

Geometry, Topology, and Applications of the Minkowski
Product and Action

by
Micah Smukler
Weiqing Gu, Advisor

Advisor: _____

Second Reader: _____

(Francis Su)

Department of Mathematics

HARVEY MUDD
COLLEGE

Abstract

Geometry, Topology, and Applications of the Minkowski

Product and Action

by Micah Smukler

The Minkowski product can be viewed as a higher-dimensional version of interval arithmetic. We discuss a collection of geometric constructions based on the Minkowski product and on one of its natural generalizations, the quaternion action. We also will present some topological facts about these products, and discuss the applications of these constructions to computer-aided geometric design.

Table of Contents

List of Figures	iii
Chapter 1: Introduction	1
Chapter 2: Useful Facts and Definitions	2
2.1 The quaternions	2
2.2 The Minkowski product	4
2.3 The quaternion action	5
2.4 Topological issues and examples	6
Chapter 3: Geometry and Topology of the Minkowski Product	9
3.1 The self-quotient	9
3.2 Results	10
Chapter 4: Curves with Singularities	14
4.1 Curves with endpoints	15
4.2 Curves with singularities	17
Chapter 5: Spline-action Surfaces	20
5.1 Background	20
5.2 Spline-action surfaces	21
5.3 Run-time comparison with tensor product surfaces	26
Chapter 6: Inversion and Generalized Actions	30

6.1	A review of inversion	30
6.2	Inverting the quaternion action	31
6.3	The generalized quaternion action	32
6.4	The universality of the generalized action	33
Chapter 7:	Quaternion Action and Inversion of Canal Surfaces	37
7.1	Inversion of a particular circle	37
7.2	Inversion of canal surfaces and specific examples	40
7.3	The virtual inversion	42
Chapter 8:	Future Work	44
8.1	New constructions	44
8.2	Refining current constructions	45
Chapter 9:	Conclusion	46
	Bibliography	47

List of Figures

2.1	A Mobius strip constructed using the quaternion action	8
3.1	The self-quotient of the plane curve $f(u) = u + 2 + \mathbf{i}(u^3 - 3u)$	11
4.1	A surface, surrounding the curve $\alpha(t) = (-t^3, t^2, 0)$, which closes off around the endpoint of the curve.	17
5.1	The curves $a(t)$ and $b(t)$ used to construct the spline surface below.	23
5.2	A spline-action surface	24
5.3	A surface surrounding the $(7, 2)$ -torus knot, and useful candidate for the spline-action surface construction	29
6.1	The generalized action $A \odot_f X$, where A , X , and f are as defined in the section.	34
7.1	A typical Dupin cyclide	41
7.2	A surface which blends two cones	41
7.3	A surface which blends two pyramids (cones over squares)	43

Acknowledgments

I would like to thank my advisor, Weiqing Gu, for all her excellent advice and helpful guidance. I would also like to thank my second reader, Francis Su, for his useful comments. Finally, I would like to thank Marshall Pierce, Victoria Krafft, and Adrian Mettler, for recovering my thesis from a sick computer at a crucial moment; Ariel Barton, for always being there when I needed her; and Maurice de Lafayette Omar, for much hard work.

Chapter 1

Introduction

In this thesis, we present two geometric constructions: the Minkowski product and the quaternion action, both of which are relatively new to the mathematical community. We will study both geometric and topological properties of these constructions, as well as a number of applications (mainly to computer-aided geometric design) and more specific constructions based on them.

In Chapter 2, we define both of these constructions, and present some of the geometric intuition and motivation behind them. In Chapter 3, we summarize some previous results about the topology of the Minkowski product. Chapter 4 discusses a particular use of the quaternion action to construct surfaces surrounding singular curves. In Chapter 5, we present the theory that results from examining quaternion actions of B-spline curves, and compare it to the standard theory of tensor product surfaces. Chapter 6 will motivate and describe a generalization of the quaternion action. Finally, Chapter 7 will return to the standard quaternion action, and examine in particular canal surfaces, their inversions, and an applications of the action to computer-aided geometric design; we then finish in Chapter 8 with a discussion of future work.

Chapter 2

Useful Facts and Definitions

We begin by briefly defining the quaternions (in which all our constructions will take place), and giving some useful facts about them. The proofs of all these facts are either trivial or can be found in [12]. We then present the definitions of the Minkowski product and of the action that we will be using, and discuss their geometric utility and some topological limitations.

2.1 The quaternions

Definition 1 *The set of all quaternions, denoted \mathbb{H} , is the four-dimensional Euclidean vector space over \mathbb{R} which is spanned by the independent vectors $\{1, \mathbf{i}, \mathbf{j}, \mathbf{k}\}$. We endow \mathbb{H} with a multiplicative structure whose identity is 1 by setting $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = -1$, $\mathbf{ij} = \mathbf{k} = -\mathbf{ji}$, $\mathbf{jk} = \mathbf{i} = -\mathbf{kj}$, and $\mathbf{ki} = \mathbf{j} = -\mathbf{ik}$, and by requiring that the multiplication be distributive.*

It follows from the definition that quaternion multiplication is noncommutative; however, it is easily shown to be associative. We also will make use of the following operations on \mathbb{H} , all of which are the analogues of similar operations on \mathbb{C} .

Definition 2 *The real part, or scalar part, of a quaternion $q = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$, denoted $\text{Re } q$, is a . The imaginary part, or vector part, of q is the quaternion $b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$, which has no real part; such a quaternion is called an imaginary quaternion. The conjugate of q , denoted \bar{q} or q^c , is the quaternion $\text{Re } q - \text{Im } q$; like in the complex case, we have*

the identity $qq^c = |q|^2$. Finally, the three-dimensional subspace of \mathbb{H} consisting of the imaginary quaternions is denoted $\text{Im } \mathbb{H}$; we frequently identify $\text{Im } \mathbb{H}$ with \mathbb{R}^3 .

The following fact about the quaternions is the source of most of their utility:

Fact 1 *Let p and q be unit quaternions (that is, $|p| = |q| = 1$). Then the mapping $f : \mathbb{H} \rightarrow \mathbb{H}$ given by $f(x) = pxq$ is a rotation (i.e., an orthogonal linear transformation of positive determinant). Moreover, all rotations of \mathbb{H} (and thus of \mathbb{R}^4) can be written in this way.*

That is, quaternion multiplication has essentially the same geometric interpretation in \mathbb{R}^4 as complex multiplication does in \mathbb{R}^2 .

