contains all the vertices of this convex hull, we see that $\widehat{\Omega}_P$ cannot lie in any halfspace bounded by a plane through the origin. This means that L_P cannot lie in any hemisphere of S^2. By Lemma 2.1, we have $\ell(L_P) > 2\pi$. Hence $\theta_P > 2\pi$. ♠

This completes the proof of the Main Theorem modulo the Hull Lemma.

2.3 Proof of the Hull Lemma

We suppose that Ω is an embedded 7-vertex polyhedral torus having 7 points on its convex hull H. We can perturb so that the points are in general position. This makes H into a triangulated solid polyhedron with 7 vertices, 15 edges, and 10 faces. One property we will use repeatedly is that Ω is *neighborly*: Every two vertices of Ω are joined by an edge in the triangulation. In particular, every edge of H is an edge of the triangulation of Ω.

We call the 15 edges in ∂H the *external edges*. We call the remaining 6 edges *internal edges*. We say that an *internal edge pattern* is a choice of 6 distinguished edges from the 1-skeleton of the triangulation, normalized (by symmetry) so that the first internal edge is (01). There are $\binom{20}{5} = 15504$ different internal edge patterns.

We say that an *external triangle* of Ω is one that lies in ∂H, and an *internal triangle* is one that does not. Each internal edge is incident to two internal triangles.

Lemma 2.4 *An internal triangle cannot be bounded by* 3 *external edges.*

Proof: Suppose that such a triangle exists. Call it τ. If all three edges of τ lie in ∂H and $\tau \not\subset \partial H$ then τ separates H into two components, both of which contain vertices of Ω in their interior. Any path in H connecting two such vertices must intersect τ. On the other hand $\Omega - \tau$ is path connected. This is a contradiction. ♠

By Lemma 2.4, the internal edge pattern determines the set of internal triangles and the set of external triangles. It is worth pointing out a triangle of ∂H is not necessarily a triangle of Ω. However, its boundary is one of the 35 three cycles contained in the 1-skeleton of the triangulation of Ω.

Since Ω is neighborly, the degree of a vertex in the triangulation of ∂H is just the number of external edges incident to that vertex, and this is the same as 6 minus the number of internal edges incident to the vertex. In particular, since H is a convex polyhedron and ∂H is triangulated, the degree of each vertex in the triangulation of ∂H is at least 3. In other words, at most 3 internal edges are incident to the same vertex of Ω. We eliminate all internal edge patterns having fewer than 3 internal edges incident to some vertex.

Now we consider the remaining internal edge patterns. For each vertex (q) we have the list $\{(q, v_i) \mid i = 1, ..., K_q\}$ of K_q external edges incident to

q. We write $K = K_q$ and order these $K = K_q$ vertices cyclically according to the link of (q) in Ω. The link of L_q of (q) in ∂H is some permutation $(w_1, ..., w_K)$ of $(v_1, ..., v_K)$.

Lemma 2.5 (Cycle Rule) *In order to be a viable candidate for the link of (q) in ∂H, the link $(w_1, ..., w_K)$ must satisfy two properties.*

1. *(w_i, w_{i+1}) must be an external edge for all i. Indices are taken cyclically*
2. *The cycle $(w_1, ..., w_K)$ must be a dihedral permutation of $(v_1, ..., v_K)$.*

Proof: The necessity of Condition 1 is obvious. Condition 2 requires some explanation. Let P be vertex (q). Consider the picture in ∂H at P. Since the points are in general position, ∂H is a proper convex cone near P. We let Π be a plane parallel to a support plane through (q) that just cuts off a small corner of H near (q). Consider the intersection $\partial H \cap \Pi$. This is a convex K-gon Δ, and the cyclic order of Δ is given by $(w_1, ..., w_K)$.

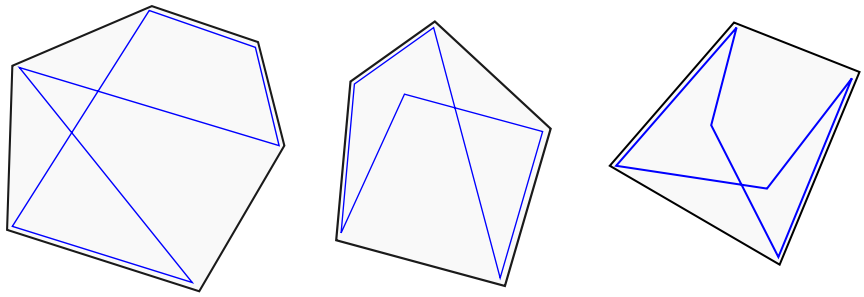

Figure 2.1: Δ (black) and γ (blue)

At the same time, and using the notation from the proof of Lemma 2.3, let $\gamma = \widehat{\Omega}_P \cap \Pi$. The polygonal loop γ is contained in the region bounded by Δ and visits the vertices of Δ in the order $(v_1, ..., v_K)$. Figure 2.1 shows some examples, with γ slightly moved off Δ to make the drawing more clear. Note that if $K < 6$ then there will be some extra vertices to γ as well; this does not matter in the argument. If the permutation is not dihedral then γ cannot be an embedded loop. But then Ω_P is not embedded, a contradiction. ♠

Using the Cycle Rule (and our computer code) we eliminate all remaining internal edge patterns except 6. These 6 are all the same up to combinatorial isomorphism. Here is one of them.

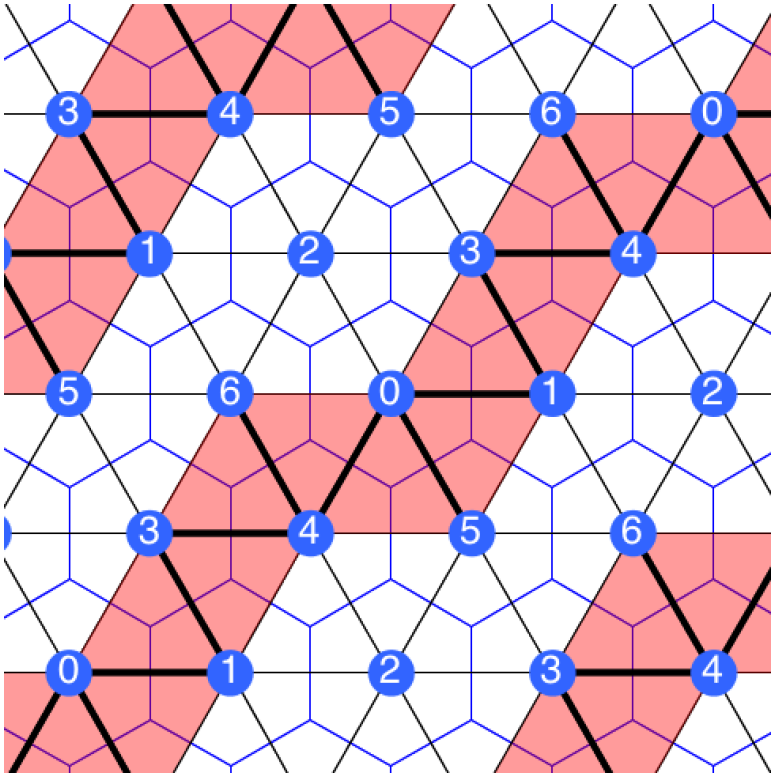

Figure 2.2: The one remaining pattern

Evidently, all 6 triangles incident to vertex (2) are external triangles. This completes the proof of the Hull Lemma.

