

Parameterizing N -holed Tori

Cindy Grimm¹ and John Hughes²

¹ Washington University in St. Louis, Box 1045,
St. Louis, MO
cmg@cs.wustl.edu

<http://www.cs.wustl.edu/cm>

² Brown University, Box 1910,
Providence, RI

jfh@cs.brown.edu <http://www.cs.brown.edu/jfh>

Abstract. We define a parameterization for an n -holed tori based on the hyperbolic polygon. We model the domain using a manifold with $2n + 2$ charts, and linear fractional transformations for transition functions. We embed the manifold using standard spline techniques to produce a surface.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling, Curve, Surface, Solid, and Object Representations, Splines, n -holed tori, hyperbolic octagon, linear fractional transformation

1 Introduction

We present a method for constructing n -holed tori for use in computer graphics. The requirements are that the construction be *explicit*: it is not sufficient to know that we have some description of a manifold known to be homeomorphic to an n -holed tori. Instead, we need explicit charts in \mathbb{R}^2 and transition functions that are easily computable.

We take a two step approach: First we take a classical description of an n -holed torus and make that description explicit. Next, we use traditional planar embedding techniques, such as splines, to create an embedding of the n -holed torus. We differ from previous approaches [RP99,WP97] in that we cover the domain with $2n + 2$ charts, instead of modeling the domain as a single entity. The embedding is then built by blending between $2n + 2$ individual, planar chart embeddings, rather than making a single, multiperiodic embedding function.

Our construction method is of interest from both a mathematical standpoint and an implementation one. Modeling arbitrary topology surfaces, especially ones with higher-order continuity, continues to be a challenging problem. The solution presented here is a computationally tractable natural parameterization of n -holed tori, and has many of the desirable features of traditional modeling techniques such as splines.

We begin with a global parameterization technique from topology [Lef49] which takes a $4n$ -sided polygon and produces an n -holed torus by associating

edges of the polygon. By recognizing (following classical complex analysis) that this polygon can be embedded in the unit disk so that the associations on the edges are those induced by a group action on the disk, we see immediately that the object produced is indeed a complex manifold: our manifold structure, with $2n + 2$ charts, has transition functions that are all linear fractional transformations (see Section 4).

Once we have the explicit description of the manifold, we embed it into \mathbb{R}^3 using standard, planar embedding techniques (see Section 5).

2 Related work

The modeling of arbitrary topology surfaces has received a great deal of attention. The approaches range from parameterized ones such as hole filling with n -sided patches [Ld89,Pet96], subdivision surfaces [CC78], and manifolds [GH95,Gri02,NG00], to implicit methods [Blo97] and other volumetric approaches [CBC⁺01,FPRJ00]. n -holed tori are easy to build using implicit methods, but they lack a parameterization.

Research in the area of hole filling and networks of patches is extensive, and beyond the scope of this paper to cover in detail. The basic approach is to apply constraints to the geometry of neighboring patches in order to ensure their continuity. These approaches are general, in that the approach works for any valid network, and hence any topology that can be modeled by that network. We differ from these approaches in that we specify the construction for each genus and do not rely on constraints to maintain continuity.

Manifolds and subdivision surfaces have both been used to build smooth, arbitrary topology surfaces without the use of constraints. Our approach differs from these in that our parameterization has few charts and reflects the standard structure for the n -holed surface, with no embedding-induced bias. Subdivision surfaces have proven to be very useful to the modeling community, but they do not easily extend to higher order continuities, nor are they parameterized, although the characteristic map can be used to derive a local parameterization [DKT98].

The mathematical foundations for hyperbolic-plane constructions of surfaces go back to the 19th century and the work of Poincaré and Klein; the topological basics are well-described in many elementary topology books [Mas77]. These ideas were first used by Ferguson and Rockwood [FR93] to build smooth surfaces with handles. Their basic approach begins with a function f over the hyperbolic disk which does not agree at the boundaries. By applying a *multiperiodic basis* to f , they create a function f' which does agree. The original paper created the function f using a very specific construction process involving Coons patches arranged symmetrically in the domain; this process was replaced by a more general approach using radial basis functions in a later paper [RP99]. In this second approach, the function f is built from a set of scattered points in the domain; the corresponding point locations in \mathbb{R}^3 define the function f . An embedded Delaunay triangulation of the points provides a control polygon.