In fact, we can be somewhat more precise about the specifics of a particular rotation. In the following definition, for a unit quaternion q , we use the shorthand $e^{q\theta} = \cos \theta + q \sin \theta$.

Fact 2 *Choose a positively-oriented orthonormal basis $q, p, q \times p$ of $\text{Im } \mathbb{H}$, the first element of which is q . The rotation induced by left multiplication of $e^{q\theta}$ is a rotation of θ in both the $1q$ -plane and the orthogonal $p(q \times p)$ -plane. The rotation induced by right multiplication is a rotation of θ in the $1q$ -plane, but of $-\theta$ in the orthogonal $p(q \times p)$ -plane.*

Since arbitrary unit quaternions are of the form $e^{q\theta}$ for some q and θ , this fact allows us to determine the precise geometric meaning of any given quaternion.

The following algebraic fact will occasionally prove useful:

Fact 3 *Let $q \in \mathbb{H}$ have nonzero imaginary part. Then $\{q, 1\}$ spans a subalgebra of \mathbb{H} which is isomorphic to \mathbb{C} . That is to say, every plane $P \subset \mathbb{H}$ which contains the real axis is a subalgebra of \mathbb{H} isomorphic to \mathbb{C} .*

Finally, the following fact is easily verified, but will be used frequently enough that we explicitly mention it here.

Fact 4 *Let $q \in \mathbb{H}$ have nonzero imaginary part. Then there are precisely two elements $p \in \mathbb{H}$ with $p^2 = q$, and both of them lie in the span of $\{1, q\}$. In other words, the quaternion square root of any non-real quaternion is unique up to sign.*

Note that this is not true, for example, of -1 , which is the square of any purely imaginary unit quaternion.

2.2 The Minkowski product

Definition 3 *Let A and B be subsets of a group G . The Minkowski product of A and B , denoted $A \otimes B$, is the set*

$$A \otimes B = \{ab \mid a \in A, b \in B\}. \quad (2.1)$$

The Minkowski product is a very general construction, as it applies to subsets of an arbitrary group. In general, however, its study under that name has been restricted to groups with some geometric character; specifically, it was originally defined by Farouki et al. as a construction in the multiplicative group of \mathbb{C} [4]. Unless otherwise indicated, we will always be considering the Minkowski product in the multiplicative group of \mathbb{H} ; this construction was originally suggested by Gu and Wiener [13].

In either \mathbb{C} or \mathbb{H} , the Minkowski product has a geometric significance, derived from the geometric significance of complex or quaternion multiplication. Suppose we have two sets A and B . It is clear that we can write their Minkowski product as

$$A \otimes B = \bigcup_{a \in A} aB = \bigcup_{b \in B} Ab, \quad (2.2)$$

where aB denotes the set formed by left-multiplying each element of B by a , and Ab is similarly defined in terms of right-multiplication. Multiplication by a constant in either \mathbb{C} or \mathbb{H} corresponds geometrically to a rotation and a scaling in \mathbb{C} or \mathbb{H} . Thus we can think of $A \otimes B$ as a union of rotated and scaled copies of B ,

where the elements of A describe the details of the rotating and scaling; similarly, it is a union of copies of A , rotated and scaled as described by B . This suggests that we might be able to use the Minkowski product to construct geometrical figures—such as a torus—which consist of rotated and scaled copies of some single other figure.

2.3 The quaternion action

Unfortunately, the Minkowski product is insufficiently general to construct the torus, or many other interesting geometric figures of this sort. This follows from the fact that one cannot obtain arbitrary rotations in \mathbb{H} through left multiplication; an arbitrary rotation in \mathbb{H} requires both left and right multiplication. Because of this, and because we it will also be convenient to allow translations of the figure, we introduce the following construction, which, in the quaternions, will remedy both of these defects of the Minkowski product at the expense of losing the symmetry between the sets involved:

Definition 4 Let $X \subset \mathbb{H}$; let $A \subset \mathbb{R} \times \mathbb{H}^3$. The action of A on X , denoted $A \odot X$, is given by

$$A \odot X = \{a_{\mathbb{R}} \mathbf{a}_{\text{left}} \mathbf{x} \mathbf{a}_{\text{right}}^c + \mathbf{a}_{\text{offset}} \mid (a_{\mathbb{R}}, \mathbf{a}_{\text{left}}, \mathbf{a}_{\text{right}}, \mathbf{a}_{\text{offset}}) \in A, \mathbf{x} \in X\}, \quad (2.3)$$

where q^c denotes the quaternion conjugate of q .

Thus $A \odot X$ consists of a union of rotated, scaled, and translated copies of X , where A describes how to rotate, scale, and translate. Moreover, unlike in the case of the Minkowski product, we can obtain arbitrary surfaces of this form through the action. For example, we can write the standard torus with core radius a and meridinal radius b as an action of the following two sets:

$$A = \{(b, \cos \frac{\phi}{2} + \mathbf{k} \sin \frac{\phi}{2}, \cos \frac{\phi}{2} + \mathbf{k} \sin \frac{\phi}{2}, a(\mathbf{i} \cos \phi + \mathbf{j} \sin \phi) \mid \phi \in [0, 2\pi)\} \quad (2.4)$$

$$X = \{\mathbf{i} \cos \theta + \mathbf{k} \sin \theta \mid \theta \in [0, 2\pi)\}. \quad (2.5)$$

Here X is a circle in the \mathbf{ik} -plane. The first component of A tells us to scale X by a factor of b ; the second and third to rotate it about the \mathbf{k} -axis by an angle ϕ ; the last, to translate it so its center lies on a circle of radius a , centered in the \mathbf{ij} -plane. Thus $A \odot X$ is a torus.

2.4 Topological issues and examples

Any topological analysis of the Minkowski product or quaternion action must begin with the following observation:

Proposition 1 *The Minkowski product $A \otimes B$ of two sets is a surjective image of their Cartesian product $A \times B$, endowed with the standard product topology, under the continuous map $\pi(p, q) = pq$. Similarly, the quaternion action $A \odot B$ of two sets is a surjective image of their Cartesian product, under the continuous map $\pi((r, q_i, q_r, q_o), p) = r q_i p q_r^c + q_o$.*

Once stated, the proof of this proposition is clear. It has a number of immediate topological corollaries: for example, it directly implies that the Minkowski product or quaternion action of any two compact sets is compact.

