The Most Efficient Origami Torus

Richard Evan Schwartz[a,c,1]

This manuscript was compiled on January 27, 2026

A paper torus, also called an origami torus, is a torus in three-dimensional space that is made from finitely many triangles which fit together in such a way that the sum of the angles around each vertex is 2π. This article announces the result that there does not exist an origami torus with 7 vertices and there does exist an origami torus with 8 vertices. These results settle the question about the most efficient way, in terms of the number of vertices, to construct an origami torus.

paper torus | origami | optimal construction | computer assisted proof

Significance Statement

This paper announces the solution of a basic and long-standing problem in computational geometry. Informally stated, the problem asks for the most efficient way to make a topological torus out of an origami construction. The efficiency here is measured in terms of the number of vertices required, or equivalently the number of folds needed. This paper also explains a computer-based experimental approach, based on something akin to supervised learning, to discovering the solution in a way that ought to resonate with a wide range of scientists.

1. Introduction

In this article I will describe the most efficient way, in terms of the number of vertices, to make a torus out of origami. My preprint (1) has all the mathematical details.

A. Flat Tori. A *flat torus* is a space you get by gluing the opposite sides of a parallelogram together by translations. The left side of Figure 1 shows a diagram for a flat torus, in the most familiar case when the parallelogram is a square. The arrows indicate how the sides are supposed to be glued together.

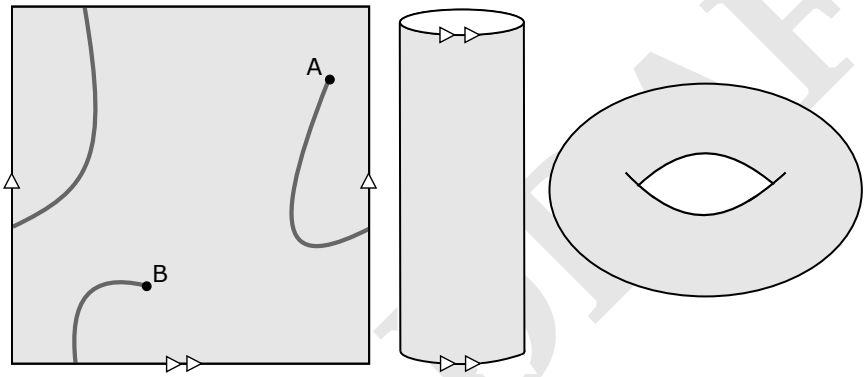

Fig. 1. The flat square torus

The path drawn in Figure 1, connecting point A to point B, is a continuous path in the flat torus. This path first goes through the right side and then pops out at the place on the left side which is meant to be glued to the right. Then the path continues up to the top and pops out the bottom before reaching B. Readers who have played the video game Asteroids will recognize that the spaceship might take such a path during the game. Essentially, a flat torus is a parallelogram-shaped screen with "wrap". We measure lengths of curves on a flat torus T using the Euclidean notion of length. Thus, the length of the curve γ connecting A to B in Figure 1 is the sum of the Euclidean lengths of the three arcs. The torus T also carries a *metric*. The distance between A and B is the arc length of the shortest continuous path connecting them.

We can attempt to put T inside three dimensional Euclidean space $\mathbf{R}^3$. The middle part of Figure 1 shows a cylinder we get by rolling up the square in space so as to implement the left-to-right gluing for T. This cylinder is an intermediate model for T. The right side of Figure 1 shows how we might complete the operation and bring the top and bottom loops together. Once we are done, we have a shape which is like the surface of a donut. If you try to do this with a piece of paper you will discover that you need to crumple the paper quite a bit to get the job done. If you use a stretchy rubber square, you could successfully implement the operations depicted in Figure 1.

Author affiliations: [a]Brown University, Department of Mathematics, 151 Thayer Street, Providence, RI, 02902

R.E.S. wrote the paper.

No competing interests

[2] Address correspondence to Richard Evan Schwartz. E-mail: Richard.Evan.Schwartz@gmail.com

B. Isometric Embeddings of Flat Tori. The discussion above raises the following question: Can we put a paper torus T in $\mathbf{R}^3$ in a way that does not locally stretch or compress the material? If we could do this, we would call it an *isometric embedding*. Formally, an isometric embedding is a continuous, one-to-one, and arc-length-preserving map

$$F : T \to \mathbf{R}^3 \qquad [1]$$

One-to-one means that $F(A) = F(B)$ only if $A = B$. *Arc-length-preserving* means that if γ is a curve on T with finite length, then $F(\gamma)$ has the same length at γ.

There are two kinds of such isometric embeddings. The first kind, called C^1 *isometric embeddings*, are those for which F is differentiable and has continuously varying partial derivatives. We mention these only for historical context. The refinement by N. Kuiper (2) of the famous Nash Embedding Theorem (3) implies that any flat torus admits a C^1 isometric embedding into $\mathbf{R}^3$. The image $F(T)$ has a wild, subtle, and often beautiful appearance. See e.g. V. Borrelli et al. (4) and the references therein for a modern treatment. The second kind of isometric embedding, the subject of this article, is what we call a *paper torus embedding*.

C. Paper Tori. The map F in Equation 1 is a *paper torus embedding* if F is an isometric embedding and it has the following additional property: T has a finite decomposition into triangles, a *triangulation*, so that the restriction of F to each triangle τ is an affine map. An *affine map* between Euclidean spaces is the composition of a linear transformation and a translation. In other words, $F : \tau \to \mathbf{R}^3$ is not only a distance preserving map, but also affine. The image $F(\tau)$ is a geometric triangle having exactly the same shape and size as τ. The image $F(T)$ is called a *paper torus*; it is made by fitting finitely many triangles together to make an embedded torus so that around each vertex the total angle of the triangles is 2π. These angle sums are called the *cone angles*. The 2π cone angle property is what we might call a *local flatness condition*. We can take the triangles which surround a single vertex and straighten them out so that they lie flat in the plane, like pizza slices fitting together to make a full pizza.

A paper torus is an extremely concrete object. At least in principle, and subject to a certain interpretation of the word *can*, you can take a paper torus embedding F as in Equation 1 and use it as instructions for folding up the flat torus T along the edges of the triangulation to make $F(T)$. For this reason, a paper torus is also sometimes called an *origami torus*.

Equation 1 is a redundant way to describe a paper torus. If we have managed to put together finitely many triangles in $\mathbf{R}^3$ so that they make an embedded torus Ω having all cone angles equal to 2π, then there automatically exists a paper torus embedding as in Equation 1 so that $\Omega = F(T)$. The reconstruction of F and T works as follows. We take an infinite supply of all the triangles used in making Ω and fit them together in the plane with the same side pairings as we see in Ω. The 2π cone angle condition says that the pattern we get will lie flat in the plane. When we do this process forever, so to speak, we build a doubly-periodic pattern of triangles of the kind shown in Figures 2,3,5,7, and 8 below. This pattern will have symmetry generated by a pair V, W of translation vectors in the sense that if we translate the plane by either V or W we get the same picture back. Our flat torus T is obtained by cutting out the parallelogram with vertices $0, V, W, V + W$ and gluing the opposite sides. Compare Figure 2. The pattern in the plane automatically wraps up to give a triangulation of T, and F carries each triangle in T to its location in Ω. We mention this property of paper tori for two reasons. First, when we construct these things, we will not necessarily want to deal explicitly with F or T. Second, as we have already indicated, many of our figures below depict objects on the torus as doubly-periodic patterns in the plane. In formal terms, we are drawing the *lifts* of the objects in the torus to the *universal cover* of the torus and then depicting the doubly-periodic pattern they make.