Spline orbifolds [WP97] are also built on the hyperbolic plane and use a spline-like embedding function. To create a tessellation they begin with a “sketch” polyhedron and use it to create a group-invariant triangulation of the domain. This sketch polyhedron is also used in a least-squares fitting approach to set the control points of the spline function.

We differ from these previous approaches in our embedding and tessellation.

3 Manifold representation, group actions, and a fundamental theorem

The standard definition of a manifold begins “Let X be a Hausdorff space, and suppose that $\{V_i\}$ is an open cover of X , and that $\phi_i : V_i \rightarrow U_i \subset \mathbb{R}^k$ is a homeomorphism....” From this definition, one builds “transition functions” and eventually completes the structure of a manifold.

Evidently the disjoint union Y of all the U_i in this standard definition have an equivalence relation on them: $p_i \in U_i$ is equivalent to $p_j \in U_j$ if and only if $\phi_i^{-1}(p_i) = \phi_j^{-1}(p_j)$. Indeed, we typically define “transition functions” $\psi_{ij} : U_{ij} \rightarrow U_{ji}$, where U_{ij} is defined as $\phi_i(V_j)$, and ψ_{ij} is just $\phi_j \circ \phi_i^{-1}$. The equivalence relation then becomes $p_i \sim \psi_{ij}(p_i)$. The manifold X is evidently in one-to-one correspondence with the quotient of Y by this equivalence relation: the point $p \in X$ corresponds to the equivalence class of $\phi_i(p)$, where $p \in V_i$.

In the event that one does not have the space X to start with, one can build a collection of open sets U_i and a collection of functions ψ_{ij} on appropriate subsets of them and then consider quotient of their disjoint union by the induced equivalence relation. Under suitable conditions, the quotient will be a manifold. This is the approach taken by Grimm and Hughes [GH95]; its advantage is that the description provided is particularly amenable to implementation. We will continue to use their notation, but in this particular case, we will know *a priori* that the space being constructed is a manifold. We will begin with the general notation and some mathematical preliminaries, and then, in Section 4.1, describe the n -holed torus manifold explicitly.

The components of the Grimm and Hughes description of a (two-dimensional) manifold are:

- A finite set, A , of n nonempty subsets $\{c_1, \dots, c_n\}$ of \mathbb{R}^2 . A is called an *atlas*. Each element $c_i \in A$ is called a *chart*.
- A set of subsets, $U_{ij} \subset c_i$ ($j = 1, \dots, n$); the subset U_{ii} must be all of c_i . These regions act as the “overlap regions” for the manifold: we will see that points of U_{ij} and U_{ji} get identified with one another. Note that U_{ij} need not be a connected set; nor need it be nonempty, except in the case $i = j$.
- A set $\Psi = \{\psi_{ij} | i, j = 1, \dots, n\}$ of n^2 functions called *transition functions*. For each (i, j) , the map $\psi_{ij} : U_{ij} \rightarrow U_{ji}$, where $U_{ij} \subset c_i$ and $U_{ji} \subset c_j$, must be smooth of some order k , and satisfy the requirements listed below. Note that because U_{ij} may not be connected, ψ_{ij} may be described by a set of functions, one for each connected component of U_{ij} .

As noted above, there is a relation \sim defined on $Y = \sqcup_{c \in AC} c$ (where \sqcup denotes disjoint union) such that if $x \in c_i$, $y \in c_j$, then $x \sim y$ iff $\psi_{ij}(x) = y$. We require that the transition functions be symmetric ($\psi_{ij} = \psi_{ji}^{-1}$), that ψ_{ii} is the identity for all i , and that they satisfy the *cocycle condition*, i.e., that $\psi_{ij} \circ \psi_{ki} = \psi_{kj}$ wherever this makes sense. These requirements ensure that the relation \sim is an equivalence relation [Gri96]. The quotient of Y by \sim is then (under certain technical conditions³) guaranteed to be a manifold of class C^k .

Note that a “point” in this quotient manifold consists of a list of all of the chart points that are equivalent under \sim , i.e., given a point p_i in a chart c_i , the corresponding manifold point is all of the tuples (c_j, p_j) such that $p_i \sim p_j$.