In many cases, particularly when we are generating surfaces for graphical display, we would like to avoid self-intersection for aesthetic reasons. In such cases, the maps alluded to in Proposition 1 become homeomorphisms onto their images. If we in addition wish to require that the sets A and B be connected curves (as we might wish to in order to avoid potentially unseemly discontinuities in the parametrization) we therefore obtain the following result:

Theorem 1 *Any set obtainable as the non-self-intersecting quaternion action (and thus also as the non-self-intersecting Minkowski product) of two connected curves A and B is homeomorphic to either an open disc, a cylinder, or a torus.*

Proof : As shown, for example, in [1], any connected curve is topologically either an open interval or a circle. In the given circumstance, $A \odot B$ is homeomorphic to

$A \times B$, by the previous argument. The three cases above follow when both A and B are homeomorphic to intervals, when one is homeomorphic to an interval and the other to a circle, and when both are homeomorphic to circles, respectively. ■

Thus, if we wish to obtain topologically interesting surfaces as the result of the quaternion action, we must violate some condition of the above theorem. We conclude this section with a pair of simple examples of how one might accomplish this.

We will begin by constructing a Möbius strip using the quaternion action. The construction is essentially identical to that of the torus given above (where $a = 2, b = 1$), with only two exceptions. First, we replace the circle on which we were acting by a straight line segment: that is, instead of the X used above, we set $X = \{tk | t \in [-1, 1]\}$. Solely making this change would result in a surface consisting of unions of vertical line segments a fixed distance away from the \mathbf{k} -axis: that is, a cylinder. In order to make the surface into a Möbius strip, we introduce a half-twist: before performing the rotation given by A , we rotate X in the \mathbf{ik} -plane by an angle $\phi/2$; we can do this by conjugating X with the unit quaternion $\cos(\phi/4) + \mathbf{j} \sin(\phi/4)$. If we make the assignment

$$q = \left(\cos \frac{\phi}{2} + \mathbf{k} \sin \frac{\phi}{2}\right) \left(\cos \frac{\phi}{4} + \mathbf{j} \sin \frac{\phi}{4}\right), \quad (2.6)$$

we therefore obtain

$$A = \{(1, q, q, 2(\mathbf{i} \cos \phi + \mathbf{j} \sin \phi)) | \phi \in [0, 2\pi]\} \quad (2.7)$$

as a set such that $A \odot X$ is a Möbius strip. A plot of this Möbius strip is shown in Figure 2.4.

Now, we can modify the above construction to obtain an embedding of the Klein bottle in $\mathbb{H} \simeq \mathbb{R}^4$. To do so, we replace X by the circle $\{e^{\mathbf{k}\theta} = \cos \theta + \mathbf{k} \sin \theta\}$. This replacement means that, instead of constructing a cylinder to which we are

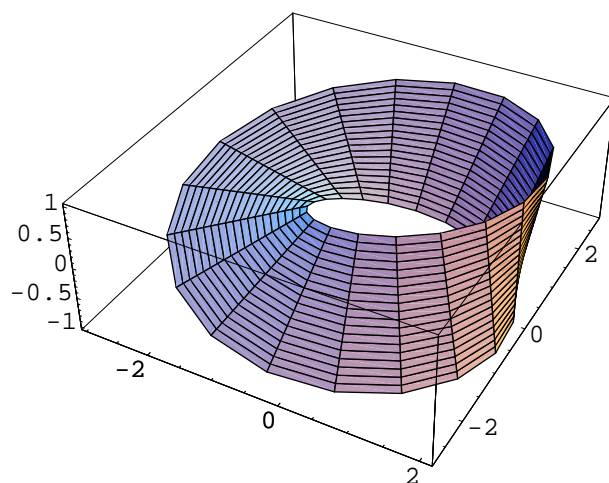


Figure 2.1: A Möbius strip constructed using the quaternion action

giving a half-twist, we are constructing a torus to which we are giving a half-twist; thus it yields a Klein bottle.

Finally, we examine the surfaces constructed in light of Theorem 1. The reason why their construction does not violate Theorem 1—despite the fact that both are quaternion actions of connected curves—is that the maps defined in Proposition 1 fail to be injective along the subsets $\phi = 0$, $\phi = 2\pi$ of the curve A . For other topological surfaces, similar constructions could be defined; however, since no other surface besides the projective plane is the surjective image of one of the three allowable topologies of $A \times B$, such constructions would in general have to involve several connected components which would be pieced together to form the surface.

Chapter 3

Geometry and Topology of the Minkowski Product

We will give a brief summary of the topological properties of the Minkowski product over \mathbb{C} . Much of the analysis extends to \mathbb{H} , and some of it to the quaternion action. For more details, the reader is invited to consult [11].

3.1 The self-quotient

As discussed in the previous section, any Minkowski product of curves is the surjective image of their Cartesian product, under the natural map π defined by $\pi(a, b) = ab$. Thus, in cases when we have sufficiently well-behaved sets (if, for example, they are compact, so injectivity of π implies that it is a homeomorphism onto its image), we need only fully understand the self-intersection of the sets in order to understand the topology of their Minkowski product. The key to studying the topological properties of the Minkowski product is therefore the *self-quotient* of a complex set A , defined as follows:

Definition 5 Let $A \subset \mathbb{C}$, and let $A^{-1} = \{a^{-1} | a \in A\}$. The self-quotient of A , denoted A^* , is the Minkowski product $A \otimes A^{-1}$.

The primary reason for its usefulness is the following theorem:

Theorem 2 Let A and B be subsets of \mathbb{C} . Then the natural map $\pi : A \times B \rightarrow A \otimes B$ is injective (and thus bijective) iff $A^* \cap B^* = \{1\}$.

Proof: By definition, π is injective iff there do not exist two distinct pairs $(a, b), (a', b') \in A \times B$ with $ab = a'b'$. But this holds iff there do not exist pairs $(a, a') \in A \times A, (b, b') \in$

$B \times B$ with $a/a' = b/b' \neq 1$ (since, if the quotients are 1, we have $a = a', b = b'$, which remains consistent with injectivity). But this clearly holds iff $A^* \cap B^* = \{1\}$, by the definition of the self-quotient. ■

The proof of the above theorem actually implies that even more is true: it implies that each point of intersection of the self-quotients A^* and B^* corresponds to at least one self-intersection of the Minkowski product $A \otimes B$ (possibly more if π is at least three-to-one at the point of intersection). Thus the self-quotients of sets provide a simple criterion for determining when the Minkowski products of those sets will self-intersect.