It is surprising that paper tori exist. Yu. D. Burago and V. A. Zalgaller (5) provided the first examples in 1960. Many years later, in 1995, the same authors proved (6) that one can realize every flat torus as a paper torus. What this means is that in Equation 1 one can produce an isometric embedding for any flat torus T (based on any parallelogram you like). By now there are a variety of constructions of paper tori. See the work of U. Brehm (7), T. Tsuboi (8), and P. Arnoux et al. (9). The recent paper of F. Lazarus and F. Tallérie (10) proves that one can use the same combinatorial structure (just varying the geometry of the triangles) to realize every flat torus as a paper torus. The study of paper tori lies at the intersection of combinatorial topology, computational geometry, and origami-informed mathematics.

D. The Minimum Vertex Problem. As with many mathematical problems, one can view paper tori through the lens of optimization. Given that they exist, how efficiently can they be made? I don't know when this question was first asked, but it seems like one of the first things you would want to know about paper tori beyond the fact of their existence. A nice measure of efficiency is given by the number of vertices. (This is the same as asking about the minimum number of triangles in the triangulation, or the minimum number of edges you fold along.) The examples in (5), (6), and (10) involve many vertices. For the "universal triangulation" in (10), the number is in the thousands.

The constructions in (7), (8) and (9), which are all based on the examples in (7), involve a much smaller number of vertices. The minimal example in these works has 10 vertices. V. Tugayé recently produced an explicit 9-vertex example, which I saw in the summer of 2025 when I visited Paris. Going from the other direction, we can say that a paper torus requires at least 7 vertices because there are no triangulations of a torus having fewer than 7 vertices. The results above leave open the possibility that the answer to the minimum vertex question for paper tori is 7, 8, or 9. This article gives the answer: 8.

2. The 7-Vertex Case

A. The Császár Torus. There is only one 7-vertex triangulation of a flat torus, up to the freedom of allowing the vertices to slide around. This triangulation has 7 vertices, 21 edges, and 14 triangles. Every vertex is connected to every other by an edge. If you cut the shaded parallelogram out of the diagram in Figure 2, then glue the opposite edges together, you will get the desired triangulation of the resulting flat torus. This depiction of the 7-vertex triangulation follows our practice, mentioned above, of depicting objects in a flat torus as doubly periodic patterns.

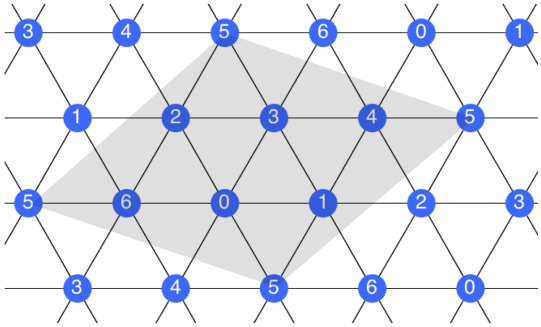

Fig. 2. The 7-vertex triangulation

Now we describe a broader class of surface that subsumes paper tori. A *polyhedral torus* is a finite collection of triangles in $\mathbf{R}^3$ which fit together to make an embedded torus. Here we are dropping the 2π cone angle condition. For the sake of comparison, we can also say that a polyhedral torus is the image $F(T)$ where $F : T \to \mathbf{R}^3$ is a *polyhedral torus embedding*. This is to say F is one-to-one and continuous and has the following additional property: T has a triangulation such that the restriction of F to each triangle is an affine map. In contrast to the paper torus case, F need not be an isometric embedding here and F need not preserve the shapes of the triangles. We will never have the occasion to bring up F and T in the context of polyhedral tori. We mention them here only to more fully compare paper tori and polyhedral tori. For emphasis, we say again that all paper tori are polyhedral tori, but only some polyhedral tori are paper tori.

In 1949, Ákos Császár (11) showed that it was possible to produce a 7-vertex polyhedral torus. This amounts to finding 7 points in space such that, when we make all 14 triangles using the distinct triples (such as 013 and 134) which arise in the triangulation, they fit together to make an embedded torus. Can Császár's construction be modified to get a 7-vertex paper torus? You could imagine starting with Császár's example and somehow jiggling the vertices to make all the cone angles add to 2π. In 2019, P. Quintanar-Cortés (12) made partial progress towards proving that this is impossible. Here I will give a complete proof.

B. Some Geometric Notions. A set H is *convex* if it has the following property: If p and q are any two points in H, then the line segment joining p to q also lies in H. The *convex hull* of a set S is the intersection of all the convex subsets that contain S. The convex hull of a polyhedral torus is the smallest convex polyhedron containing the vertices of the torus. A *sector* in space is the planar region between two rays which have a common endpoint and make an angle of less than π. The sector consists of all the rays that interpolate linearly between the two outermost rays. A *fan* is a finite union of sectors joined end to end in a cyclic pattern. Given a vertex p of a polyhedral torus Ω, the union of the triangles which surround p is called the *flower at* p. The flower at p in turn defines the fan F_p. This is the union of all rays emanating from p whose initial parts lie in the flower.

C. Proof Outline. Let ℓ denote arc-length on a unit sphere. Let Ω be a polyhedral torus. Let H be the convex hull of Ω and let ∂H be its boundary. In Császár's original example, 5 vertices lie in ∂H and 2 vertices lie in the interior $H^o = H - \partial H$. In 1991, J. Bokowski and A. Eggert (13) proved that, in general, at least one vertex p of Ω lies in H^o. (Below we will discuss in detail a variant of this result.) The vertex p is joined by an edge to every other vertex, including all the vertices in ∂H. Consider the fan F_p. The intersection of F_p with the unit sphere centered at p is a spherical polygon P_p, a loop made from arcs of great circles. The length $\ell(P_p)$ coincides with the cone angle at p. As we explain below, if $\ell(P_p) \leq 2\pi$ then P_p is contained in some hemisphere. What this means is that there is a plane Π through p such that the F_p is contained in one of the (closed) halfspaces bounded by Π. But then all the points of ∂H lie in this halfspace. This is impossible if $p \in H^o$. The convex hull has to "surround" p, so to speak. So, the cone angle at p exceeds 2π. This means that Ω is not a paper torus.