3 Robust Embedding

3.1 The Example

Here is the list of faces in the triangulation shown in Figure 1.2.

$$
\begin{array}{lll}
0 & 1 & 2 \\
0 & 1 & 3 \\
0 & 2 & 4 \\
0 & 3 & 5 \\
0 & 4 & 6 \\
0 & 5 & 6 \\
1 & 2 & 5 \\
1 & 3 & 4 \\
1 & 4 & 7 \\
1 & 5 & 7 \\
2 & 4 & 7 \\
2 & 5 & 6 \\
2 & 6 & 7 \\
3 & 4 & 6 \\
3 & 5 & 7 \\
3 & 6 & 7
\end{array}
$$

Let Ω be the example from Equation 1 scaled up by 10^{32}. Here are the coordinates, with row i describing vertex i for $i = 0, ..., 7$.

$$
\begin{bmatrix}
+7550\,\texttt{(28 zeros)} & +6500\,\texttt{(28 zeros)} & +9805\,0571\,5859\,7793\,5561\,6538\,2008\,5693 \\
-4550\,\texttt{(28 zeros)} & +3450\,\texttt{(28 zeros)} & +9902\,8162\,4334\,3054\,2934\,3176\,1585\,8328 \\
-1700\,\texttt{(28 zeros)} & +11400\,\texttt{(28 zeros)} & +9765\,3883\,4703\,1231\,7624\,1842\,4567\,2434 \\
+4550\,\texttt{(28 zeros)} & -3450\,\texttt{(28 zeros)} & +9902\,8162\,4334\,3054\,2934\,3176\,1585\,8328 \\
-7550\,\texttt{(28 zeros)} & -6500\,\texttt{(28 zeros)} & +9805\,0571\,5859\,7793\,5561\,6538\,2008\,5693 \\
-0900\,\texttt{(28 zeros)} & +6650\,\texttt{(28 zeros)} & +0000\,0000\,0000\,0000\,0000\,0000\,0000\,0000 \\
+1700\,\texttt{(28 zeros)} & -11400\,\texttt{(28 zeros)} & +9765\,3883\,4703\,1231\,7624\,1842\,4567\,2434 \\
+0900\,\texttt{(28 zeros)} & -6650\,\texttt{(28 zeros)} & +0000\,0000\,0000\,0000\,0000\,0000\,0000\,0000
\end{bmatrix}
$$

We let $\Omega(z_0, z_1, z_2)$ be the example obtained from Ω by replacing the first 3 entries in the third column by z_0, z_1, z_2 and then making the symmetric replacements for the entries z_4, z_3, z_6. (The symmetries $z_0 = z_4$ and $z_1 = z_3$ and $z_2 = z_6$ are discussed in the paragraph after Equation 1.) So, in our example, Ω is obtained by setting these variables to the respective 32-digit integers we have listed.

There is one point we need to be clear on. We do not get $\Omega(z_0, z_1, z_2)$ by scaling up what we called $T(z_0, z_1, z_2)$ in the introduction. We get it by scaling up $T(sz_0, sz_1, sz_2)$ where $s = 10^{-32}$.

3.2 Main Result

Let z_0, z_1, z_2 be as they are for Ω. Given any $\Omega' = \Omega(z_0', z_1', z_2')$ we let

$$|\Omega' - \Omega| = \max(|z_0' - z_0'|, |z_1' - z_1|, |z_2' - z_2|). \tag{5}$$

We do not insist that z_0', z_1', z_2' are integers.

Definition: We say that Ω is *K-robustly embedded* if and only if Ω' is embedded whenever $|\Omega' - \Omega| \le K$.

Here is our main result.

Lemma 3.1 (Robust Embedding) *Ω is 10^{28}-robustly embedded.*

Proof: We have $120 = 24 + 72 + 24$ triangle pairs to consider. The first 24 are pairs having 0 vertices in common. The next 72 are pairs having 1 vertex in common. The last 24 are pairs having 2 vertices in common. Given one of these pairs (T_0, T_1) we let $k(T_0, T_1)$ be the number of vertices the pair have in common.

We let $T_0^\sharp$ denote the convex hull of the vertices of T_0 that do not belong to T_1. Thus $T_0^\sharp$ is respectively a triangle, a segment, and a point when $k(T_0, T_1) = 0, 1, 2$. We define $T_1^\sharp$ in the same way.

Given a vector L, let $|L|$ denote the max of the absolute values of its coordinates. In practice we work with 16-digit integer vectors, giving us $|L| < 10^{17}$. We say that (T_0, T_1) is *λ-separated* if there exists a vector L with $|L| < 10^{17}$ such that one of the following two equations holds

$$\max L \cdot T_0^\sharp \le \min L \cdot T_1^\sharp - \lambda, \qquad \max L \cdot T_1^\sharp \le \min L \cdot T_0^\sharp - \lambda. \tag{6}$$

To be clear, we are taking the maximum and minimum dot product with L over the triangles. That is,

$$\max L \cdot T_0^\sharp = \max_{p \in T_0^\sharp} L \cdot p,$$

and likewise for the other terms. By convexity, it suffices to take this maximum over the vertices. This makes for finite computations.

If $|\Omega' - \Omega| \le K$ then no vertex of T_0 or T_1 moves by more than K, and at most one of the coordinates changes. Hence, each dot product above changes

by at most $10^{17}K$. Hence, if our original pair is λ separated then the new pair (T_0', T_1') is still λ' separated, provided that

$$\lambda' := \lambda - (2 \times 10^{17})K > 0.$$

We show by direct calculation that all pairs are λ-separated, where

$$\lambda = 5112923953186065818296799992123421121146292290. \tag{7}$$

This amounts to producing 120 specific vectors $L_0, ..., L_{119}$ and then verifying one of the two sides of Equation 6 when L_i is tested against the (i)th pair. In the next section we will explain how we produced these vectors, even though logically speaking it does not matter how we produced them.

Our separation calculation shows that Ω is K-robustly embedded as long as

$$K < \frac{\lambda}{2 \times 10^{17}} = 2.55... \times 10^{28}.$$

Therefore Ω is 10^{28}-robustly embedded. ♠

Remark: The Robust Embedding Lemma conforms pretty well to our pictures in Figures 1.4 – 1.6. Abusing our terminology slightly, we would say that our pictures indicate that the example T from Equation 1 is a better than 10^{-4}-robustly embedded. Here we are scaling up by 10^{32}.

3.3 Producing the Vectors

Let (T_1, T_2) be one of the 120 triangle pairs. Our method depends on $k(T_1, T_2)$. We do one thing when $k = 2$ and something else when $k = 0, 1$. In all cases, the output of our construction is a vector L.