The other thing about the self-quotient which it is useful to note is that it frequently suffers from a particular kind of self-intersection at $\{1\}$; namely, $\{1\}$ has the behavior that any neighborhood of it becomes disconnected upon its removal. We refer to such a point as a *pinch point*. This behavior is illustrated in Figure 3.1, which shows the self-quotient of the curve given by $f(u) = u + 2 + \mathbf{i}(u^3 - 3u)$.

3.2 Results

We will now give a summary of some results that can be derived about planar Minkowski products. In the interests of space and clarity, we omit the proofs of these results; they can however be found in [11].

As it happens, the curves which are easiest to classify are those which are the injective images of the interiors of compact sets: that is, they are either the injective image of a circle, or the injective image of the interval $(0, 1)$ via a continuous map that can be extended to an image of the closed interval $[0, 1]$. For convenience, we will call such curves *well-behaved*.

The first really useful theorem we can prove is the following, which essentially says that the only singularities in the interiors of Minkowski products are those caused by self-quotients:

minkowskiexamples.nb

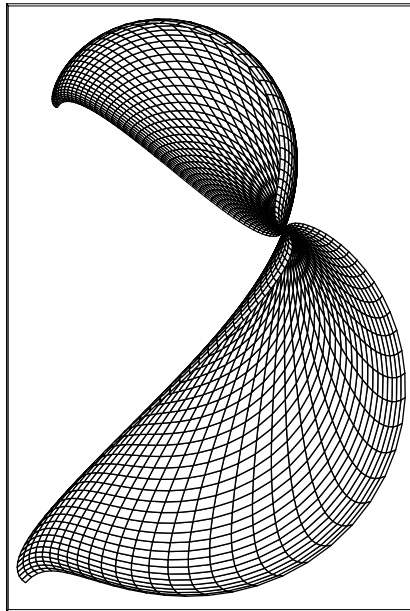


Figure 3.1: The self-quotient of the plane curve $f(u) = u + 2 + \mathbf{i}(u^3 - 3u)$.

Theorem 3 *Let A and B be well-behaved curves such that there do not exist neighborhoods U in A and V in B with $U = cV^{-1}$ for some $c \in \mathbb{C}$. Then there is a basis for the topology of $A \otimes B$ which consists solely of sets which are either open in $A \otimes B$ or homeomorphic to a closed half-plane.*

It is a fairly easy consequence of this fact that any such Minkowski product is homeomorphic to either an open disc or an open n -holed annulus for some n , with the possible addition of some of its boundary.

The obvious question is then: given a particular pair of well-behaved curves A and B , how do we compute the number of holes in $A \otimes B$? A complete (but not necessarily practical) characterization of this is given as follows:

Theorem 4 *Let A and B be well-behaved curves. Consider the subset M of $A^2 \times B^2$ given by*

$$M = \{(a, a', b, b') \mid a/a' = b/b'\}.$$

Then M has $2k + 1$ path-components, where k is the rank of the fundamental group of $A \otimes B$.

Moreover, we can put a more readily calculable upper bound on the rank of the fundamental group. Let the *projected boundary* of $A \otimes B$ be the set $\partial A \otimes B \cap A \otimes \partial B$; let the *envelope* of $A \otimes B$ be the set of all points which have a preimage p for which any open set containing p maps to a set homeomorphic to an open half-plane¹. It is clear that the boundary of a Minkowski product is readily calculable; as it happens, calculating the envelope of a Minkowski product is also relatively straightforward. Moreover, the following theorem holds:

¹This is not the standard definition of the envelope of a family of plane curves; however, it is essentially equivalent if the envelope is the Minkowski product of two well-behaved curves and is easy to define; thus we will use it in the following theorem.

Theorem 5 *The rank of the fundamental group is at most equal to one-half the number of points where either the projected boundary or the envelope self-intersects, or they intersect each other.*

In some special cases, we can determine a large amount of information about the Minkowski product even without avoiding self-quotients (and thus pinch points). For example, if A and B are well-behaved curves, both of which admit a real-analytic parametrization, then they are completely determined by any neighborhoods U and V ; it follows that, if there are any neighborhoods of A and B which locally appear to be inverses, then A and B are globally inverses. In particular, then, we can restrict the possible positions of multiple pinch points. Recall that a *logarithmic spiral* can be thought of as a 1-parameter subgroup of the multiplicative group of \mathbb{C} , and is parametrized as $f(t) = z^t$ for some nonzero complex z .

Theorem 6 *Suppose A and B are well-defined curves, both of which admit a real-analytic parametrization. Then there exists a (possibly degenerate) logarithmic spiral L on which every pinch point of $A \otimes B$ lies. In particular, if either A or B is not itself a logarithmic spiral, then, for any two pinch points p_1 and p_2 , there is some branch \log_* of the complex natural logarithm such that*

$$\frac{\log_* p_1}{\log_* p_2} \in \mathbb{Q}. \quad (3.1)$$

Chapter 4

Curves with Singularities

We now turn to a consideration of the quaternion action. Because of the greater complexity of the quaternion action, it is both more difficult to prove general topological results about it and easier to use it for specific constructions; most of our work involving it will therefore be of the second form.

The quaternion action will, given two curves, produce a surface in $\mathbb{H} \simeq \mathbb{R}^4$. In most cases, however, we will want to obtain a surface that lies in some three-dimensional space. There are two fundamentally different ideas we can use to do so.

First, we can carefully ensure that the surface produced by the action really lies in a three-dimensional space. For example, if our set $X \subset \mathbb{H}$ is purely imaginary, and our set $A \subset \mathbb{R} \times \mathbb{H}^3$ consists solely of points of the form $(a_{\mathbb{R}}, \mathbf{a}_0, \mathbf{a}_0, \mathbf{a}_{\text{offset}})$, where $\mathbf{a}_{\text{offset}}$ is purely imaginary, then the action $A \odot X$ will itself be purely imaginary. This method, which was used to produce the pictures in the previous section has the advantage of completely preserving the geometric intuition behind the action: the surfaces that result will consist of rotated and scaled copies of A , where X describes how we are rotating and scaling. However, this approach is rather restrictive: it can only be used to describe this very specific type of surface. It should be noted that the example shown in Chapter 2 was of this type.