D. Spherical Polygons. Now I will explain why the bound $\ell(P_p) \leq 2\pi$ forces P_p to lie in some hemisphere. We will show this when $\ell(P_p) < 2\pi$. The case of equality follows by taking a limit. We apply the famous *Crofton formula* (14): $\ell(P_p)/\pi$ equals the average number of times (taken with respect to uniform measure) that a great circle intersects P_p. If $\ell(P_p) < 2\pi$ then some great circle C intersects P_p at most once. Being an embedded loop, P_p would have to intersect C at least twice if P_p contained

points in the interior of both hemispheres bounded by C. Hence P_p is contained in one or the other hemisphere bounded by C. Let us give a painless proof of the version of Crofton's formula that we use. Suppose more generally that γ is a finite union of spherical arcs with $\ell(\gamma) < 2\pi$. Let $A(\gamma)$ be the average number of times a great circle intersects γ. We want to show that $A(\gamma) < 2$. By symmetry, $A(\gamma)$ does not change if we take one of the arcs of γ and slide it around by an isometry of the unit sphere. So, we can slide all the arcs of γ so that they fit in a single great circle γ' and leave some gaps. But every great circle except γ' intersects γ' exactly twice. (For the purposes of computing averages, we are free to ignore finitely many exceptional great circles.) This gives $A(\gamma') = 2$, and so $A(\gamma) < 2$.

E. Flowers and Convexity. Below I will prove what I call the *Flower Result*: If Ω has all 7 vertices in ∂H then Ω also has a flower in ∂H. The Flower Result is a bit weaker than the Bokowski-Eggert result cited above, but we can use it in place of the Bokowski-Eggert result in our impossibility proof, and at the same time we can sketch a self-contained explanation for why the Flower Result is true.

If Ω has a vertex in H^o then we proceed with our proof, as above. If not, then Ω has a flower $F \subset \partial H$. The cone angle of F is 2π and also F matches the boundary of a convex polyhedron in the vicinity of p. The only way this can happen is that the 6 triangles of F lie in a union $\Pi_1 \cup \Pi_2$ of 2 half-planes which meet along a line through p. The picture looks kind of like a V-shaped valley. But then at least 3 of the triangles in the flower, and hence at least 5 points of Ω, lie in one of the half planes. Since Ω is embedded, and every edge of the triangulation is an edge of Ω, we would have 5 vertices in the same plane, joined each to each by edges which meet only at the endpoints. In other words, we would have a planar embedding of K_5, the complete graph on 5 vertices. This contradicts the famous graph theory result, due to Leonhard Euler, that there are no planar embeddings of K_5. This contradiction shows that Ω does not exist.

F. The Flower Result. Now we sketch the proof of the Flower Result. Suppose Ω has all 7 vertices in ∂H. Recall that Ω has 7 vertices and 21 edges and 14 faces. Using the Euler characteristic (namely that in a triangulation "faces plus vertices minus edges equals 2"), and the fact that every vertex on ∂H is connected to every other by an edge, we compute that the triangulation of ∂H has 7 vertices and 10 faces and 15 edges – at least after we perturb slightly so that the points are in general position. (Since we do not care about the cone angles here, this perturbation is harmless.) One heuristic way to get this count is to start with a tetrahedron, which has 4 vertices, 6 edges, and 4 faces. When we add a point we increase the number of edges by 3 and the number of faces by 2. This moves us from the (vertex,edge,face) triple $(4, 6, 4)$ all the way to the triple $(7, 15, 10)$.

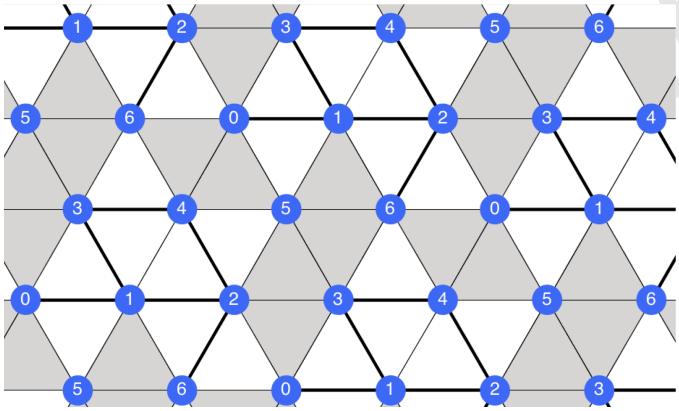

Fig. 3. Internal edges in triangulation

Exactly $6 = 21 - 15$ edges do not lie in ∂H. We call these edges *internal* and the other 15 *external*. Figure 3, a decorated version of Figure 2, shows one combinatorial possibility. In Figure 3 the 6 internal edges are 01, 12, 13, 24, 34, 26. Our argument, which we explain by way of example, hinges on comparing the picture in Ω with what we can infer about the picture in ∂H from the diagram of internal edges. Consider Vertex 6 in Figure 3. Vertex 6 is incident to 5 external edges, connecting to Vertices $1, 5, 3, 4, 0$. If we stand on ∂H at Vertex 6 and twirl around we will see these other vertices in some order. We cannot see them in the order listed: If Vertices 3 and 4 appear consecutively, then Edge 34 is external. This is a contradiction.

Suppose that our 5 vertices wrap around Vertex 6 on S in some other order, say $1, 5, 3, 0, 4$. (This ordering at least avoids the previous contradiction.) To rule out this order (and indeed any other) we compare the fans $F_6\partial$ and $F_6\Omega$ at Vertex 6, respectively defined in terms of ∂H and Ω. Because H is a convex polyhedron, $F_6\partial$ is the boundary of a convex pentagonal cone. If we slice $F_6\partial$ by a plane Π that is offset a bit from Vertex 6 we will see a convex pentagon. The fan $F_6\Omega$ intersects Π in some hexagonal path that is contained in the region bounded by the pentagon. The slice of $F_6\Omega$ visits the pentagon vertices in the order $1, 5, 3, 0, 4$. Figure 4 shows how this looks.

In Figure 4, the blue edges near the boundary are supposed to exactly trace over the black edges. We have offset them for the sake of illustration. Notice that the order of the points forces this blue polygon to cross itself. But then the fan $F_6\Omega$ is not embedded, as it would be when Ω is a polyhedral torus. The same argument rules out every other order besides the initial one, which we have already ruled out.

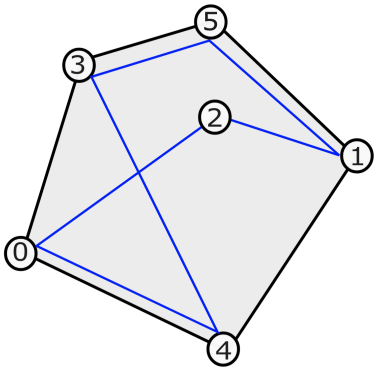

Fig. 4. Planar slice of $F_6\partial$ in black and $F_6\Omega$ in blue

In short, we have examined the pattern of internal edges in our example and we have seen that this pattern cannot actually come from a polyhedral torus with all 7 vertices on the convex hull boundary. There are $\binom{21}{6}$ different patterns we need to analyze like this and, after using symmetry, just $\binom{20}{5} = 15504$. Our proof enumerates all possible diagrams like this and systematically rules out all but 6 patterns. The 6 remaining patterns correspond to cases where there is a flower in the convex hull boundary. See (1) for a more detailed treatment. Also, you can download my software at **http://www.math.brown.edu/∼res/Java/SevenTorus.tar** and run the program yourself.