Our method is easier when $k = 2$. We describe that first. For $j = 0, 1$ let n_j be one of the two unit vectors perpendicular to the plane containing T_j. It turns out that n_0 and n_1 are never multiples of each other. There is usually a descrepancy of at least 10^{-3}. We then have the two vectors

$$V_\pm = \frac{n_0 \pm n_1}{\|n_0 \pm n_1\|}. \tag{8}$$

We then set

$$L = \text{floor}(10^{32} V_\pm). \tag{9}$$

We leave it to the reader to figure out why this would (very, very probably) work. The rounding is the only step which introduces the miniscule chance of the method not working. As we said above, our formal proof does not require a justification for why the vectors work. They just have to work.

Remark: The fact that $n_1 \approx n_2$ in some cases suggests that perhaps there is a more symmetric example fairly near to ours. In this particular construction, we are benefitting from the slight asymmetry in our example. We would need to treat the case $n_1 = \pm n_2$ specially because one of the two vectors in Equation 8 is not defined.

Now we turn to the cases when $k(T_0, T_1) < 2$. These cases are more fun. We define

$$\phi(x) = 2x' - 1, \qquad x' = x - \text{floor}(x). \tag{10}$$

The range of ϕ is $[-1, 1]$. Next, we define

$$V_n = \frac{W_n}{\|W_n\|}, \qquad W_n = (n\sqrt{2}, n\sqrt{3}, n\sqrt{5}). \tag{11}$$

The countable collection $\{V_n\}$ is dense in the unit sphere. This is nice because we are looking for a good direction to use in order to separate T_0 from T_1, and the list $V_0, V_1, V_2, ...$ eventually samples all directions up to an arbitrarily small error. Another nice feature of these pseudo-random vectors is that they are easy to name explicitly.

For each pair (T_0, T_1) we try to verify a numerical version of Equation 6 for our original example in Equation 1. We take $\lambda = .005$ when $k(T_0, T_1) = 0$ and $\lambda = .01$ when $k(T_0, T_1) = 1$. We found these "tolerances" by trial and error. We wanted a healthy separating vector but we did not want to wait too long (or forever) for our search to terminate.

For each pair, the search terminated with success. We then chose

$$L = \text{floor}(10^{32} V_n), \tag{12}$$

where V_n is the first vector on our list that satisfies our Equation 6 with the given tolerance.

For the reader who is interested in reproducing our vectors without recoding everything, we list out the indices of the vectors which worked for us. Here is what we got for the 24 pairs having no vertices in common.

1	4	4	10515	32202	32202
1409	8	7	10	32	32
1	1	326	15725	24	170004
1	1	183100	183100	24	1

Here is what we got for the 72 pairs having one vertex in common.

1	32202	32202	32202	1409	1409
1409	1	7910	8	10515	1
32202	10515	32202	7910	7910	32202
8	10515	1	1	1	32202
32202	1409	1	7910	1	1
8	32202	7910	1	1	10515
32202	32202	32202	10515	1	8
8	1	32202	10515	32202	32202
8	32202	7910	8	1	10515
8	32202	8	10515	32202	32202
32202	1	1	1	8	10515
10515	32202	32202	7910	32202	7910

These lists are meant to be read across and then down, like a page in a book.

Remark: The repetition of some of the numbers on these lists is another suggestion that perhaps there is a more symmetric example than ours that is fairly nearby. In my opinion, the structure of these lists somewhat reflects the geometry depicted in Figures 1.4 –1.6.

4 Near Flatness

4.1 The Main Result

Let Ω be our integer flat torus from §3.1. Let θ_i denote the cone angle of Ω at the (i)th vertex. We say that Ω is ϵ-*flat* if $\max_i |\theta_i - 2\pi| < \epsilon$. The maximum is taken over all 8 cone angles. We have the exact coordinates of Ω and we also have a high-precision version of the square root function and the arccos function. Using 100 digits of accuracy in a direct angle calculation, we check that Ω is (7.26×10^{-32})-flat. For readers who are not satisfied with such a calculation, we prove, just using rational arithmetic, the following slightly weaker bound.

Lemma 4.1 (Near Flat) Ω *is* 10^{-30}-*flat.*

4.2 Discussion

Here we discuss the idea behind the proof. Look again at Figure 1.3. This shows a plot of the intrinsic flat structure on Ω up to an error of about 10^{-16}. Around each vertex V_i we have 6 triangles $T_{i0}, ..., T_{i5}$. These triangles do not necessarily fit perfectly in the plane, but the error is so small that you cannot tell from the naked eye. We compute these triangles numerically using dot products, the square-root function, and the arccos function. We do not care about the rigor of our calculation. Call the union of these 6 triangles a *flower*. The (i)th flower corresponds to the (i)th star in Ω, though only approximately.

For our proof, we re-do the calculation of our 8 flowers using high precision arithmetic and scaling, as follows:

1. We get 100 digit precision coordinates for the triangles.
2. We scale up the vertices by 10^{32}.
3. We translate so that the central vertex is the origin.
4. We round all coordinates down to integers by taking floors.

This gives us integer models which approximate the geometry of the 8 stars in Ω.

Now something very nice happens. For our scaled-up models, these big integer triangles fit together perfectly around a vertex. Put another way, for each flower F there are 6 integer vectors $V_0, ..., V_5$ such that each of the 6 triangles in F is defined by 2 cyclically consecutive vectors on the list. We then compute the mismatch between the various dot products associated to Ω and the corresponding dot products associated to the integer flowers, and this gives us the Near Flat Lemma.

We will include all our flowers at the end of the chapter. For the purposes of our proof, our method for producing the integer flowers is irrelevant. Logically, all that matters is that we simply have these 8 lists of vectors which we can compare to the picture in Ω. Rather than compare the angles, we compare the dot products which go into the calculation of the angles. We finish off the argument by using the Lipschitz properties of the function that converts between the relevant dot product ratio and the angle.

4.3 Rational Arithmetic

We perform integer arithmetic calculations using the BigInteger class. These are done exactly. Out of a slight laziness, and a desire for a transparent argument over a tedious exercise that anyone can perform, we will avail ourselves of the BigDecimal class in Java, which is perfectly capable of computing the ratios of 64 digit integers up to 100 digits of precision. Let us at least explain how we might have implemented a pure integer verification.

Suppose we want to establish a bound like

$$\left|\frac{a}{b} - \frac{a^*}{b^*}\right| < \epsilon. \tag{13}$$

where a, b, a^*, b^* are positive integers and $\epsilon > 0$. We would choose the smallest integer N such that $1/N < \epsilon$ and then try to check that

$$N|ab^* - ba^*| < bb^*. \tag{14}$$

In practice we would compute ϵ using the BigDecimal class, then round in the correct direction to get N, and then use Equation 14. Logically, it does not matter how we find N. All that matters is the existence of N. We did not do this, but if some readers of this article protest, we will add it to our code.

4.4 The Proof

Let τ_{ab} be the (b)th triangle in the (a)th star of Ω and let τ^*_{ab} be the (b)th triangle in the (a)th integer flower. Let V_{ab} and W_{ab} be the two vectors that τ_{ab} defines. By this we mean that V_{ab} points from vertex a to the first other vertex of τ_{ab} and W_{ab} points from vertex a to the second other vertex of τ_{ab}. (The torus is oriented, so we have no trouble defining "first" and "second" here.) We define V^*_{ab} and W^*_{ab} with respect to τ^*_{ab} in a similar manner. Let θ_{ab} and θ^*_{ab} respectively denote the angles of τ_{ab} and τ^*_{ab} at vertex (a).