Second, we can in some way project the surface $A \odot X$ into a three-dimensional subset of \mathbb{H} (generally $\text{Im } \mathbb{H}$). Depending on the specific projection used, the properties of the resulting classes of surfaces will of course vary. However, in each case, the class of surfaces that we can produce will be more general than the class of

surfaces producible by the first method.

In this chapter, we will discuss what is perhaps the most obvious method of projection, and some of its applications. The essential problem is simple. Given a piecewise differentiable curve γ , it is not generally possible to produce a pipe surface surrounding the curve, without the introduction of self-intersections in the surface near the singularity of the curve. However, by using our action and the technique of projection described above, we can easily construct a type of *generalized pipe surface*, consisting of a 1-parameter family of closed curves surrounding the curve, which does not in general intersect itself near the singularity, and which, when “far away” from the singularity, approximates a pipe surface.

4.1 Curves with endpoints

We will begin with a slightly simpler construction which illustrates the basic idea. Suppose we are given a differentiable curve γ with one endpoint. We wish to produce a generalized pipe surface surrounding γ , which behaves nicely at the endpoint. That is, unlike a regular pipe surface, which simply stops at the endpoint, we wish this surface to close there.

We will proceed as follows. Give γ some parametrization $\alpha : [0, \infty) \rightarrow \mathbb{R}^3$, such that $\lim_{t \rightarrow 0} \alpha'(t) = 0$, but $\alpha'(t) \neq 0$ for all positive t . Identify \mathbb{R}^3 with $\text{Im } \mathbb{H}$, and construct a parametrized curve $\bar{\alpha} : [0, \infty) \rightarrow \mathbb{H}$, such that $\lim_{t \rightarrow 0} \bar{\alpha}'(t)$ is nonzero, and $\text{Im } \bar{\alpha}(t) = \alpha(t)$ for all t . The most straightforward way to do this is to set $\bar{\alpha}(t) = t + \alpha(t)$.

Using the action, we then construct a surface surrounding $\overline{\alpha(t)}$. Set

$$\beta(t) = \cos\left(\frac{\theta(t)}{2}\right) \cos\left(\frac{\phi(t)}{2}\right) + \mathbf{i} \cos\left(\frac{\theta(t)}{2}\right) \sin\left(\frac{\phi(t)}{2}\right) - \mathbf{j} \sin\left(\frac{\theta(t)}{2}\right) \sin\left(\frac{\phi(t)}{2}\right) + \mathbf{k} \sin\left(\frac{\theta(t)}{2}\right) \cos\left(\frac{\phi(t)}{2}\right), \quad (4.1)$$

and

$$\eta(t) = \cos\left(\frac{\psi(t)}{2}\right) + \mathbf{i} \sin\left(\frac{\psi(t)}{2}\right), \quad (4.2)$$

where $\theta(t)$ and $\phi(t)$ are the Euler angles in \mathbb{R}^3 required to rotate the \mathbf{ij} -plane so that it is orthogonal to $\alpha'(t)$, and $\psi(t)$ is defined by

$$\sin \psi(t) = \frac{|\alpha'(t)|}{|\bar{\alpha}'(t)|} \quad (4.3)$$

$$\cos \psi(t) = \frac{\operatorname{Re} \bar{\alpha}'(t)}{|\bar{\alpha}'(t)|}. \quad (4.4)$$

Now, we act on the set

$$X = \{\mathbf{i} \cos t + \mathbf{j} \sin t | t \in [0, 2\pi)\} \quad (4.5)$$

(that is, a circle in the \mathbf{ij} -plane) by the set

$$A = \{(\mathbf{r}, \eta(t)\beta(t), \beta(t)\eta(t)^c, \overline{\alpha(t)}) | t \in [0, \infty)\}. \quad (4.6)$$

The resulting surface will be a 1-parameter family of circles of radius \mathbf{r} , centered on $\overline{\alpha(t)}$. Moreover, by the way we constructed η and β , these circles will lie in a space whose angle with $\operatorname{Im} \mathbb{H}$ is $\psi(t)$, and their projections into $\operatorname{Im} \mathbb{H}$ will be orthogonal to $\alpha(t)$.

Now, consider the natural projection of this surface into $\operatorname{Im} \mathbb{H}$ obtained by setting the real coordinates of each of its points to zero. It will necessarily be a generalized pipe surface of the type described above, surrounding $\alpha(t)$. The projection of each circle will be an ellipse of semimajor axis \mathbf{r} , the eccentricity of which at a given point is determined by $\psi(t)$: as $\psi(t)$ increases, so will the eccentricity. When $\psi(t) = \pi/2$, the ellipse will have eccentricity 1 and will therefore be a line segment. By the definition of ψ , this will occur precisely when the original curve $\alpha(t)$ has zero derivative; thus it will occur only when $t = 0$. So the resulting surface will smoothly curve in on itself as it approaches the end of the curve. If we choose $\operatorname{Re} \bar{\alpha}'(t)$ such that its derivative is small relative to $\alpha'(t)$, then $\phi(t)$ will be small for large t ; thus the surface will approximate a pipe surface in this case.

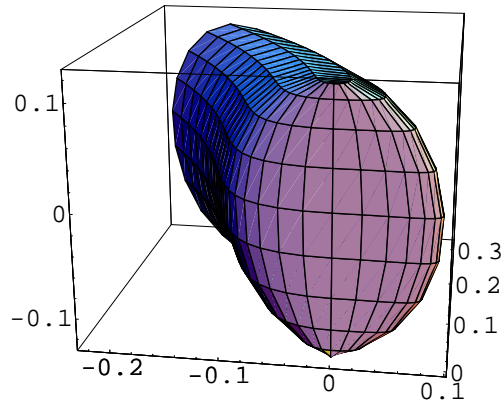


Figure 4.1: A surface, surrounding the curve $\alpha(t) = (-t^3, t^2, 0)$, which closes off around the end-point of the curve.

A surface constructed as above is shown in Figure 4.1. We selected the original curve $\alpha(t) = (-t^3, t^2, 0)$, lifted it to the curve $\bar{\alpha}(t) = t - \mathbf{i}t^3 + \mathbf{j}t^2$, and applied this construction to it using $r = 1/4$.