3.1 Group actions on manifolds

A group G is said to act on a manifold M if for each $g \in G$, there’s an associated homeomorphism $f_g : M \rightarrow M$ with the following properties:

- If e is the identity element of G , then f_e is the identity on M .
- If g and h are two elements of G , and $k = gh$, then $f_k = f_g \circ f_h$.

If the maps f_g are diffeomorphism, we say that “ G acts by diffeomorphisms on M .” If M has a metric structure and the maps are isometries, we say that G acts by isometries on M .”

Such a group action defines an equivalence relation: two points m_1 and m_2 of M are equivalent if there’s an element g of G such that $f_g(m_1) = m_2$.

As an example, the group of integers acts on the real line by translation: $f_n : \mathbb{R} \rightarrow \mathbb{R} : x \mapsto x + n$. The set of equivalence classes for this action is in one-to-one correspondence with the interval $[0, 1]$, except that the equivalence classes for 0 and 1 are the same; in other words, the quotient space is a circle. (The topology on the quotient is that of the circle as well, but this requires careful working out of details beyond the scope of this paper.) The interval $[0, 1]$ is called a *fundamental domain* for the group action: it is a subset K of the manifold with the property that $f_g(K) \cap K$ is entirely within the boundary of K (for every g in G), and such that $M = \cup_{g \in G} f_g(K)$.

The situation exemplified in the previous paragraph is not uncommon: frequently the quotient of a manifold by a group action is another manifold.

One general result [CDGM00] is all that we will need in this paper:

Theorem 1. *Suppose G is a discrete group acting by isometries on a smooth manifold M , and for every compact subset K of M , the set*

$$G(K; K) = \{g \in G \mid gK \cap K \neq \emptyset\}$$

³ The technical conditions have to do with being Hausdorff: consider two copies of the real line, U_1 and U_2 , and the maps $\psi_{12} : U_1 - \{0\} \rightarrow U_2 - \{0\} : x \mapsto x$, $\psi_{21} = \psi_{12}^{-1}$, $\psi_{11} = id_R$ and $\psi_{22} = id_R$. These satisfy all the requirements of the description above, but the quotient space, “the line with two origins,” is not actually a manifold, since the two copies of the point “0” cannot be separated by open sets.

is finite. Further suppose that the only element of G that fixes a point of M is the identity. Then M/G is a smooth manifold as well, and the map $M \rightarrow M/G$ is a local diffeomorphism.

We will apply this to the case where M is the open unit disk in the complex plane, and G is a group of fractional linear transformations on the disk. A fractional linear transformation can be represented by a matrix:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix},$$

where the numbers a, b, c , and d are complex. The associated transformation takes the complex number z to $\frac{az+b}{cz+d}$. One can easily check that the matrix:

$$Q = \begin{bmatrix} -1 & i \\ 1 & i \end{bmatrix}$$

takes the upper-half plane (the set of complex numbers with positive imaginary part) to the interior of the unit disk in the complex plane.

Furthermore, the matrices with *real* entries and determinant one (called $SL_2(\mathbb{R})$) send the complex plane one-to-one onto itself. That means that the matrices QMQ^{-1} (for $M \in SL_2(\mathbb{R})$) send the unit disk one-to-one onto itself, so we can say that $SL_2(\mathbb{R})$ acts on the unit disk as well.

Serre [Ser90] describes this action and its properties in detail. In this paper, we will build a discrete subgroup of $SL_2(\mathbb{R})$; the quotient of the disk by this action will be an n -holed torus. This will actually be evident by the construction: a fundamental domain for the action will be a regular $4n$ -gon in the hyperbolic disk, and the group action will be one that identifies edges with edges in the way described in elementary topology books. The fact that $G(K;K)$ is finite and that the only transformation fixing a point is the identity will be obvious from the construction. Hence the theorem above tells us that the resulting space is, in fact, a manifold as required.

4 Building a manifold

From topology [Lef49], we know that an n -holed torus can be built from a $4n$ -sided polygon by associating the edges of the polygon as shown in Figure 1. We use this structure to determine how many charts to use and how they overlap. We also use it to define the transition functions, by first mapping from chart c_i to the polygon, then from the polygon to chart c_j . This two-step process is used to illustrate how the transition functions are built; the actual transition functions map directly from one chart to the other.