When $r > 0$ we have

$$\theta = \phi(r), \qquad \phi(\cdot) = \arccos(\sqrt{(\cdot)}), \qquad r = \frac{(V \cdot W)(V \cdot W)}{(V \cdot V)(W \cdot W)}, \tag{15}$$

We take the positive branch of the square root function here.

We have $48 = 8 \times 6$ pairs of triangles to check. We get the following bounds which work uniformly for all these pairs:

1. $r, r^* \in I := [.00112, .77714]$ for all a, b.
2. $|r - r^*| < 1.07 \times 10^{-32}$.
3. $V \cdot W > 0$ and $V^* \cdot W^* > 0$.

Lemma 4.2 *ϕ is 15-lipschitz on I.*

Proof: We compute that

$$\left|\frac{d\phi}{dt}\right| = \frac{1}{2\sqrt{t(1-t)}}.$$

Establishing our bound is the same as establishing that

$$225 - \frac{1}{4t(1-t)} > 0, \qquad \forall t \in I$$

Establishing this bound is a routine exercise in algebra (or calculus) which we omit. ♠

Combining Lemma 4.2 with our calculations, and remembering that our flowers have 6 triangles in them, we see that Ω is ϵ-flat when

$$\epsilon = 6 \times 15 \times 1.07 \times 10^{-32} < 10^{-30}.$$

This completes the proof of the Near Flat Lemma.

4.5 The Integer Flowers

$$\begin{bmatrix} 124788644050109876602825361295674 & 0 & 0 \\ 77712297261878079581593486900079 & 701299238253016147793542532242 91 & 0 \\ -704359812948700118983113598452 07 & 1863860846174647282699241636445 24 & 0 \\ -148148278556748091479904846745286 & 116256160792163113490569910420232 & 0 \\ -980958396509123484174281864597 04 & -844607650422341154582626638280 23 & 0 \\ 534159637249392321001669501547 53 & -891512226680094342232541236906 15 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 844653102916722215274504411882 66 & 0 & 0 \\ 695503688111756229536572813416 42 & 1036096128845893925770241282368 71 & 0 \\ -342413162807242592800299817223 40 & 1089473829853815615710260489796 21 & 0 \\ -103791685091899882233687263063983 & 533777010079216899400192074 2749 & 0 \\ -883387938606681727027040049538 84 & -123183133805664642986337651698562 & 0 \\ 440458018509797838125352157646 31 & -101107834979521394561120290119996 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 104677635440723967179988723412878 & 0 & 0 \\ 120349982575376878832149699114 30 & 836035134405817547813215620534 77 & 0 \\ -943170811720764779320173480084 92 & 544134097076563844184994054215 07 & 0 \\ -191630826036432582723537220778797 & -128131286236400952463466737043786 & 0 \\ -973137448643561047915198727703 05 & -182544695944057336881966142465293 & 0 \\ 320976620011992189573675149386 00 & -185561707945141208999189687748482 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 151584388889297911866446233050267 & 0 & 0 \\ 56495441368939829023723230650622 & 947158490405411844361828196410 71 & 0 \\ -492238263379653601955027464283 17 & 686396646503921741128435861460 43 & 0 \\ -124728889560435721214027036255644 & -386132009624326109066287288214 9 & 0 \\ -685799082068823463711716965509 92 & -913170093155464186068804711820 19 & 0 \\ 561489813535533748426553397046 51 & -874556892193031575162175982998 71 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 199251097863976649156276288186992 & 0 & 0 \\ 161120818626134715906111075736407 & 974858408524300042447869505525 76 & 0 \\ -443301676525334001645971594447 53 & 121619286399136814961872050344 16 & 0 \\ -102277745880424119662590582022408 & -184517853197030808754346791671 89 & 0 \\ -441132354719880207515695426909 15 & -116731435954750256516337414942322 & 0 \\ 381302792378419332501652124505 84 & -974858408524300042447869505525 77 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 20686355637674156862294612970 6945 & 0 & 0 \\ 364748131594131286307826367 8069 & 108826590670126188269802629065 616 & 0 \\ -757891723825758422096711878581 45 & 801175904811553070701152642600 41 & 0 \\ -116999860561305386939755847152992 & -657573769902289845475912516041 89 & 0 \\ -412106881787295447300846592948 48 & -145874967471384291617706515864231 & 0 \\ 582160014203865741959132384856 30 & -115617091154199741307099293537955 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 188317304997339823174158679929020 & 0 & 0 \\ 17841734520938330693404444524334 & 1031459154341337348432797053188 12 & 0 \\ -803288653594906850975011357296 66 & 261086581946542544134026783013 33 & 0 \\ -696930270590789451509531873464 48 & -836624942024550643631503187935 64 & 0 \\ 935939415252107683278414713888 03 & -210666025048595406864753841069280 & 0 \\ 163286968584289713478794658735251 & -127003530846140342501603522275717 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 151584388889297911866446233050267 & 0 & 0 \\ 954354075357445370237908933456 15 & 87455689219303157516217598299870 & 0 \\ -562392313604980578852401519560362 & 199072046789130364493850160222280 & 0 \\ -105719267706905189219225977078940 & -260761843901490102323339233495028 & 0 \\ -564954413689398290237232306506 23 & -947158490405411844361828196410 72 & 0 \\ 950889475203580828427230023996 44 & -947158490405411844361828196410 72 & 0 \end{bmatrix}$$

5 Effective Surjectivity

Here we prove a result which will help us use the Inverse Function Theorem to promote our nearly flat embedded example Ω to an exactly flat embedded example.

Let $G : \boldsymbol{R}^3 \to \boldsymbol{R}^3$ be a smooth map. Let $dG|_p$ denote the differential of G at p. We call G *cleanly expansive* on a ball B if

1. $dG|_p$ is an expanding map at each $p \in B$.

2. For any vector $V \in \boldsymbol{R}^3$ and any $p, q \in B$ we have $dG|_p(V) \cdot dG|_q(V) > 0$.

Here we are interpreting V as being simultaneously in the tangent space at p and at q.

Lemma 5.1 (Effective Surjectivity) *Let B be the unit ball centered at $(0,0,0)$ and suppose that $G(0,0,0) = (0,0,0)$. Suppose that G is cleanly expansive on the unit ball B. Then $B \subset G(B)$.*

Proof: Let $G : \boldsymbol{R}^3 \to \boldsymbol{R}^3$ be a map which is cleanly expansive on the unit ball B. We note first of all that G is locally invertible in an open neighborhood of B. This follows from the Inverse Function Theorem and continuity. The point here is that dG is nonsingular not just on B but in a neighborhood of B.

We claim that G is injective. Suppose not. Suppose that $p_0 \neq p_1 \in B$ are such that $G(p_0) = G(p_1)$. Then we can find a line segment in B which G maps to an arc γ that starts and ends at the same point. This is impossible, because all the tangent vectors of γ (taken as γ moves in one direction) have positive dot product with the initial one $\gamma'(0)$. This contradiction shows that G is injective.