4.2 Curves with singularities

Having performed the above construction, we now return to the original problem of the chapter: to construct a generalized pipe surface surrounding a piecewise differentiable curve, without self-intersection. The construction is similar to the one in the previous section; however, the requirement that we avoid self-intersection makes it slightly more delicate. We proceed as follows:

Let γ be continuous and piecewise differentiable. We will assume for simplicity that our parametrization γ has only one singularity; however, the construction extends straightforwardly to cases where γ has multiple (but only finitely many) singularities. For reasons which will soon become apparent, we will also restrict

our attention to curves γ where the right- and left-hand limits of the unit tangent vector at the singularity are well-defined.

Given the above assumptions, we can suppose without loss of generality that γ admits a differentiable parametrization α for which the singularity occurs at $\alpha(0)$ (and thus $\alpha'(0) = 0$). By translating γ , we can also assume that $\alpha(0) = 0$. Moreover, since the unit tangent vector has well-defined limits at the singularity, γ has a (not necessarily unique) tangent plane at the singularity; thus we can rotate γ so that tangent plane is the \mathbf{ij} -plane, and so there is some neighborhood $(-\epsilon, \epsilon)$ of 0 such that, if $|t| < \epsilon$, then $\alpha(t)$ has positive x -coordinate iff t is positive.

We proceed almost identically as before. First, we lift α to some $\bar{\alpha} \in \mathbb{H}$ with nonzero derivative at 0. Then we construct a one-parameter family of circles centered at the points on $\bar{\alpha}$. We will find it useful to construct this family through successive rotations, as follows:

Begin by placing a circle of some fixed radius R parallel to the xy -plane at each point on $\bar{\alpha}(t)$.

Rotate each of these circles by an angle $\psi(t)$ in the $\mathbf{1i}$ -plane, where $\cos \psi = \frac{\operatorname{Re} \bar{\alpha}'(t)}{|\bar{\alpha}'(t)|}$, and $\sin \psi = \frac{|\alpha(t)|}{|\bar{\alpha}'(t)|}$.

Rotate each of them by an angle $\phi(t)$ in the \mathbf{ik} -plane, where $\cos \phi(t) = \frac{\alpha'_3(t)}{|\alpha'(t)|}$, $\sin \phi(t) = \frac{\sqrt{\alpha'_1(t)^2 + \alpha'_2(t)^2}}{|\alpha'(t)|}$.

Note that, up until this point, we have proceeded precisely as we did before (since the ϕ defined above is easily shown to be the ϕ used in Section 4.1 above. For the previous method, we would now rotate each circle by some angle $\theta(t)$ in the \mathbf{ij} -plane, calculated to make their projections in \mathbb{R}^3 lie orthogonally to the curve γ . However, if we do so, there is no guarantee that the resulting surface will not have self-intersections. We avoid these self-intersections as follows.

Rotate each circle by an angle $\vartheta(t)$ in the \mathbf{ij} -plane, where $0 \leq \vartheta \leq \pi/2$, and $\cos \vartheta(t) = \frac{|\alpha_1(t)|}{R + |\alpha_1(t)|}$. Since $0 \leq \cos \vartheta(t) \leq 1$, we will always be able to define ϑ uniquely in this way. Furthermore, by the way ϑ is chosen $R \cos \vartheta$ will always be

less than $|\alpha_1(t)|$. Since $R \cos \vartheta$ gives an upper bound on the x -distance of a point on the circle from $\bar{\alpha}(t)$, it follows that all the points on the circle centered at $\bar{\alpha}(t)$ will have the same sign as t . Thus there is some neighborhood of 0 in which these circles don't intersect, so we can choose R small enough that they do not intersect in general.

It follows that, when we project back down into $\text{Im } \mathbb{H} \simeq \mathbb{R}^3$, the resulting curve is a generalized pipe surface which has the desired properties.

Chapter 5

Spline-action Surfaces

5.1 Background

In order to be able to discuss our next topic, the spline-action surface, we will briefly review the definitions of B-spline curves and of tensor product surfaces.

A *B-spline curve* of degree m is a piecewise polynomial map $f : \mathbb{R} \rightarrow V$ of maximum degree m , where V is a vector space over \mathbb{R} . To each B-spline curve we associate a nondecreasing *knot sequence* $U = \{u_i\}_{i=-\infty}^{\infty}$. We refer to the number of times that a particular $u_i \in U$ is repeated as the *multiplicity* of u_i , and require the following two conditions on f :

1. The restriction of f to any half-open interval $[u_i, u_{i+1})$ is polynomial.
2. At each u_i , f and its first $m - k$ partial derivatives are continuous, where k is the multiplicity of u_i .

For the purposes of computing B-spline curves, we can associate to any B-spline curve f an infinite sequence of points in \mathbb{R}^n , known as the *de Boor control points* of the curve, such that evaluating f at any point is simply a matter of repeated linear interpolation between the control points. Computing a point on the curve in this way is known as *De Boor's algorithm*; to perform it on a degree m curve in d -dimensional space requires $d \binom{m+1}{2}$ linear interpolations.

A *tensor product surface* of degree (m, n) is a piecewise polynomial map $f : \mathbb{R}^2 \rightarrow V$, of maximum degree m in the first coordinate and maximum degree n in the second coordinate, where V is a vector space over \mathbb{R} , and where there exists some

collection of horizontal and vertical lines such that f is polynomial in each region bounded by those lines.

It is often convenient to think of a tensor product surface as a 1-parameter family of degree m B-spline curves, such that the de Boor control points of any one of these curves themselves lie on degree n B-spline curves.

For a more detailed description of B-spline curves, tensor product surfaces, and de Boor's algorithm, the reader is invited to consult Gallier [6].

5.2 Spline-action surfaces

5.2.1 Polynomial surfaces

Since both \mathbb{H} and $\mathbb{R} \times \mathbb{H}^3$ are vector spaces over \mathbb{R} , it is reasonable to talk about B-spline curves in each of them. This fact motivates the following definition:

Definition 6 *A polynomial spline-action surface $S \subset \mathbb{H}$ is a surface given by the Minkowski action $A \odot X$, where A is the image of a spline curve $f(s)$ in $\mathbb{R} \times \mathbb{H}^3$, and X is the image of a spline curve $g(t)$ in \mathbb{H} .*

Note that we can directly use this definition to compute points on some spline-action surface S as follows. Suppose we have some pair of B-spline curves, parametrized as $f(s) = (f_{\mathbb{R}}(s), \mathbf{f}_{\text{left}}(s), \mathbf{f}_{\text{right}}(s), \mathbf{f}_{\text{offset}}(s)) \in \mathbb{R} \times \mathbb{H}^3$, $\mathbf{g}(t) \in \mathbb{H}$. Then S inherits a natural parametrization $[f \odot g](s, t)$ from f and g . Under this parametrization, we can compute an arbitrary point $[f \odot g](s_0, t_0)$ on S as follows: using two applications of De Boor's algorithm, compute $f(s_0)$ and $g(t_0)$. We then use the formula

$$[f \odot g](s_0, t_0) = f_{\mathbb{R}}(s_0) \mathbf{f}_{\text{left}}(s_0) \mathbf{g}(t_0) \overline{\mathbf{f}_{\text{right}}(s_0)} + \mathbf{f}_{\text{offset}}(s_0), \quad (5.1)$$

which follows directly from the definition of the action.