We first recall the description of how an n -holed torus is constructed from a $4n$ -sided polygon and how we cover the polygon with $2n+2$ charts (Section 4.1). In Section 4.2 we describe the hyperbolic $4n$ -gon. The charts and maps to and from the polygon are defined in Section 4.3.




Fig. 1. An abstract $4n$ -sided polygon with the associated edges labeled.

4.1 Building the N -holed Tori

An n -holed torus can be built by associating the edges of a $4n$ -sided polygon; the two-holed torus case is shown in Figure 2. Some observations about this construction:

- Each hole is represented by a group of four consecutive edges.
- The first pair of associated edges (a, a^{-1}) correspond to a loop that goes around the hole. The second pair of associated edges (b, b^{-1}) corresponds to a loop that goes through the hole (once the torus is embedded in 3-space by our particular embedding; some other embedding might swap the roles.)
- All of the vertices of the polygon correspond to a single point on the final surface. Each loop begins and ends at this point, in the order shown in Figure 2. Note that this is **not** the same order as the corners in the polygon.

We use $2n + 2$ charts to cover the polygon (see Figure 4). The first chart, termed the “inside” chart, is a disk that covers much of the interior of the polygon. The second chart, termed the “vertex” chart, is also a disk. It is centered on the polygon vertices and covers the corners of the polygon. The remaining n charts are termed “edge” charts and are unit squares placed so that the midline of the chart covers most of the edge loop.

The vertex chart will cover the $4n$ corners of the polygon; this implies that each $2\pi/(4n)$ wedge of the vertex chart should map to its corresponding corner of the polygon. This mapping is greatly simplified if the corners of the polygon each have an angle of $2\pi/(4n)$, which is clearly not possible with Euclidean geometry. We therefore use a hyperbolic polygon.




Fig. 2. Left: An 8-sided polygon with edges and vertex corners labeled. Right: A sketch of a 2-holed torus with the loops and vertex corners labeled.

4.2 Hyperbolic polygon

In Figure 3 we see $4n$ circles arranged evenly around a unit disk. Each circle meets the unit circle orthogonally. Neighboring pairs of circles intersect at an angle of $2\pi/(4n)$. The part of the unit disk that is not inside any of the circles is therefore a regular $4n$ -gon (with curved edges) and with vertex angles each $2\pi/(4n)$.

A brief digression about this polygon is in order: the unit disk is one of several models for hyperbolic geometry. In this model, the geodesics (“straight lines”) are (a) the diameters of the disk, and (b) circle-arcs that meet the boundary orthogonally; the isometries in this geometry are fractional linear transformations that map the disk to the disk. Therefore, what we have constructed is a regular $4n$ -gon in the hyperbolic plane; the edges of this $4n$ -gon, although curved from a Euclidean perspective, are straight lines in the hyperbolic perspective. Just as one can take an equilateral triangle in the Euclidean plane and flip it over one edge, then over another, and so on, and thus “fill out” the entire plane with a tiling by equilateral triangles, one can take this hyperbolic $4n$ -gon and “flip” it over its edges repeatedly to fill out the entire hyperbolic disk. The “flips” are not reflections in the Euclidean sense, but rather are fractional linear transformations that map the disk to itself, leaving the “flip edge” fixed.

To return to the construction: the small circles are defined by their distance from the origin (h) and their radius (r). Their angular spacing is:

$$\omega = \frac{2\pi}{4n} \quad (1)$$




Fig. 3. The hyperbolic polygon for the two-holed torus, labeled with α (corner angle), ω ($2\pi/(4n)$), β (edge-chart interior corner), h , r , and γ (length to polygon boundary).

The constraint that the circles meet the disk-boundary orthogonally introduces the following relationship between h and r :

$$r = \sqrt{h^2 - 1} \quad (2)$$

By changing r (and consequently h) we can create nearly any desired angle between the intersecting circles. The choice of h that makes the angle ω is:

$$h = \sqrt{\frac{\cos(\omega) + 1}{\cos(\omega) + 1 - 2\sin^2(\omega/2)}} \quad (3)$$

We will also need the distance to the vertex (α) and to the middle of the edge (γ).

$$\alpha = h \cos(\omega/2) - \sqrt{(h \cos(\omega/2))^2 - 1} \quad (4)$$

$$\gamma = h - r \quad (5)$$

Finally, we need the point β , which will be used to set the vertex chart size and in the tessellation (Section 5.2).