Now suppose, for the sake of contradiction, that $B \not\subset G(B)$. Let $\delta < 1$ be the infimum of radii r such that $B_r(0) \subset G(B)$. By the Inverse Function Theorem, we have $\delta > 0$. By compactness, there exists a point $p \in \partial B$ such that $\|G(p)\| = \delta$. Let $q = G(p)$.

Consider the line segment β which connects q to $(0,0,0)$, oriented so that β starts at q and ends at $(0,0,0)$. By the Inverse Function Theorem, we can find small open neighborhoods U_p, U_q about p and q respectively so that $G : U_p \to U_q$ is a diffeomorphism. The intersection $U_q' = U_q \cap G(B)$ locally lies on one side of $G(\partial B)$. Since $G(B)$ contains all of $B_\delta(0)$, the set U_q'

contains all points of $U_q \cap \beta$ sufficiently close to q. Thus, if we choose some initial arc $\beta^* \subset \beta$ contained entirely in U_q' we have an inverse image $\alpha^* \subset B$ that starts at p and immediately moves into the interior of B.

We want to extend the length of β^* so that its inverse image α^* is still defined. Since G is a local homeomorphism at all points of B, this extension is possible as long as α^* remains in B. Could it happen that α^* ever exits B? Well, as we trace along β we come ever closer to the origin. So, if α^* ever exits B we would have produced a new point $p^* \in \partial B$ such that $\|G(p^*)\| < \delta$. This is impossible. Hence α^* cannot exit B. But now we can continue the process all the way, finding an arc $\alpha \subset B$ such that $G(\alpha) = \beta$.

One endpoint of α is p. Consider the other endpoint p' of α. We have $G(p') = (0, 0, 0)$. Since G is injective on B this forces $p' = (0, 0, 0)$. This means that α connects a point on the unit sphere ∂B to the origin. The length of α is 1 and the length of β is $\delta < 1$. But then G would decrease the arc length of α, contradicting the fact that G is cleanly expansive on B. ♠

6 Healthy Expansion

6.1 The Proof Modulo the Expansion Lemma

In this section, we reduce the proof of Theorem 1.3 to an auxiliary lemma we call the Expansion Lemma. After this, we reduce the Expansion Lemma to one final bound on the derivatives of the matrix dF. We call this final bound the Crude Bound Lemma.

Let us recall the set-up from the previous chapter. Let $G : \boldsymbol{R}^3 \to \boldsymbol{R}^3$ be a smooth map. Let $dG|_p$ denote the differential of G at p. We call G *cleanly expansive* on a ball B if

1. $dG|_p$ is an expanding map at each $p \in B$.
2. For any vector $V \in \boldsymbol{R}^3$ and any $p, q \in B$ we have $dG|_p(V) \cdot dG|_q(V) \geq 0$.

For the purposes of proving Theorem 1.3, we generalize the definition slightly. We say that G is λ-*cleanly expansive* on a ball B if the scaled map λG is cleanly expansive on B. Let $B_r(v)$ denote the ball of radius r centered at a point v.

Corollary 6.1 *Let $p, q \in \boldsymbol{R}^3$ be such that $G(p) = q$. Suppose that G is λ-cleanly expansive on $B_r(p)$. Then $B_s(q) \subset G(B_r(p))$ when $s = r/\lambda$.*

Proof: This follows from the Effective Surjectivity Lemma by translation and scaling. ♠

Remark: Let us make sure that we have scaled things the right way. If λ is large then we have to scale G up a lot for it to cover a reasonable-sized ball. This means that it is harder for the image of G to cover balls of a given size. Our scaling law reflects this (rather than the opposite) phenomenon.

Now we turn our attention back to the proof of Theorem 1.3. Let Ω be our example from §3.1. In the notation established in §3.1, we define

$$\Phi(z_0, z_1, z_2) = 10^{32}(\theta_0, \theta_1, \theta_2). \tag{16}$$

Here $(\theta_0, \theta_1, \theta_2)$ are the cone angles at vertices $0, 1, 2$ respectively. Let

$$p = (z_0, z_1, z_2) \in \boldsymbol{R}^3$$

denote the point corresponding to our example Ω.

Lemma 6.2 (Expansion) *The map Φ is 100-cleanly expansive in the ball of radius 10^5 centered at p.*

Following this section we reduce the Expansion Lemma to one final bound, which we call the Crude Bound Lemma.

Let $q = \Phi(p) \in \boldsymbol{R}^3$. By Corollary 6.1, we have

$$B_{1000}(q) \subset \Phi(B_{10^5}(p)). \tag{17}$$

But Ω is 10^{-30}-flat. This means that $(2\pi, 2\pi, 2\pi)$ lies in the cube of side-length 2×10^{-30} centered at $(\theta_0, \theta_1, \theta_2)$. Basic geometrical properties of the cube now imply that

$$\|(\theta_0, \theta_1, \theta_2) - (2\pi, 2\pi, 2\pi)\| < 10 \times 10^{-30} = 10^{-29}.$$

Scaling up, we see that

$$\|q - 10^{32}(2\pi, 2\pi, 2\pi)\| < 10^3, \qquad q = \Phi(p).$$

But now Equation 17 gives

$$10^{32}(2\pi, 2\pi, 2\pi) \subset B_{1000}(q) \subset \Phi(B_{10^5}(p)).$$

This tells us that there is some point $p^* \in B_{10^5}(p)$ such that, with the obvious notation, $(\theta_0^*, \theta_1^*, \theta_2^*) = (2\pi, 2\pi, 2\pi)$. But then the torus Ω^* corresponding to p^* is flat. Since Ω is 10^{28}-robustly embedded, we see that Ω^* is also embedded.

This reduces the proof of Theorem 1.3 to the Expansion Lemma.

6.2 Cleaning off the Problem

It remains only to prove the Expansion Lemma. So far we have been working with the scaled up torus Ω defined in §3.1, but at this point in the proof Ω has served its purpose.

We would like to revert to our torus T defined in Equation 1 and we would prefer to use the map F defined in connection with T. Let us discuss how these objects compare. First, recall that p is the point corresponding to our example Ω. Let $S : \boldsymbol{R}^3 \to \boldsymbol{R}^3$ be the map which dilates by 10^{32} about the origin. The point

$$\overline{p} = S^{-1}(p)$$

corresponds to T.

By construction

$$F = S^{-1} \circ \Phi \circ S. \tag{18}$$

This means that the Expansion Lemma is true if and only if the following lemma is true.

Lemma 6.3 (Expansion II) *F is 100-cleanly expansive in $B_{10^{-27}}(\overline{p})$.*

Notice that we are working in an extremely small ball about $\overline{p}$. Heuristically, we expect dF to be almost constant in this ball. Let us first analyze the picture right at $\overline{p}$. We repeat the formula for $dF_{\overline{p}}$ given in Equation 3, a formula in which the first two digits of each entry are correct.

$$dF_{\overline{p}} = \begin{bmatrix} -0.91\ldots & +0.74\ldots & +0.39\ldots \\ +0.74\ldots & -1.92\ldots & +1.14\ldots \\ +0.39\ldots & +1.14\ldots & -0.06\ldots \end{bmatrix} \tag{19}$$

This is a routine calculation. Given the extremely low precision we need for our proof, we will not rigorously justify our 2-digit calculation of $dF_{\overline{p}}$.