3. The 8-Vertex Case: Experimentation

A. First Attempt Fails. Unlike in the 7-vertex case, where there is a combinatorially unique way to triangulate the torus, there are 7 combinatorially distinct ways to triangulate a torus with 8 vertices. Frank H. Lutz has compiled a wealth of information about triangulations of surfaces and higher dimensional manifolds. You can visit his "manifolds page" (15) and download any of the triangulations he has compiled, including the seven 8-vertex triangulations of a torus. This is what I did. Motivated by my proof in the 7-vertex case I first tried to use a combinatorial analysis, as above, to rule out the existence of an 8-vertex polyhedral torus having all 8 vertices, and no flowers, in its convex hull boundary. I analyzed about 10 million patterns and ruled out all but several thousand. Unable to rule out everything, I changed my mind and decided to look for 8-vertex paper tori.

B. Birds in Abundance. Let me say that a *bird* is an 8-vertex polyhedral torus with all 8 vertices on the convex hull boundary, and a *Phoenix* is a bird that is also a paper torus. My idea was to produce a large supply of birds and then to use a hill-climbing algorithm to mutate one of them into a Phoenix. The basic idea is to choose one of the triangulations, then place 8 points at random on the sphere and build the polyhedral triangulation. If this does not produce a polyhedral torus we simply repeat. For six of the seven 8-vertex triangulations, this recipe turns out to produce birds in abundance. In particular, this shows that my effort to rule out the existence of birds was doomed.

C. Hill Climbing. Generally speaking, a *hill climbing algorithm* involves a sample space X and an objective function $\Phi : X \to \mathbb{R}$. One attempts to find the maximum value of Φ on X by the following simple procedure. Keep track of some point $p \in X$, initially chosen at random. Choose a nearby point $q \in X$ and replace p by q if $\Phi(q) > \Phi(p)$. Repeat. Whether or not this algorithm meets with success has everything to do with the fine points in its design.

In a typical situation, the space X will be a nice subset of Euclidean space, like a ball, and Φ will be a differentiable function. As long as the dimension of X is not too large, the hill-climbing algorithm roughly mimics gradient flow. If you pick points at random very near to the point p, you don't have to wait long before you get a difference $q - p$ which roughly points in the direction of the gradient $\nabla\Phi$. Then Φ will likely increase when you move from p to q. The beauty of this procedure is that you essentially "do gradient flow" without having to compute the gradient! In cases when Φ is a complicated function whose gradient is difficult to compute symbolically, this feature is a life-saver.

D. Finding the Phoenix. As I said above, my basic idea is to start with a bird and try to mutate it into a Phoenix *via* hill climbing. Within this framework, I take my space X to be the space of birds, and Φ to be the function $\Phi = -\max_i |2\pi - \theta_i|$. Here θ_i is the sum of the angles of the triangles around Vertex i, and $i = 0, 1, 2, 3, 4, 5, 6, 7$. The minus sign serves to turn this into a maximization problem rather than a minimization problem. The maximum of Φ would be 0 if a Phoenix exists amongst the birds. Eventually I got this to work. I will describe some design choices which helped with the success.

Normal/Tangential Separation: The first good design choice is to separate the random evolution into a "tangential direction" and a "normal direction". If you have 8 points in convex position and you move the points at random, you are fairly likely to destroy the convexity. The way I deal with this problem is to consider pairs (S, E) where S is a collection of 8 points on the surface E of an ellipsoid. I first vary the points within E and then I apply a random near-identity linear transformation to vary the shape of E (and the points) slightly. I think of the first variation as "tangential" and the second as "normal". This method preserves the convexity property. Of course, I only accept the variation if it also preserves the embedding property.

Low Face Number: Define the *face number* of a bird to be the number of triangles it has in its convex hull boundary. To be clear, the convex hull boundary (generically) has 12 triangular faces, but not all of them need belong to the bird. In discussing the 7-vertex impossibility result, we already saw that the existence of a flower in the convex hull boundary kills any chance for a bird to be a Phoenix. Relatedly, I noticed experimentally that birds with high face number often have a convex fan at some vertex even if the entire flower surrounding that vertex did not live in the convex hull boundary. I decided to only accept birds for the algorithm which have face number at most 6.

Degeneration Guards: As the hill climbing process operates on a bird, there is a danger that the bird somehow gets worse in other ways even though its cone angles approach 2π. Here are the main issues: The ratio of largest distance between the vertices to smallest distances might blow up; some of the dihedral angles (the angles measuring how faces fold) might tend to zero; the ellipsoid containing the points might be getting extremely eccentric. To avoid these kinds of problems, I added thresholds to the algorithm, so that it automatically rejects birds if any of the above measurements gets too extreme. I found that setting the minimum dihedral angle at about 10^{-5} and the max/min ratio, both for the distances and the ellipse axes, to about 5 worked very well.

Controlled Randomness: To avoid having the evolving bird fall into a kind of rut, or local minimum that is not a Phoenix, I introduce a parameter σ as follows. At each step of the evolution, I generate a random number $r \in (0, 1)$. If $r < \sigma$ I accept the next bird in the process regardless of whether it does better with the objective function. When $\sigma = 0$ this feature is never used. When σ is small it is used only rarely. The closer σ is to 1, the more of a randomizing effect this feature has.

The Best Triangulation: Of the 6 bird-producing 8-vertex triangulations of the torus, I have a favorite one. In this one, each vertex is connected to 6 others. Succinctly, the triangles in this triangulation have the form $(j, j+1, j+3)$ and $(j, j+2, j+3)$ with indices taken mod 8. Figure 5 shows this triangulation. The thickly drawn polygon in Figure 5, while not a parallelogram, can be glued together to make a flat torus. Initially I tried to find a Phoenix using each of the 6 fruitful triangulations but eventually I narrowed my focus and stuck to the triangulation shown in Figure 5. The symmetry and regularity of the triangulation appealed to me.

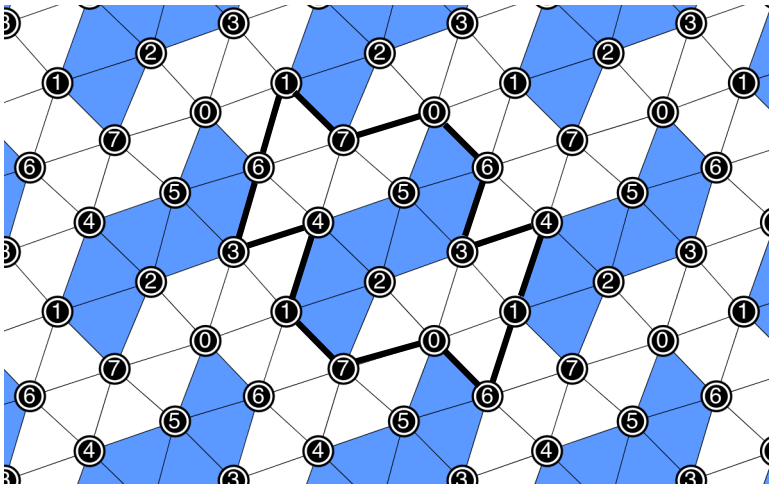

Fig. 5. The best 8-vertex triangulation of the torus

Two Fold Symmetry: With the features above put in place, my program produced a few Phoenixes up to numerical precision. Some had face number 4 and some had face number 6. After I told Samuel Leliévre about my findings, he asked if I could make one with 2-fold symmetry. I was not sure at first, but then I found a Phoenix with a rough kind of rotational symmetry and also the external face pattern shown in Figure 5. Once I saw that I knew what to do. Rather than specify all 8 points, I specified the 4 vertices numbered $0, 1, 2, 3$ and then I arranged that vertices $7, 6, 5, 4$ respectively are the 180 degree rotations of $0, 1, 2, 3$ around the Z-axis. This way of doing things makes sense because there is a combinatorial isomorphism of the pattern in Figure 7 which has the action $j \to 7 - j$. This symmetry simplified the algorithm, and made it run faster and more reliably.