The derivations of these equations and of β are given in Appendix A.

4.3 Mapping to the polygon

As mentioned above, the isometries of the hyperbolic plane model are all linear fractional transformations (LFTs). We will be describing our transition functions in terms of these.




Fig. 4. From left to right: The inside chart, an example edge chart, and the vertex chart. The edge chart is split in two, with the left half mapped to the lower left edge, the right half mapped to the lower right edge. Each wedge of the vertex chart is mapped to a different corner.

If T and S are LFTs with matrices K and L , then the transformation with matrix KL is $T \circ S$. If K is the matrix $\begin{bmatrix} a & b \\ c & d \end{bmatrix}$ then the matrix

$$\begin{bmatrix} d & -b \\ -c & a \end{bmatrix} \tag{6}$$

represents the transformation T^{-1} .

We use two types of LFT. The first rotates and scales, the second translates:

$$S(p) = \begin{bmatrix} p & 0 \\ 0 & 1 \end{bmatrix} \quad T(p) = \begin{bmatrix} 1 & p \\ p & 1 \end{bmatrix} \tag{7}$$

The domain for the inside and vertex charts is the disk given in polar coordinates by $0 \leq r \leq 0.5$, i.e., the set of complex numbers $\{r \cos \theta + \mathbf{i}r \sin \theta | 0 \leq r < \frac{1}{2}, \theta \in [0, 2\pi]\}$. The domain for the edge charts is a square centered at the origin, i.e., $\{x + \mathbf{i}y | -\frac{1}{2} < x, y < \frac{1}{2}\}$.

The inside chart is mapped to the polygon by scaling the disk until it almost touches the edges of the polygon:

$$S(\gamma)(a + \mathbf{b}\mathbf{i}) \tag{8}$$

The edge charts are mapped so that the chart is split vertically, with the left half mapping to the part of the disk near edge a , and the right half mapping to edge a^{-1} (see Figure 4). Each piece of the chart is rotated, then translated out to the matching edge. The right half chart is rotated by π before translating, so as to orient the mid-line in the opposite direction. The edge charts are first

scaled by $7/8$ so they cover most, but not all, of the edge (the remainder will be covered by the vertex chart). Let j be the chart number (going counter-clockwise around the circle). Recall that the edges come in groups of four, $aba^{-1}b^{-1}$, so that the a charts overlap edges $4j$ and $4j + 2$, while the b charts overlap edges $4j + 1$ and $4j + 3$.

$$p_l = \cos((4j + j \bmod 2)\omega) + i \sin((4j + j \bmod 2)\omega) \quad (9)$$

$$p_r = \cos((4j + 2 + j \bmod 2)\omega) + i \sin((4j + 2 + j \bmod 2)\omega) \quad (10)$$

$$p_f = \cos \pi + \mathbf{i} \sin \pi \quad (11)$$

$$S(p_l)T(h - r)S(7/8)(a + \mathbf{bi}) \quad a < 0.5 \quad (12)$$

$$S(p_r)T(h - r)S(7/8)S(p_f)(a + \mathbf{bi}) \quad a \geq 0.5 \quad (13)$$

The vertex chart is mapped to the polygon with $4n$ maps, one for each vertex of the polygon. Note that the order of the wedges in the chart is **not** the same as the order of the polygon corners. The order can be determined by following the associated edges around the polygon (see Figures 2 and 4). The mapping first rotates the chart by (R) , then translates it to the appropriate corner C so that the correct chart wedge is facing inwards. The vertex chart is scaled by δ so that the boundary passes through β (see Appendix A). The general form of the mapping is:

$$p_c = \cos((C + 1/2)\omega) + \mathbf{i} \sin((C + 1/2)\omega) \quad (14)$$

$$p_r = \cos(-R\omega) + \mathbf{i} \sin(-R\omega) \quad (15)$$

$$S(p_c)T(\alpha)S(p_r)S(\delta) \quad (16)$$

The wedges in the vertex chart are indexed clockwise, with the first wedge centered on the negative real axis. The polygon corners are indexed counter-clockwise, with the first corner at $\omega/2$ (or $C = 0$). As we walk around the wedges in the vertex chart (incrementing R by 1) then C is updated by:

$$C = (C - 3) \bmod 4n \quad (C \bmod 4) = 1, 2$$

$$C = (C + 1) \bmod 4n \quad (C \bmod 4) = 0, 3$$

In other words, we either move to the next corner, or back three.