6.3 Reduction to a Crude Bound

We now define the matrix

$$M = \begin{bmatrix} -0.91 & +0.74 & +0.39 \\ +0.74 & -1.92 & +1.14 \\ +0.39 & +1.14 & -0.06 \end{bmatrix} \tag{20}$$

Let $\mathcal{M}$ denote the set of matrices such that each entry is within $1/20$ of the corresponding entry of M. Thus, for instance, the upper left entry lies in $[-.93, -.89]$.

At this point, we let B denote the ball of radius 10^{-27} centered at $\overline{p}$. This is the final domain of interest to us. Below we will prove the following bound.

Lemma 6.4 (Crude Bound) *Throughout B, we have*

$$\left\| \frac{\partial F}{\partial z_i \partial z_j} \right\| < 10^{24}, \qquad i, j \in \{0, 1, 2\}.$$

Corollary 6.5 *$dF|_q \in \mathcal{M}$ for all $q \in B$.*

Proof: Let $m_{ij}(q)$ denote the (i,j)th matrix entry of $dF|_q$ Equation 19 says that $m_{00}(\overline{p}) \in [-.92, -.91]$, etc. Integrating the bound from the Crude Bound Lemma along a rectilinear path that joins $\overline{p}$ to q we see that

$$|m_{ij}(q) - m_{ij}(\overline{p})| < 3 \times 10^{-27} \times 10^{24} = \frac{3}{1000} < \frac{1}{20}.$$

This bound implies our lemma. ♠

Lemma 6.6 *For all points $q_1, q_2 \in B$ and all vectors V we have*

$$dF_{q_1}(V) \cdot dF_{q_2}(V) > 0.$$

Proof: Suppose this lemma is false. By continuity we can find a triple q_1, q_2, V such that V is a unit vector and

$$dF_{q_1}(V) \cdot dF_{q_2}(V) = 0.$$

Recalling that $\|\cdot\|_\infty$ denotes the maximum absolute value of a matrix entry, we write

$$dF|_{q_k} = M + E_k, \qquad \|E_k\|_\infty < \frac{1}{20}.$$

We have

$$(M(V) + E_1(V)) \cdot (M(V) + E_2(V)) = 0.$$

Rearranging this, we have

$$\|M(V)\|^2 = -(E_1(V) + E_2(V)) \cdot M(V) - E_1(V) \cdot E_2(V).$$

Now we observe that $\|M\|_\infty < 2$. Combining our bounds on the entries of E_k with the triangle inequality and the Cauchy-Schwarz inequality, we have

$$\|M(V)\|^2 \le \frac{1}{10}\|M(V)\| + \frac{1}{400} \le \frac{1}{10}\|M\|_\infty + \frac{1}{400} < \frac{1}{5} + \frac{1}{400} < \frac{1}{4}.$$

The vector $W = M(V)$ is such that $\|W\| \le 1/4$ and $\|M^{-1}(W)\| = 1$. But this is only possible if one of the eigenvalues of M^{-1} exceeds 4 in absolute value. This is absurd. Since (as one readily computes) all eigenvalues of M are greater than $1/2$ in absolute value, all eigenvalues of M^{-1} are less than 2 in absolute value. ♠

Now we study the matrices in $\mathcal{M}$.

Lemma 6.7 *For every matrix $M^* \in \mathcal{M}$, the differential dM^* expands distances by at least a factor of* $1/100$.

Proof: As in the previous proof, we write

$$M^* = M + E, \qquad \|E\|_\infty < \frac{1}{20}.$$

It suffices to prove that the eigenvalues of $M + E$ are all greater than $1/100$ in absolute value. We suppose this is false and derive a contradiction. So, there is some unit vector V such that

$$M(V) + E(V) = \lambda V, \qquad |\lambda| < \frac{1}{100}.$$

But then, by the triangle inequality,

$$\|M(V)\| \leq \frac{1}{100} + \frac{1}{20} < \frac{1}{4}.$$

Now we have the same contradiction as in the previous proof. ♠

Lemma 6.6 and Lemma 6.7 combine to prove the Expansion Lemma II, which in turn is equivalent to the Expansion Lemma. The only thing we have not proved is the Crude Bound Lemma. Therefore, we have reduced Theorem 1.3 to the Crude Bound Lemma. We prove the Crude Bound Lemma in the next chapter.

7 The Crude Bound Lemma

7.1 Discussion

To finish the proof of Theorem 1.3, we only need to prove the Crude Bound Lemma. Before getting to the proof, I want to discuss a certain robustness in our argument. We have arranged things so that we need a bound of 10^{24} on some partial derivatives. Had we worked with a 64 digit integer example built from a 10^{-60} flat object – something we certainly could have done – we would have left ourselves needing to get a bound of about 10^{48} on the same partial derivatives. Or we could have used a 128-digit example, and so on.

We mention this because we don't want to give the reader the impression that somehow Theorem 1.3 hinges on some fine points of our analysis in this chapter. Our 32-digit example simultaneously gives us something that does not require us to sweat too much over the final estimate and something which lets us have relatively short data files.

7.2 Breaking Down the Terms

In this section we will break down the partial derivatives we need to estimate into simpler expressions.

Recall that $F(z_0, z_1, z_2) = (\theta_0, \theta_1, \theta_2)$ is the map which computes how each of the angles $\theta_0, \theta_1, \theta_2$ depends on the values of z_0, z_1, z_2. We write the (ij)th term in the Jacobian dF as $\partial_i \theta_j$.

We are considering symmetric variations. We do not just move (z_0, z_1, z_2) but also we move (z_4, z_3, z_6) in a symmetric way. Thus, for instance

$$\partial_0 \theta_j = \frac{\partial \theta_j}{\partial z_0} + \frac{\partial \theta_j}{\partial z_4}.$$

The other expressions expand in a similar way. When we take second derivatives we get 4-term expressions. For instance

$$\partial_{01} \theta_j = \frac{\partial^2 \theta_j}{\partial z_0 z_1} + \frac{\partial^2 \theta_j}{\partial z_4 z_1} + \frac{\partial^2 \theta_j}{\partial z_0 z_3} + \frac{\partial^2 \theta_j}{\partial z_4 z_3}$$

Each θ_j is a 6-term expression of terms of the form θ_{ijk}, where θ_{ijk} is the angle of triangle (i, j, k) at vertex i. Hence, $\partial_{ij} \theta_j$ is a 24-term sum of expressions of the form

$$\partial(a, b, c, d, e) = \frac{\partial^2 \theta_{cde}}{\partial z_a \partial z_b}. \tag{21}$$

These expressions vanish when $\{a,b\} \not\subset \{c,d,e\}$.

We call $\partial(a,b,c,d,e)$ *nice* if $c \notin \{a,b\}$. Otherwise we call the term *mean*. Our terminology indicates our preference.

Lemma 7.1 *Every mean term is a sum of at most* 2 *nice terms.*

Proof: First consider terms of the form $\partial(a,a,a,b,c)$. Using the relation $\theta_a + \theta_b + \theta_c = \pi$ we have

$$\partial(a,a,a,b,c) = -\partial(a,a,b,c,a) - \partial(a,a,c,a,b).$$

The terms on the right are nice.