E. Success and Digit Cleaning. With all these design features in place, my algorithm fairly reliably turns birds into (approximate) Phoenixes. These approximate Phoenixes have cone angles equal to 2π up to an error of about 10^{-15}. Subsequently, I implemented a high precision version of Newton's method. This method holds all of the coordinates fixed except the Z-coordinates of vertices $1, 2, 3, 4, 5, 6$, with the symmetry constraints $z_1 = z_6$ and $z_2 = z_5$ and $z_3 = z_4$. Using Newton's method, I improved a few of my examples so that the cone angles are within 10^{-400} of 2π. This amply satisfied me that these numerical approximations are close to true Phoenixes. Finally, I cleaned up my examples by simplifying all the numbers which do not vary during Newton's method. Here is my best example in this regard. All coordinates are exact except z_1, z_2, z_3, which are truncated after 32 digits.

$$\begin{array}{ccc} +0.64 & -0.20 & 1 \\ -1.09 & +0.38 & z_1 \\ -0.25 & +0.51 & z_2 \\ +0.78 & +0.62 & z_3 \\ -0.78 & -0.62 & z_3 \\ +0.25 & -0.51 & z_2 \\ +1.09 & -0.38 & z_1 \\ -0.64 & +0.20 & 1 \end{array}$$

$$\begin{aligned} z_1 &= 0.0206\ 6632\ 6669\ 8443\ 6159\ 8992\ 3371\ 8861 \\ z_2 &= 0.0048\ 5312\ 7706\ 5192\ 8720\ 4090\ 7479\ 6169 \\ z_3 &= 0.0082\ 2752\ 1455\ 6137\ 1645\ 5791\ 2547\ 8661 \end{aligned} \quad [2]$$

F. Some Pictures. Based on its 3-dimensional appearance I call my example a *pup tent* (rather than a Phoenix). Peter Doyle suggested this name after he made some 3-dimensional plots. This term does not refer just to a single example, but rather to a family of examples. In (1) I list specific properties that an 8-vertex paper torus must satisfy in order to qualify as a pup tent, but these axioms do not seem important here because I am just focusing on one example. Figure 6 shows 2 Mathematica (16) plots of the pup tent from Equation 2, one view from the side and one view from the top.

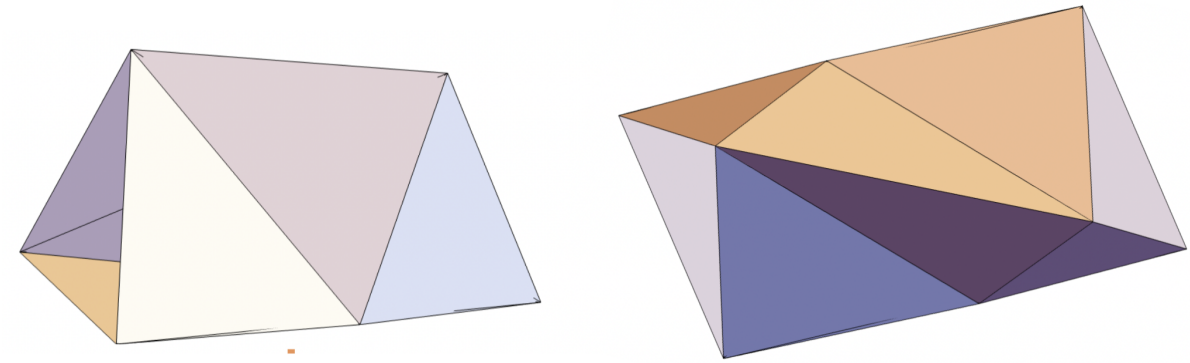

Fig. 6. Two Mathematica 3D plots of the pup tent

These 3D projections do not really show the complicated and fragile way that the pup tents are embedded in space. I would say that these pup tents barely exist. Figure 7 shows a subtle rendition of the pup tent. Referring to Equation 1, the pup tent is $\Omega = F(T)$. To generate Figure 7, I picked a parallel family of planes and sliced Ω by a plane in this family. The right side of Figure 9 shows this slice. The black polygon on the boundary is the slice of Ω itself, and the yellow interior is the slice of the solid body in space, often called a *handlebody*, bounded by Ω. I have included a close-up of the top part of the right-hand picture, to show that it really is embedded. The left side of Figure 7 shows the doubly periodic pattern in the plane that encodes various objects on T, including the triangulation. The thick black polygons on the left correspond to the curve in T that F maps to the slice on the right. We can completely slice Ω by parallel planes in our family, and this will give a partition of Ω into curves. The fingerprint curves on the left show these same curves on T. Thus, the left side of Figure 7 is something like a topographical map of Ω, with one elevation highlighted corresponding to the slice on the right.

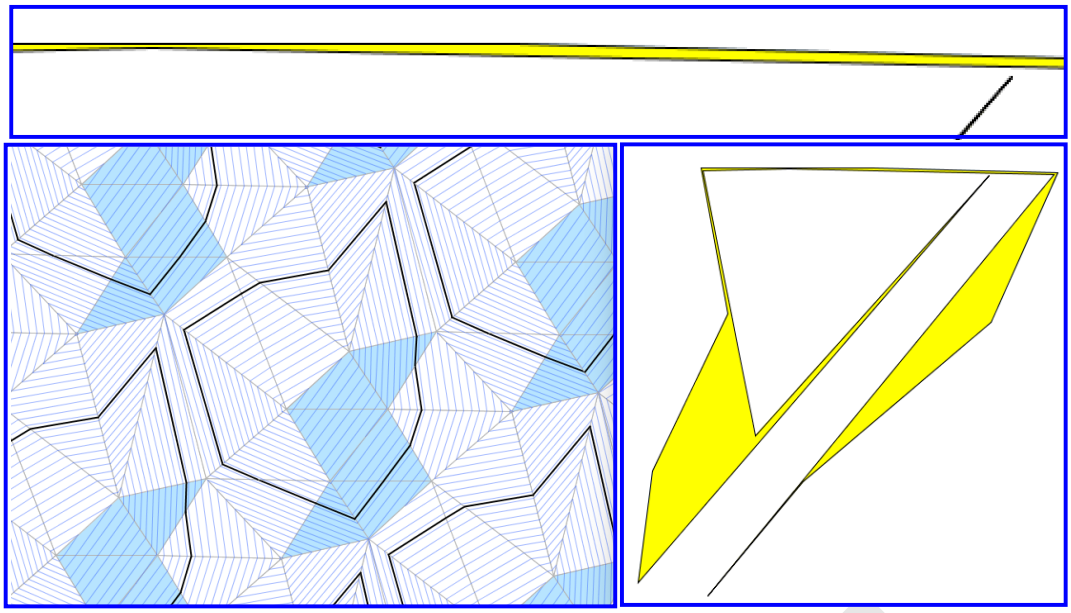

Fig. 7. A topographical map of the pup tent with respect to some parallel family of planes

Figure 8 shows the same thing, but with respect to a different parallel family of planes. This time we take planes parallel to the XZ-coordinate plane. The figure on the right is bounded by two embedded polygonal loops and they bound the yellow annulus. Once again, this annulus is barely embedded: The spike at the top just misses the original part on top. In the doubly-periodic picture on the left, the corresponding loops appear as bi-infinite paths. The symmetry comes from the fact that the XZ plane contains the Z-axis, which is the axis of symmetry for Ω.