4.4 Alternative description

To describe the charts differently, we begin by describing a group of isometries of the disk. For each identified edge-pair (e.g., a and a^{-1}) there is an orientation-preserving isometry of the disk that carries one edge to the other. In the case of the octagon, this map is a rotation by $\pi/2$ about the origin, followed by a (hyperbolic!) rotation of $\pi/2$ about the midpoint of the edge labelled a^{-1} . Similar maps take edge b to edge b^{-1} , and so on. If one writes the maps out explicitly, it is clear that their only fixed points are on the unit circle, hence outside the domain we are considering. By taking all possible repeated compositions of these isometries, we get a tiling of the hyperbolic disk by hyperbolic octagons.

Each point of the disk is in exactly one of these polygons (or perhaps on the boundary of two or more of them). Thus we can define a chart that sends some domain into the disk, partly inside and partly outside of the fundamental octagon; the part that is outside the fundamental octagon corresponds (under one of the isometries) to some set of points that's *inside* another part of the octagon. In figure 4, we have drawn the charts this way, showing, for example, an edge chart having an image that lies partly outside the fundamental octagon. We have also drawn the edge-chart followed by an isometry, and the part that lay outside is now within, and vice-versa.

4.5 The transition functions




Fig. 5. The overlap regions U_{ij} for the different cases.

The transition functions are the composition of chart i 's map to the polygon with chart j 's inverse map. Some care must be taken to ensure that the correct pairs are combined, since each connected overlap region will have a different LFT. Refer to Figure 5:

- The inside chart overlaps with each edge chart in two regions, one in the left half and one in the right half.
- Each edge chart overlaps with two different edge charts. The overlap region has one, two, or three connected components each covering one of the corners.
- The vertex chart overlaps the inside chart in $4n$ places.
- The vertex chart overlaps each edge chart in two disconnected regions. Note that there are two possible constructions for these transition functions, but both yield the same LFT.

With the charts described and the transition functions described, we have all the components necessary for the Grimm-Hughes description of a manifold. The only question that remains is whether the resulting quotient object is, in fact, a manifold. In general, that requires proving the quotient is Hausdorff, but in this case the answer is simple: the group of isometries described above evidently acts without fixed points (except for the identity transformation), and because the regular $4n$ -gon is a fundamental region for the group action, it is clear that for any compact subset K of the disk, the set $G(K; K)$ of all transformations T such that $T(K)$ intersects K is finite. Hence the theorem from Section 3 applies, and the quotient space is actually a manifold, indeed, a complex one-dimensional manifold.

5 Embedding

In this section we describe how to embed our manifold in 3-space and how to tessellate the domain in order to produce a mesh for rendering.

5.1 Defining an embedding

To embed the manifold we first define a spline patch [BBB87] $E_c : C \rightarrow \mathbb{R}^3$ and a blend function $b_c : C \rightarrow \mathbb{R}$ for each chart C . The blend function and its derivatives must be zero by the boundary of the chart. The blend functions must also form a partition of unity; fortunately, C^k blend functions, for any desired k , can be built from spline basis functions [GH95]. The surface embedding is then:

$$S(p) = \sum_{c \in A} b_c(\alpha_c(p)) E_c(\alpha_c(p)) \quad (17)$$

where α_c is the point corresponding to p in chart c , if there is one. If p does not correspond to a point for chart c then b_c is defined to be zero. (Note that “corresponding to” in this case is simple: a point p in the manifold is an equivalence class of points in the charts, i.e., a list of $2n + 2$ or fewer chart-points.)

For the edge charts we use the tensor product of two spline basis functions, each of which has a support of $[0, 1]$. For the vertex and inside charts we use the distance from the chart center and a single spline basis function whose support is the diameter of the chart. This produces a radially symmetric blend function.




Fig. 6. Example embeddings. The mesh in the upper left hand corner is the n -holed mesh we fit to initially, to produce the surface in the lower left hand corner. The right two surfaces were created by editing this initial surface.

If the blend and embedding functions are C^k continuous then the resulting surface is C^k . To produce a visually pleasing surface it is best if the spline patches agree where they overlap, *i.e.*, $E_i(\alpha_i(p)) = E_j(\alpha_j(p))$; this ensures that the blending above is purely a formality.