Given that $\partial(a,b,c,d,e) = \partial(b,a,c,d,e)$, the only remaining mean terms we have to treat are those of the form $\partial(b,a,a,b,c)$ and $\partial(b,a,a,c,b)$. We will treat the first one. The second one has the same treatment.

Thinking of our partial derivatives as operators on the function θ_{cde}, we have the basic formula

$$\frac{\partial}{\partial z_c} + \frac{\partial}{\partial z_d} + \frac{\partial}{\partial z_e} = 0.$$

Geometrically, the variation we are considering simply translates the triangle up or down along the Z-axis and changes none of the angles. Using this relation, we have

$$\partial(b,a,a,b,c) = -\partial(b,b,a,b,c) - \partial(b,c,a,b,c).$$

Both terms on the right are nice. ♠

In summary each partial derivative $\partial_{ij}\theta_k$ is a sum of K nice terms of the form $\partial(a,b,c,d,e)$ where $K \leq 24 \times 2 = 48$. In view of this breakdown, it suffices to show that

$$|\partial(a,b,c,d,e)| < 10^{22}$$

for all nice terms.

7.3 Normalized Pairs of Vectors

To avoid overloading the variable names, we write vectors as $V = (u,v,w)$. We say that a pair (V_1, V_2) of vectors is *normalized* if they have the form

$$V_1 = (0, v_1, w_1), \qquad V_2 = (u_2, v_2, w_2). \tag{22}$$

We say that a torus T' is *relevant* if it corresponds to a point in the ball of radius 10^{-27} about the point describing our example T in Equation 1. We say that a pair of normalized vectors is *relevant* if it is obtained from a triangle (P_0, P_1, P_2) in a relevant torus by the following procedure.

- We set $V_1' = P_1 - P_0$ and $V_2' = P_2 - P_0$.
- We rotate (V_1', V_2') about the Z-axis to get a normalized pair (V_1, V_2).

We mean to make this construction for all permutations of these points. Thus each triangle in a relevant torus gives rise to 6 pairs of relevant vectors.

Let $\vartheta(V_1, V_2)$ be the angle between V_1 and V_2. If we have an arbitrary pair of vectors we can rotate about the Z-axis so that they are normalized. This rotation does not change the geometry of the triangle defined by these vectors and it does not change the dependence of this geometry on the Z-coordinate. For this reason, each nice term discussed in the previous section has the form

$$\partial_{ij}\vartheta = \frac{\partial^2 \vartheta}{\partial w_i \partial w_j}, \qquad i,j \in \{1,2\}. \tag{23}$$

To finish the proof of the Crude Bound Lemma, we just need to prove for all relevant pairs that

$$|\partial_{ij}\vartheta| < 10^{22}, \qquad (i,j) = (1,2), (1,1), (2,2).$$

We don't need to deal with the case $(i,j) = (2,1)$ because $\partial_{21}\vartheta = \partial_{12}\vartheta$.

7.4 Concrete Formulas

Finally, we are at the point where we can get formulas. We have

$$\vartheta = \arccos\left(\frac{V_1 \cdot V_2}{\sqrt{(V_1 \cdot V_1)(V_2 \cdot V_2)}}\right) \tag{24}$$

This angle formula is more general than the one we considered in Equation 15. The expressions in Equation 23 are square roots of rational functions. We find it convenient to set $g_{ij} = (\partial_{ij}\vartheta)^2$ and to show that $g_{ij} < 10^{44}$ for all relevant pairs and indices. Our bounds below will be much better than this.

Lemma 7.2 *For each relevant pair (V_1, V_2) of normalized vectors (in the notation above) we have $\|V_1\|, \|V_2\| \in [0.8, 2.4]$ and $|u_2| > 0.005$.*

Proof: The norm bounds are straightforward. We just list out the possibilities in our example T, then observe that any other relevant torus differs from T by less than 10^{-27} in each coordinate.

The bound on u_2 is more subtle. For each "raw pair" (V_1', V_2') associated to our triangle we compute

$$a(V_1', V_2') = |u_1' v_2' - u_2' v_1'|.$$

This quantity measures the area of the parallelogram spanned by our vectors when we project them into the plane. In particular $a(V_1', V_2') = a(V_1, V_2)$ when (V_1, V_2) is the normalized pair associated to (V_1', V_2'). We check directly that $|a(V_1', V_2')| > 0.02$ for all relevant triangles. Hence

$$|u_2| = \frac{|u_2||v_1|}{|v_1|} = \frac{|a(V_1, V_2)|}{|v_1|} = \frac{|a(V_1', V_2')|}{|v_1|} > \frac{0.02}{3} > 0.005.$$

This completes the proof. ♠

Now we have three cases, which we treat in order of complexity. Let $g_{ij} = (\partial_{ij}\vartheta)^2$.

Case 1: We have

$$g_{12} = \frac{u_2^4 v_1^4}{\left(u_2^2(v_1^2 + w_1^2) + (v_2 w_1 - v_1 w_2)^2\right)^3}. \tag{25}$$

Write this expression as N/D. By Lemma 7.2 we have

$$N < (2.4)^8 < 10^5, \qquad D > (0.005)^6 (0.8)^6 > 10^{-15},$$

Hence $g_{12} < 10^{20}$.

Case 2: Next, we have

$$g_{11} = \frac{v_1^2 \left(-2w_1(v_2 w_1 - v_1 w_2)^3 + u_2^2(v_1^2 + w_1^2)(v_1^2 v_2 - 2v_2 w_1^2 + 3v_1 w_1 w_2)\right)^2}{(v_1^2 + w_1^2)^4 \left(u_2^2(v_1^2 + w_1^2) + (v_2 w_1 - v_1 w_2)^2\right)^3} \tag{26}$$

Using the same notation as in Case 1 and using the bounds from Lemma 7.2 we find that

$$D \geq (0.005)^6 (0.8)^7 > 10^{-16}.$$

We expand N out completely and find that it has 32 terms, all coefficients less or equal to 80 in absolute value, and maximum degree 16. This gives us

$$|N| < 32 \times 80 \times (2.4)^{16} < 10^{10}.$$

Putting together our bounds, we have $g_{11} < 10^{26}$.

Case 3: Finally, we have

$$g_{22} = \frac{\begin{pmatrix} 2v_2w_2(v_2w_1 - v_1w_2)^3 + u_2^4(v_1^3v_2 + 3v_1^2w_1w_2 + 2w_1^3w_2) \\ + u_2^2\big(4v_2^2w_1^3w_2 - 6v_1v_2w_1^2w_2^2 + 3v_1^2w_1w_2(v_2^2 + w_2^2) + v_1^3(v_2^3 - v_2w_2^2)\big) \end{pmatrix}^2}{(u_2^2 + v_2^2 + w_2^2)^4 \left(u_2^2(v_1^2 + w_1^2) + (v_2w_1 - v_1w_2)^2\right)^3} \tag{27}$$

We use the same notation as above. Essentially the same calculation as in Case 2 gives $D > 10^{-16}$. When we expand N out completely, we see that N has 55 terms, all coefficients less or equal to 124 in absolute value, and maximum degree 16. Hence

$$|N| < 5 \times 124 \times (2.4)^{16} < 10^{10}.$$

Hence $g_{22} < 10^{26}$.