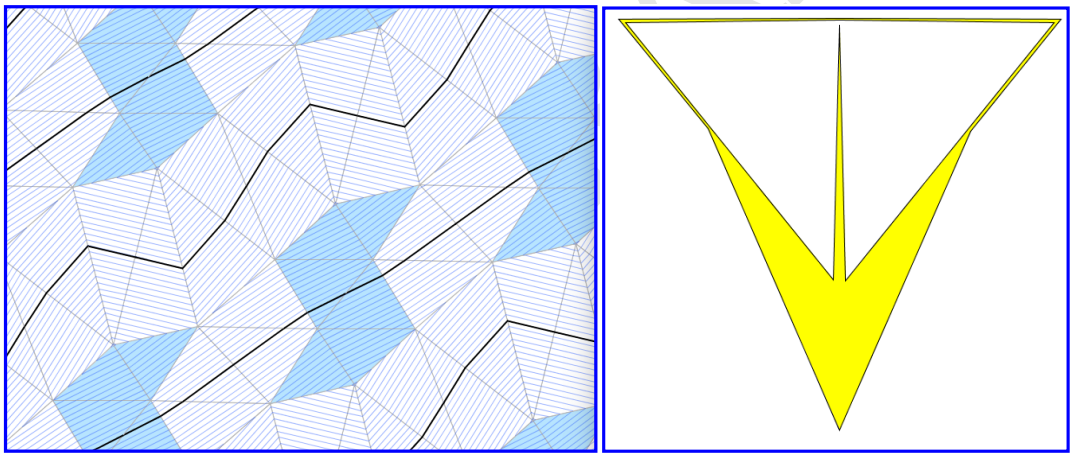

Fig. 8. A topographical map of the pup tent with respect to the XZ-parallel planes

4. The 8-Vertex Case: Sketch of a Rigorous Proof

In this chapter I will sketch the rigorous mathematical proof that an 8-vertex paper torus exists. In (1) I work out all the details. Briefly, the idea is to take the example from Equation 2, which is only off by a tiny bit, and then prove that one can jiggle the points a bit to make the cone angles exactly 2π without ruining the embedding property.

A. The Three Step Approach. Let $\Omega = \Omega(z_1, z_2, z_3)$ be the example in Equation 2, with the variables z_1, z_2, z_3 specified there. Here we are thinking of Ω concretely as a union of 14 triangles in $\mathbf{R}^3$. Let $(\theta_1, \theta_2, \theta_3)$ be the corresponding triple of cone angles. Let $\Omega' = \Omega(z_1', z_2', z_3')$ denote an example like the one in Equation 2 except that we have changed z_k to z_k' for $k = 1, 2, 3$. Let $(\theta_1', \theta_2', \theta_3')$ be the corresponding cone angles.

Step 1: We prove that Ω is 10^{-4}-*robustly embedded*. What we mean is that any choice of Ω' is embedded as long as

$$\max(|z_1 - z_1'|, |z_2 - z_2'|, |z_3 - z_3'|) \leq 10^{-4}.$$

This step gives us some wiggle room. We know that we can jiggle the points a little bit without ruining the embedding property.

Step 2: We prove that Ω is 10^{-31}-*flat*. What we mean is that $\max(|\theta_1 - 2\pi|, |\theta_2 - 2\pi|, |\theta_3 - 2\pi|) < 10^{-31}$. This step tells us that we are quite close to a paper torus. Given this very tight bound, we shouldn't need to change the z coordinates much to get a true pup tent.

Step 3: We introduce the map $\Theta(z_1, z_2, z_3) = (\theta_1, \theta_2, \theta_3)$. Let B denote the ball of radius 10^{-12} centered at $p = (z_1, z_2, z_3)$. Let B' be the ball of radius 10^{-14} centered at $\Theta(p) = (\theta_1, \theta_2, \theta_3)$. We prove that $B' \subset \Theta(B)$.

Combining Steps 2 and 3, and noting that $(2\pi, 2\pi, 2\pi)$ is less than 10^{-14} away from $\Theta(z_1, z_2, z_3)$, we have $(2\pi, 2\pi, 2\pi) \in \Theta(B)$. But then there is some point $(z_1', z_2', z_3') \in B$ such that $\Theta(z_1', z_2', z_3') = (\theta_1', \theta_2', \theta_3') = (2\pi, 2\pi, 2\pi)$. At this point we let $\kappa_j' = 2\pi - \theta_j'$. The quantity κ_j' is called the *vertex curvature* of Ω' at Vertex j. The famous "combinatorial Gauss-Bonnet Theorem" tells us that the sum of the vertex curvatures of any polyhedral torus is 0. Thanks to this fact, and to symmetry, the fact that $\kappa_1' = \kappa_2' = \kappa_3' = 0$ readily implies that $\kappa_i' = 0$ for all i. This proves the existence of an 8-vertex paper torus.

B. Details of Step 1. To prove that any given torus Ω' is embedded we just have to prove that all the triangles intersect as dictated purely by their pattern of vertices. Suppose that (Δ_0', Δ_1') is a pair of triangles. We have to check the following implications.

1. If Δ_0' and Δ_1' have no vertices in common then $\Delta_0' \cap \Delta_1' = \emptyset$. The triangles have empty intersection.
2. If Δ_0' and Δ_1' share a single common vertex v then $\Delta_0' \cap \Delta_1' = v$. The triangles only intersect at their common vertex.
3. If Δ_0' and Δ_1' share a single common edge e then $\Delta_0' \cap \Delta_1' = e$. The triangles only intersect at their common edge.

It turns out that if we can check the first two conditions for all triangle pairs that satisfy them, then the third condition follows automatically. The reason is that a bad intersection of the third kind above "drags the neighboring triangles into the mix" and creates a bad intersection of the second kind. Now, there are 24 pairs of disjoint triangles in our triangulation and 72 pairs of triangles having one vertex in common. This leaves us with 96 things to check. We don't want to check these things for each individual polyhedral torus Ω', so we make a more robust kind of check just for our original Ω and then explain how it covers all the tori of interest to us.

No Common Vertex: Suppose (Δ_0, Δ_1) is a pair of triangles for our original example Ω and that these triangles have no common vertices. Given a vector L we define

$$m_j(L) = \min_{v \in \Delta_j} v \cdot L, \qquad M_j(L) = \max_{v \in \Delta_j} v \cdot L. \tag{3}$$

By convexity, we can compute these quantities just by computing the dot products on the vertices.