To create our initial embedding we first build an n -holed mesh by taking a sphere and punching n holes in it. We then split this mesh apart along boundary lines such as the ones shown in Figure 2. We flatten this mesh into the hyperbolic polygon using Floater’s [Flo97] to produce a bijection between the mesh vertices and the points on the manifold. We then solve for control points for each patch in a least-squares [FB91] fashion:

$$b_c(\alpha_c(p_i))E_c(\alpha_c(p_i)) = P_i \quad (18)$$

Each mesh vertex produces one linear constraint on some subset of the spline patches.

Once the initial surface is created we let the user edit it by direct manipulation. The user clicks on a point of the surface, which selects a point on the manifold. When the user moves the mouse we solve for the smallest change in the control points that moves the surface point to the new location (again using equation [FB91]). Figure 6 shows some an example mesh and initial surface for the two-holed case, and two edited surfaces.

5.2 Tessellation

We tessellate the domain in order to produce a mesh to render. The tessellation has two parameters, t_e and t_r , which control the resolution of the mesh. The

coarsest level tessellation ($t_e = 1, t_r = 1$) partitions the domain into $2n + 2$ regions, one for each chart (see Figure 7).

The lowest resolution tessellation places a square in each edge chart. The size of this square is chosen so that the corners of the squares of neighboring edge charts map to the same point in the polygon, which is the point β in Figure 3. The edges of these squares form $4n$ consecutive arcs in the vertex and inside charts, enclosing a disk region in the center of the chart.

The first tessellation parameter, t_e , specifies the number of divisions in the square. This produces a grid inside each edge chart, and a wheel with the inside and vertex charts. The second parameter, t_r , specifies the number of divisions along the spokes of the wheel.




Fig. 7. Upper left: The coarsest tessellation. Upper right: Subdividing the tessellation. Bottom: The tessellation in the interior of each chart type.

6 Conclusion

We have presented a technique for modeling n -holed tori, based upon a natural parameterization for n -holed tori, suitable for implementation in computer graphics.

Several areas remain open, however. The group of fractional linear transformations that we have used can probably be adjusted to become a subgroup of the modular group (2×2 matrices with integer entries and determinant one); this would result in a description of the n -holed torus as a natural branched cover of the unit sphere S^2 . Meromorphic functions on these n -holed tori might well be used to define maps into the k -fold product $(S^2)^k \subset (\mathbb{R}^3)^k$, which might then be projected to give particularly nice maps from the n -holed torus to \mathbb{R}^3 , rather than the somewhat ad-hoc spline-based embeddings we have described. We hope to pursue this in the future.

A Derivations

To derive $r = \sqrt{h^2 - 1}$ we note that the triangle formed by the center of the two circles and their intersection point is a right-angled one (by construction). Therefore $r^2 + 1^2 = h^2$.

Let $\omega = \cos(\omega) + \mathbf{i} \sin(\omega)$ be the unit vector from the origin through the point where the first two circles intersect. Let $\alpha\omega$ be the intersection point. Then

$$\cos(\omega/2) = \omega \cdot \frac{(\alpha\omega - h)^\perp}{r} \quad (19)$$

If z and w are unit vectors, then their dot product is $\operatorname{Re}(z\bar{w})$. Applying this, and $z^\perp = \mathbf{i}z$, we have:

$$\cos(\omega/2) = \operatorname{Re}(\omega \mathbf{i} \frac{(\alpha\omega - h)^\perp}{r}) \quad (20)$$

$$= \pm \frac{h}{r} (\sin \omega) \quad (21)$$

Applying $\cos(2x) = 2 \cos^2 x - 1$:

$$\cos(\omega) = \frac{2h^2 + \sin^2 \omega + 1}{h^2 - 1} \quad (22)$$

Solving for h we get:

$$h = \sqrt{\frac{\cos(\omega) + 1}{\cos(\omega) + 1 - 2 \sin^2(\omega/2)}} \quad (23)$$

To solve for β we note that we are looking for the number s such that:

$$m_0(-s + \mathbf{i}s) = m_1(-s + \mathbf{i} - s) \quad (24)$$

where m_0 and m_1 are the LFT that take the upper half of chart 0 and chart 1 to the polygon. This reduces to:

$$\frac{m_0 0, 0(-s + \mathbf{i}s) + m_0(0, 1)}{m_0(1, 0)(-s + \mathbf{i}s) + m_0(1, 1)} = \frac{m_1 0, 0(-s + \mathbf{i} - s) + m_1(0, 1)}{m_1(1, 0)(-s + \mathbf{i} - s) + m_1(1, 1)} \quad (25)$$

which reduces to a quadratic equation in s . Once we have s we can derive δ in a similar manner.