An example of the construction of a spline-action surface follows. We begin with two curves: one, $a(s)$, which approximates a helix, and the other, $b(t)$, which

approximates a circle, and we wish to form a screw-like surface (by rotating and scaling copies of the circle in a way which corresponds to the helix); see Figure 5.2.1 for plots of the specific curves used. To do so, we calculate $a'(s)$ and select a pair of vectors orthogonal to $a'(s)$. Ideally, we would apply to $b(t)$ a rotation which took the x - and y -axes to this pair of vectors. However, even if these vectors are themselves piecewise polynomial functions of s , the quaternions required to perform this rotation may not be.

In order to keep our rotation function piecewise polynomial, we therefore must make more approximations. Instead of rotating by conjugating with a unit quaternion, we will rotate by conjugating with a quaternion of the form $1 + kv$, where $v \in \text{Im } H$ is some predetermined vector perpendicular to $a'(s)$. In this case, since $a'(s)$ never has zero z -component, we select $v = (a'_2(s), -a_1(s), 0)$, where $a'(s) = (a'_1(s), a'_2(s), a'_3(s))$. Some calculation shows that such a conjugation will rotate the plane $z = 0$ so that it is perpendicular to $a'(s)$ iff $v = \frac{1}{a'_3(s)}$. In this case, $a'_3(s)$ happens to be identically $1/2$; thus the quaternion by which we wish to conjugate is in fact polynomial. The resulting surface (with an appropriate normalization of $b(t)$ to avoid self-intersection) is shown in Figure 5.2.1 below.

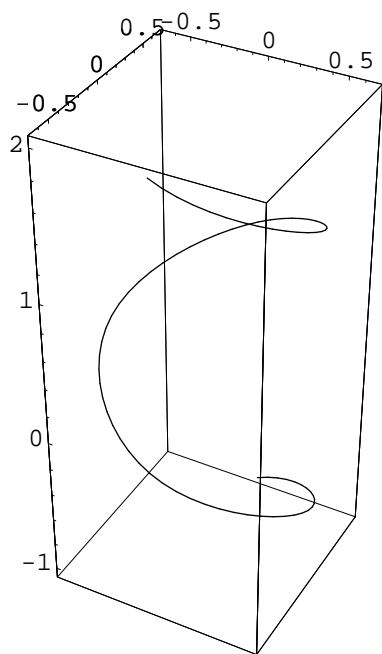
5.2.2 Rational surfaces

Definition 7 A rational spline-action surface $S \subset \mathbb{R}^3$ is a surface whose parametrization is given by $g(s, t) = \frac{\text{Im } f(s, t)}{\text{Re } f(s, t)}$, where $f(s, t)$ parametrizes a polynomial spline-action surface, and $\text{Im } \mathbb{H}$ is identified with \mathbb{R}^3 in the natural way.

We can thus think of a rational spline-action surface as being the result of looking at the projection of a spline surface in $\mathbb{H} - 0$ into its quotient space $\mathbb{R}P^3$, and then examining the resulting surface as a subset of $\mathbb{R}^3 \subset \mathbb{R}P^3$.

A substantial advantage of the spline-action method of surface construction is its geometrically intuitive character: that is, that it allows the creation of surfaces

spline.nb



spline.nb

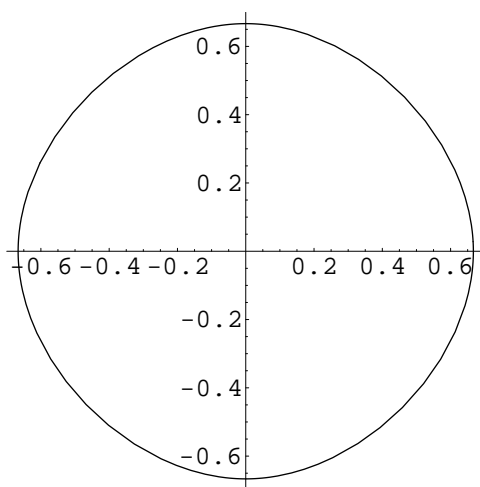


Figure 5.1: The curves $a(t)$ and $b(t)$ used to construct the spline surface below.

spline.nb

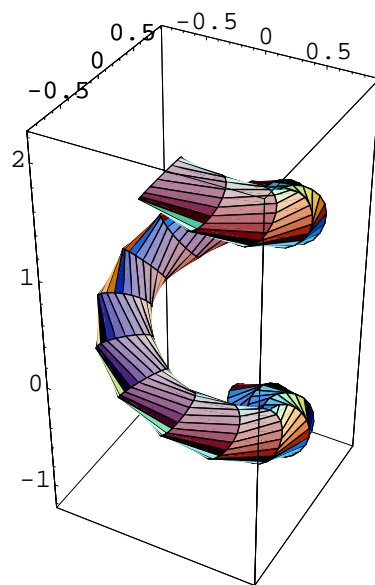


Figure 5.2: A spline-action surface

by simply rotating and scaling curves in a pre-determined manner. It is therefore a natural question to ask whether the above quotient operation preserves this geometric intuition. As the following proposition shows, the answer is a qualified yes:

Proposition 2 *Let $S = \overline{f \odot g}$ be a rational spline-action surface, formed by the projection of $f \odot g$ as outlined above. Then S will consist of rotated and scaled copies of some rational spline curve $A \subset \mathbb{R}^3$ if \mathbf{f}_{left} and $\mathbf{f}_{\text{right}}$ are equal and f_{offset} is purely imaginary.*