This completes the proof of the Crude Bound Lemma, and thereby completes the proof of Theorem 1.3. An 8-vertex paper torus exists!

8 Appendix: Proof of the Hull Theorem

We only needed the Hull Lemma for the proof of Theorem 1.1, but our analysis gets us pretty close to the (unnamed) result from [**BE**] that a 7-vertex embedded polyhedral torus cannot have all 7 vertices on its convex hull boundary. We call this result the Hull Theorem. We build on the argument from §2.3 and prove the Hull Theorem. We emphasize that this argument is not needed for the proof of any other results in this paper. We include it so as to give a proof of the beautiful result from [**BE**].

Given the work above, we just have to prove that the internal edge pattern shown in Figure 2.2 does not correspond to an actual embedded 7-vertex polyhedral torus Ω. We repeat a small copy of the picture here for convenience.

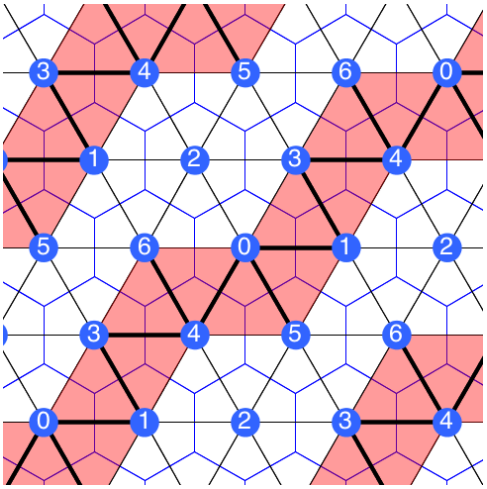

Figure 2.2: The one remaining pattern

Suppose for the sake of contradiction that Ω exists.

We think of Ω as a subset of projective space $\boldsymbol{P}^3$. We can apply a projective transformation so that vertex (2) moves to the point $[0:0:1:0]$ at infinity in $\boldsymbol{P}^3$. This point is "infinitely far away" along the Z-axis. We also can arrange that the (now) rays $(2j)$ start at (j) and move downward (rather than upwards) along the Z-axis, limiting on (2). We can do all this by a projective transformation that maps $H - \{(2)\}$ into the affine patch, which we identify with $\boldsymbol{R}^3$. Essentially, our normalization is a limit of examples in $\boldsymbol{R}^3$ in which (2) moves farther and farther down the Z axis.

Now, $H - \{(2)\}$ is still a convex subset in $\boldsymbol{R}^3$. It is like a hexagonal prism that has been truncated at one end. Here is the crucial observation: since H is convex, the projection of the hexagon (603541) into the XY-plane is a convex hexagon. We can further normalize so that the projections of (15) and (36) into the XY-plane are parallel line segments. We do this by mapping the line $[215] \cap [236]$ to a line at infinity which contains (2). Here $[abc]$ is the plane containing the points (a) and (b) and (c).

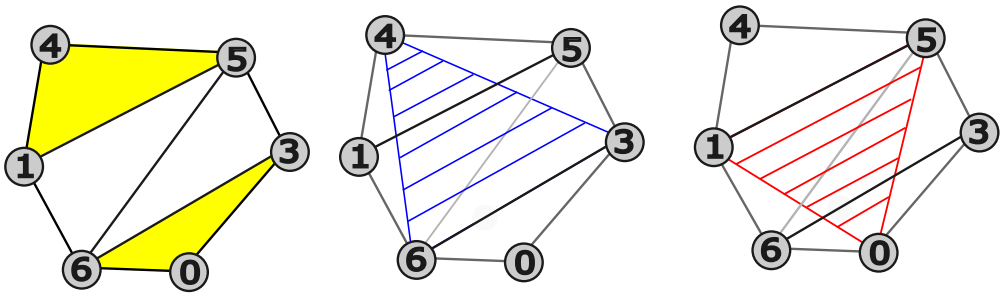

Figure 2.3: Projection of part of ∂H into the XY-plane

The 3-cycles (036) and (145) are both triangles on ∂H. These are not triangles of Ω. Every other triangle in ∂H does belong to $\partial\Omega$. By convexity, these two 3-cycles are bent downward. What we mean is that the plane containing (036) has the rest of H beneath it. The same goes for the plane containing (145).

The blue triangle in Figure 2.3 is (346). This is an internal triangle. We can foliate this triangle by parallel line segments as shown in the figure. These segments are parallel in space and they project to parallel segments in $\boldsymbol{R}^2$. The red triangle in Figure 2.3 is (015). We make all the same constructions for this internal triangle. We call these foliations *red* and *blue*, as in the figure.

Say that a *special plane* is a plane in $\boldsymbol{R}^3$ whose projection to $\boldsymbol{R}^2$ is a line parallel to the projections of the red and blue foliations. One special plane contains (15). In this plane, the red foliation is above the blue foliation. Another special plane contains (36). In this plane, the blue foliation is above the red foliation. So, by the intermediate value theorem, there is a special plane for which these two foliations coincide. But then (015) and (346) intersect. This contradicts the fact that Ω is embedded.

This completes the proof of the Hull Theorem.

9 References

[**ALM**] P. Arnoux, S. Lelievre, A. Malaga, *Diplotori: a family of polyhedral flat tori.* in preparation

[**Br**], U. Brehm, Oberwolfach report (1978)

[**BE**] J. Bokowski and A. Eggert, *All realizations of Möbius' Torus with 7 Vertices*, topologie structurale núm 17, (1991)

[**BZ1**] Y. D. Burago and V. A. Zalgaller, *Polyhedral realizations of developments* (in Russian) Vestnik Leningrad Univ. 15, pp 66–80 (1960)

[**BZ2**] Y. D. Burago and V. A. Zalgaller, *Isometric Embeddings of Two Dimensional Manifolds with a polyhedral metric into* $\boldsymbol{R}^3$, Algebra i analiz 7(3) pp 76-95 (1995) Translation in St. Petersburg Math Journal (3)3, pp 369–385

[**Cs**] Á. Császár, "A polyhedron without diagonals,"
Acta Sci. Math. (Szeged), vol. 13, pp. 140–142, 1949.

[**E**] M. Ellison, "Realizing abstract simplicial complexes with specified edge lengths" arXiv 2312.05376 (2023)

[**G**], M. Gardner, *Mathematical Games: On the remarkable Császár polyhedron and its applications in problem solving*, Scientific American **232**, 5 (1975) pp 102-107.

[**HLZ**], S. Hougardy, F H. Lutz, and M. Zelke, *Polyhedral Tori with Minimal Coordinates*, arXiv 0709.2794

[**LT**] F. Lazarus, F. Tallerie, *A Universal Triangulation for Flat Tori*, CS arXiv 2203.05496 (2024)

[**L**] F. H. Lutz, *Császár's Torus*, Electronic Geometry models (2002) 2001.02.069

[**QC**], P. Tanessi Quintanar Cortés, *Plongements polyédriques du tore carré plat*, PhD thesis, Université Claude Bernard Lyon 1, 2019. Available at `http://www.theses.fr/2019LYSE1354`.

[**T**] V. Tugayé, personal communication (2025)