References

- [BBB87] R. Bartels, J. Beatty, and B. Barsky. *An Introduction to Splines for Use in Computer Graphics and Geometric Modeling*. Morgan Kaufmann, 1987.
- [Blo97] Jules Bloomenthal, editor. *Introduction to Implicit Surfaces*. Morgan Kaufmann, 1997.
- [CBC⁺01] Jonathan C. Carr, Richard K. Beatson, Jon B. Cherrie, Tim J. Mitchell, W. Richard Fright, Bruce C. McCallum, and Tim R. Evans. Reconstruction and Representation of 3D Objects With Radial Basis Functions. *Proceedings of ACM SIGGRAPH 2001*, pages 67–76, 2001.
- [CC78] Ed Catmull and J. Clark. Reursively generated B-spline surfaces on arbitrary topological meshes. *Computer Aided Design*, 10:350–355, September 1978.
- [CDGM00] Virginie Charette, Todd Drumm, William Goldman, and Maria Morrill. Complete Flat Affine and Lorentzian Manifolds. In *Proceedings of the Workshop “Crystallographic Groups and their Generalizations II,” Contemporary Mathematics 262*, pages 135–146. American Mathematical Society, 2000.
- [DKT98] Tony D. DeRose, Michael Kass, and Tien Truong. Subdivision Surfaces in Character Animation. *SIGGRAPH 98*, pages 85–94, 1998.
- [FB91] B. Fowler and R. H. Bartels. Constraint based curve manipulation. *Siggraph course notes 25*, July 1991.
- [Flo97] Michael S. Floater. Parametrization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14(3):231–250, 1997. ISSN 0167-8396.
- [FPRJ00] Sarah F. Frisken, Ronald N. Perry, Alyn P. Rockwood, and Thouis R. Jones. Adaptively Sampled Distance Fields: A General Representation of Shape for Computer Graphics. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, pages 249–254. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, July 2000. ISBN 1-58113-208-5.
- [FR93] H. Ferguson and A. Rockwood. Multiperiodic functions for surface design. *Computer Aided Geometric Design*, 10(3):315–328, August 1993.
- [GH95] Cindy Grimm and John Hughes. Modeling Surfaces of Arbitrary Topology using Manifolds. *Computer Graphics*, 29(2), July 1995. Proceedings of SIGGRAPH '95.
- [Gri96] C. Grimm. *Surfaces of Arbitrary Topology using Manifolds*. PhD thesis, Brown University (in progress), 1996.
- [Gri02] Cindy Grimm. Simple Manifolds for Surface Modeling and Parameterization. *Shape Modelling International*, May 2002.

- [Ld89] Charles Loop and Tony deRose. A multisided generalization of bezier surfaces. *ACM Transactions on Graphics*, 8(3):204–234, July 1989.
- [Lef49] S. Lefschetz. *Introduction to Topology*. Princeton University Press, Princeton, New Jersey, 1949.
- [Mas77] William S. Massey. *Algebraic Topology: An Introduction*. Springer Verlag, New York, 1977.
- [NG00] J. Cotrina Navau and N. Pla Garcia. Modeling surfaces from meshes of arbitrary topology. *Computer Aided Geometric Design*, 17(7):643–671, August 2000. ISSN 0167-8396.
- [Pet96] Jorg Peters. Curvature continuous spline surfaces over irregular meshes. *Computer-Aided Geometric Design*, 13(2):101–131, February 1996.
- [RP99] A. Rockwood and H. Park. Interactive design of smooth genus N objects using multiperiodic functions and applications. *International Journal of Shape Modeling*, 5:135–157, 1999.
- [Ser90] Jean Pierre Serre. *A Course in Arithmetic*. Springer Verlag, New York, 1990.
- [WP97] J. Wallner and H. Pottmann. Spline orbifolds. *Curves and Surfaces with Applications in CAGD*, pages 445–464, 1997.