We introduce one more piece of notation. We let $\|L\|_\infty$ denote the maximum absolute value of a coordinate of L. We call (Δ_0, Δ_1) λ-*separated* if there exists a vector L with $\|L\|_\infty \leq 1$ such that one of the two equations holds:

$$M_0(L) + \lambda < m_1(L), \qquad M_1(L) + \lambda < m_0(L), \tag{4}$$

We have $\Delta_0 \cap \Delta_1 = \emptyset$ as long as this pair is 0-separated.

We show by direct calculation that all 24 pairs are 2×10^{-4} separated. By suitably scaling all the vectors, we make an exact integer calculation for this. The details are done carefully in (1). Consider the implications for any other torus Ω' of interest to us. Recall that we have $|z_k' - z_k| \leq 10^{-4}$ for each $k = 1, 2, 3$. Consider the triangle pair (Δ_0', Δ_1') of Ω' that corresponds to (Δ_0, Δ_1). In moving from Δ_j to Δ_j', no vertex moves by more than 10^{-4}, and at most one of the coordinates changes per vertex. Hence, each dot product above changes by at most 10^{-4}. By the triangle inequality, (Δ_0', Δ_1') is 0-separated. Hence $\Delta_0' \cap \Delta_1' = \emptyset$.

One Common Vertex: Now we consider when (Δ_0, Δ_1) have a single common vertex v. Let $\Delta_j^\sharp$ be the edge of Δ_j opposite v. Define

$$m_j^\sharp(L) = \min_{w \in \Delta_j^\sharp} w \cdot L, \qquad M_j^\sharp(L) = \max_{w \in \Delta_j^\sharp} w \cdot L. \tag{5}$$

This time we call (Δ_0, Δ_1) λ-*separated* if there exists a vector L with $\|L\|_\infty \leq 1$ such that one of the two equations holds:

$$M_0^\sharp(L) + \lambda < v \cdot L < m_1^\sharp(L) - \lambda, \qquad M_1^\sharp(L) + \lambda < v \cdot L < m_0^\sharp(L) - \lambda, \tag{6}$$

If (Δ_0, Δ_1) is 0-separated then $\Delta_0 \cap \Delta_1 = v$. Using the same scaling trick, we show by direct calculation that all 72 pairs are 2×10^{-4}-separated. But now the same argument as in the no-common-vertex case shows that for Ω' as above all such pairs are 0-separated. Hence $\Delta_0' \cap \Delta_1' = v'$.

Vector Production: Let Q be the set of integer vectors V such that $\|V\|_\infty = 300$. There are several million vectors in Q, all lying on the surface of a cube. In each of the 96 cases, I list out the vectors in Q, scaling down by $1/300$, until I find one which satisfies Equation 4 or 6, as the case may be. This is the vector I call L in each case. Of course, I came up with this easy scheme only after a lot of trial and error.

C. Details of Step 2. The angle $\vartheta(V_1, V_2)$ between two vectors V_1 and V_2 is

$$\vartheta(V_1, V_2) = \arccos\left(\frac{V_1 \cdot V_2}{\sqrt{(V_1 \cdot V_1)(V_2 \cdot V_2)}}\right) \tag{7}$$

Assuming that our 8 points for Ω are given as $P_0, ..., P_7$, each cone angle θ_i is a 6-fold sum of expressions like this. For instance, inspecting Figure 6 we see that

$$\theta_1 = \vartheta_{124} + \vartheta_{143} + \vartheta_{130} + \vartheta_{106} + \vartheta_{167} + \vartheta_{172}, \qquad \vartheta_{ijk} = \vartheta(P_j - P_i, P_k - P_i). \tag{8}$$

Plugging our torus Ω into Mathematica, which is perfectly capable of computing Equation 8 to high enough precision, I find that $|\theta_i - 2\pi| < 1.28 \times 10^{-32}$ for $i = 1, 2, 3$.

D. An Effective Inverse Function Theorem. Before launching into the details of Step 3, let me discuss some features of equation solving. Suppose $A : \mathbb{R}^3 \to \mathbb{R}^3$ is a linear transformation. We are interested in solving the equation $Av = w$ for some vector w. The classic guarantee that this is possible is that A is an invertible matrix. In this case our solution is simply $v = A^{-1}(w)$. When A is not invertible, all bets are off. Geometrically, a non-invertible A maps $\mathbb{R}^3$ into a lower dimensional subspace. If w does not lie in this subspace, we are out of luck.

Now consider an intermediate situation where A is invertible but the inverse A^{-1} is enormous. What is happening geometrically is that A is massively shrinking some directions down. If we choose w just in the wrong direction then the solution $A^{-1}(w)$ might be enormous. In other words, we might have to look *far and wide* for the desired solution. In linear algebra, such a matrix is often called *badly conditioned*. In contrast, when A^{-1} is small, we know that our solution $A^{-1}(w)$ will not be too far from the origin.

The *Inverse Function Theorem*, a famous result from real analysis, gives results like the invertibility result above for nonlinear maps like our coords-to-angles map Θ. The *Jacobian* of Θ at a point $p \in \mathbb{R}^3$ is the 3×3 matrix $d\Theta_p = (\partial\theta_i/\partial z_j)$. The Inverse Function Theorem says that if $d\Theta_p$ is an invertible matrix then, at least on a small ball centered at $q = \Theta(p)$, we can invert Θ. That is, we can solve the equation $\Theta(p^*) = q^*$ as long as q^* is sufficiently close to q. Let us sketch the proof of an effective version of this result which serves our purposes.

Let B be a ball of radius ρ centered at some point p. Let B' be a ball of radius $\lambda\rho$ centered at $\Theta(p)$. Suppose we know two things for any unit vector V:

$$\|d\Theta_q(V)\| > 2\lambda, \qquad \angle d\Theta_p(V), d\Theta_q(V) < \frac{\pi}{3}. \tag{9}$$

The second equation is about the angles between these vectors. Under these conditions we have $B' \subset \Theta(B)$. Here is a proof. Let $r \in \partial B$ and let γ be the line segment connecting p to r. Then $\Theta(\gamma)$ is a curve whose tangent vectors do not vary by more than $\pi/3$ from the initial tangent vector at $\Theta(p)$ and the arc length of $\Theta(\gamma)$ is at least 2λ times the arc length of γ. From these properties we deduce that the distance between $\Theta(p)$ and $\Theta(r)$ exceeds $\rho\lambda$. In other words, $\Theta(\partial B)$ is *entirely outside* B'. Moreover, given that the paths $\Theta(\gamma)$ are all pretty straight, the $\Theta(\partial B)$ "wraps around" B' much in the way that a balloon might enclose a baseball, and this implies that $B' \subset \Theta(B)$. To say this more formally, the straightness properties of $\Theta(\gamma)$, considered for all such paths, imply that the map $\Theta : \partial B \to \mathbf{R}^3$ is homotopic to the identity map in $\mathbf{R}^3 - B'$ and this in turn implies that $\Theta(\partial B)$ represents the generator of the homology group $H_2(\mathbf{R}^3 - B')$. This homological property is the formal way of saying that $\Theta(\partial B)$ encloses B' the way a balloon might enclose a baseball, and it implies that $B' \subset \Theta(B)$.

E. Details of Step 3. Even though Equation 7 involves a transcendental function, the derivative of this transcendental function is a rational function. This means that the entries of the matrix $d\Theta$ are rational functions of parameters z_1, z_2, z_3. They are messy, but we can work them out explicitly.

A direct calculation shows that where

$$\|d\Theta_p - M\|_\infty < \frac{1}{100}, \qquad M = \begin{bmatrix} -1.66 & +1.46 & -0.02 \\ +1.46 & +0.04 & -1.13 \\ -0.02 & -1.13 & +1.91 \end{bmatrix} \quad [10]$$

What we mean is that each entry of $d\Theta_p - M$ has absolute value less than $1/100$. In other words, M is the 2-digit truncation of $d\Theta_p$. The matrix M is symmetric and all its eigenvalues exceed $3/10$ in absolute value. Hence M does not shrink any vector by a factor smaller than $3/10$. This means that M^{-1} does not expand any vector by a factor of more than $10/3$. The matrix M is very well conditioned indeed.

We are not precisely interested in M, but rather in $d\Theta_q$ for $q \in B$, the ball of radius 10^{-12} around p. So, we need some derivative bounds to control how $d\Theta_q$ differs from both $d\Theta_p$ and from M.

Using the formulas directly (but in an efficient way) I establish the following bound:

$$\left| \frac{\partial^2 \theta_k}{\partial z_i \partial z_j} \right| < 10^9, \quad [11]$$

throughout B and for all indices i, j, k. Integrating this crude bound on the path of length at most 10^{-12} connecting p to $q \in B$, I get the following bound for all $q \in B$.

$$\|d\Theta_p - d\Theta_q\|_\infty < \frac{1}{450}. \quad [12]$$

Combining Equations 10 and 12 with the properties of M, I derive the following two crude statements, which hold for any unit vector V and any $q \in B$:

$$\|d\Theta_q(V)\| > \frac{1}{50}, \qquad \angle d\Theta_p(V), d\Theta_q(V) < \frac{\pi}{3}. \quad [13]$$

Equation 13 combines with our effective version of the Inverse Function Theorem to show that $B' \subset \Theta(B)$ for our choices of B and B'.

5. Discussion

I have proved that there does not exist a 7-vertex paper torus and I have also sketched the proof that there does exist an 8-vertex paper torus. These results settle the minimum-vertex question for paper tori. Put more colorfully, these results establish the most efficient possible way to make an embedded torus out of origami. Some readers might like to actually make a pup tent. My paper (1) has a link to a template you can copy and, with effort, fold up into a pup tent. I have to admit that I cannot successfully fold my own template, but my origami-skilled friends can do it easily.

It is worth commenting on the robustness of the proof sketched in the 8-vertex case, because otherwise all these orders of magnitude might leave the reader with the idea that the argument is somehow finely balanced and hanging by a thread. In fact, we are trying to hit a barn door with a shotgun from five feet away. I have been working with 32 digit precision for convenience and ease of presentation. I might have worked with an example having a million digits of precision. (Just for fun, Peter Doyle refined my numerical example to 1 million digits of precision.) Working with an example that was much closer to flat, I could get away with much looser technical estimates. For instance, I could work within a tiny ball of radius $10^{-100000}$, skate by with a bound of 10^{10000} in Equation 11, and replace the angle $\pi/3$ above with 10^{-1000}. In other words, I can remove any subtlety at all by working with a very high precision example.

It would be nice to have an explicit description of an 8-vertex paper torus rather than an existence proof based on an implicit function theorem. For theoretical reasons, the coordinates z_1', z_2', z_3' are roots of algebraic equations. Thus, they are explicit in some sense. However, the degrees of these equations seem to be so high that finding the explicit equations is too difficult. On a positive note, my implicit proof has some other virtues. Given the nature of the proof, it still works if we perturb the X and Y coordinates slightly. Thus, for each nearby choice of X-coordinates and Y-coordinates, subject to the imposed symmetry, we get an 8-vertex paper torus. This is to say that we really have an 8-parameter family of 8-vertex paper tori. Up to similarities of $\mathbf{R}^3$. there is a 6-parameter family of inequivalent examples. This observation opens the door to studying this parameter space in detail.

ACKNOWLEDGMENTS. I thank Samuel Leliévre and Alba Malaga for telling me all about flat tori (over a period of some years) and in particular telling me about the minimum vertex 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, Fabian Lander, Noah Montgomery, Stepan Paul, and Saul Schleimer for helpful discussions about this material. I thank Frank H. Lutz for making his manifolds webpage, which is both inspiring and useful. I thank the Institut des Hautes Études Scienfiques, the Hamilton Institute, and the Isaac Newton Institute, where I started working on this paper. This work is supported by the National Science Foundation, a Simons Sabbatical Fellowship, and the Mazur Chair at IHES.

1. RE Schwartz, Vertex-minimal paper tori (2025) arXiv:2507.14998 [math.MG].
2. NH Kuiper, On c^1-isometric imbeddings i. *Indagationes Math. (Proceedings)* **58**, 545–556 (1955).
3. J Nash, C^1 isometric imbeddings. *Annals Math.* **60**, 383–396 (1954).
4. V Borrelli, S Jabrane, F Lazarus, B Thibert, Flat tori in three-dimensional space and convex integration. *Proc. Natl. Acad. Sci.* **109**, 7218–7223 (2012).
5. YD Burago, VA Zalgaller, Polyhedral realizations of developments. *Vestnik Leningr. Univ.* **15**, 66–80 (1960) In Russian.
6. YD Burago, VA Zalgaller, Isometric embeddings of two-dimensional manifolds with a polyhedral metric into $\mathbb{R}^3$. *Algebr. i Analiz* **7**, 76–95 (1995) Translation in St. Petersburg Mathematical Journal 7(3):369–385.
7. U Brehm, Oberwolfach report (1978).
8. T Tsuboi, On origami embeddings of flat tori (arXiv preprint arXiv:2007.03434) (2020).
9. P Arnoux, S Lelièvre, A Malaga, Diplotori: a family of polyhedral flat tori. In preparation (year?).
10. F Lazarus, F Tallerie, A universal triangulation for flat tori (arXiv preprint arXiv:2203.05496) (2024).
11. Á Császár, A polyhedron without diagonals. *Acta Sci. Math. (Szeged)* **13**, 140–142 (1949).
12. PT Quintanar Cortés, "Plongements polyédriques du tore carré plat," PhD thesis, Université Claude Bernard Lyon 1 (2019). URL http://www.theses.fr/2019LYSE1354.
13. J Bokowski, A Eggert, All realizations of möbius' torus with 7 vertices. *Topol. Struct.* **17** (1991).
14. LA Santaló, *Integral Geometry and Geometric Probability.* (Addison–Wesley, Reading, MA), (1976) Reprint of 1953 Spanish original.
15. FH Lutz, The Manifold Page (https://www3.math.tu-berlin.de/IfM/Nachrufe/Frank_Lutz/stellar/) (2017) Last modified April 26, 2017.
16. I Wolfram Research, Mathematica, version 14.3 (year?) Champaign, IL, 2025.