Proof : Fix the argument of f at s_0 . The resulting curve in $\mathbb{R}^3 \simeq \text{Im } H$ is given by

$$a_{s_0}(t) = \frac{\text{Im}[f \odot g](s_0, t)}{\text{Re}[f \odot g](s_0, t)}. \quad (5.2)$$

Since \mathbf{f}_{left} and $\mathbf{f}_{\text{right}}$ are equal, $g(t)$ has the same real part as $\mathbf{f}_{\text{left}}(s)g(t)\overline{\mathbf{f}_{\text{right}}(s)}$. Since $\mathbf{f}_{\text{offset}}$ is purely imaginary, it follows that $\text{Re } g(t) = \text{Re}[f \odot g](s, t)$ for all s, t . It follows that $a_{s_0}(t) = [f \odot h](s_0, t)$, where

$$h(t) = \frac{\text{Im } g(t)}{\text{Re } g(t)}; \quad (5.3)$$

thus S' is just $[f \odot h]$. ■

5.2.3 Relationships between spline-action surfaces and tensor product surfaces

Using Equation 5.1, we can see that any spline-action surface is automatically a B-spline tensor product surface in \mathbb{H} : it is a map of \mathbb{R}^2 into \mathbb{H} , which is polynomial on the interiors of some collection of rectangles. Specifically, it is polynomial on the interiors of the rectangles whose vertices are of the form (u_i, v_j) , where $\{u_i\}$ is the knot sequence of the spline curve $f(s)$ and $\{v_j\}$ the knot sequence of the spline curve $g(t)$.

Moreover, there clearly exist tensor product surfaces which are not spline-action surfaces: for any spline-action surface, the curves formed by fixing s must be

achievable from each other by rotation, translation, and scaling. This translates into the algebraic condition that, for any real coordinates x_0, x_1, y_0, y_1 , the condition

$$[f \odot g](x_1, y_1) - [f \odot g](x_1, y_0) = q ([f \odot g](x_0, y_1) - [f \odot g](x_0, y_0)) q' \quad (5.4)$$

for some $q, q' \in \mathbb{H}$ which depend only on y_0 and y_1 .

The above condition clearly implies that the set of all surfaces that are realizable as spline-action surfaces is a proper subset of those realizable by the more common tensor-product surfaces. However, the spline-action construction is still useful, for two reasons.

First, it is a much more geometrically intuitive way to construct those surfaces for which it can be used. That is, if we wish to construct a surface consisting of rotated, scaled, and translated copies of some specific curve, the spline-action surface provides a geometrically obvious way to do it. Since many important types of curves (most notably, canal surfaces and approximations to canal surfaces) are of this form, the spline-action surface clarifies how to construct such a surface.

Second, it provides a significant gain in efficiency in those cases where it does in fact apply. We will analyze this further in the next section.

5.3 Run-time comparison with tensor product surfaces

As previously claimed, in the event that a given surface is calculable as a spline-action surface, the applications of De Boor's algorithm needed to do so are substantially faster than those applications required to compute the surface as a tensor product surface. The following theorem makes this notion precise:

Theorem 7 *Let $S = f \odot g$ be a spline-action surface. If we calculate it as such, using de Boor's algorithm, the resulting computation will require fewer linear interpolations as would the computation of S as a tensor product surface. In the limit, if f is of large degree,*

it will require only $13/36 \simeq 0.36$ times as many linear interpolations for any fixed degree of g .

Proof : Suppose f is a B-spline curve of degree m , and g a B-spline curve of degree n . De Boor's algorithm requires $d \binom{n+1}{2}$ linear interpolations to compute a point on a curve of degree n in d -dimensional space. Whether we calculate a point on S as a spline-action surface or as a tensor product surface, it will require two applications of De Boor's algorithm.

To compute S as a tensor product surface, we initially use the relation

$$[f \odot g](s, t) = f_{\mathbb{R}}(s) \mathbf{f}_{\text{left}}(s) \mathbf{g}(s) \mathbf{f}_{\text{right}}(s) + \mathbf{f}_{\text{offset}}(s) \quad (5.5)$$

to obtain a piecewise polynomial parametrization of $f \odot g$. We then apply de Boor's algorithm twice to compute a point on this surface. The surface is in a 4-dimensional space, and is of degree $3m$ in s and n in t . Thus we require $4 \binom{3m+1}{2} + 4 \binom{n+1}{2}$ linear interpolations to compute S in this case.

To compute S as a Minkowski product surface, on the other hand, we apply de Boor's algorithm separately to compute f and g . We are then computing one 13-dimensional instance of de Boor's algorithm, of degree m , and one 4-dimensional instance, of degree n . It follows that $13 \binom{m+1}{2} + 4 \binom{n+1}{2}$ terms are required.

It is clear that we need only compare the first term in each sum, and not difficult to verify that it is always smaller in the second case than in the first. ■

It is even possible to improve upon this result in a number of the cases discussed above, which lower the dimension of the auxiliary space without decreasing the degree of the curve. The following two theorems demonstrate this fact:

Theorem 8 *Let $S = f \odot g$ be a purely imaginary spline-action surface, where $g \in \text{Im } \mathbb{H}$ and the action of f is to conjugate by some quaternion and shift the resulting curve by some purely imaginary quantity. Then the results of Theorem 7 hold, except that, in the limit,*

the computation of S as a spline-action surface requires only $8/27 \simeq 0.30$ times as many linear interpolations as are required for the tensor-product calculation.

Theorem 9 *Let $S = f \odot g$ be a rational spline-action surface, where g is restricted as in Theorem 8 (so the results of Proposition 2 hold). Then the results of Theorem 7 hold, except that, in the limit, the computation of S as a spline-action surface requires only $8/36 \simeq 0.22$ times as many linear interpolations as are required for the tensor-product calculation.*

The proofs of these theorems are omitted, as they are essentially identical to that of Theorem 7.

The asymptotic estimates in the three theorems above hold when the degree of f increases while the degree of g remains fixed. Indeed, the amount of computation corresponding to g remains the same for either method. Thus we can conclude that the spline-action method is likely to yield the greatest improvement over the tensor-product method in the case where the degree of f is larger than the degree of g . In more concrete terms, this will occur when we are given a curve to be rotated and translated, and one of two situations arises. Either the curve used has an exact rational parametrization of low degree, or we care much more about the particulars of how the rotation and translation occurs than we do about the exact shape of the curve itself. In many applications one of these is likely to be the case. For example, a circle admits such a rational parametrization; thus if we consider spline-action surfaces on circles (which are canal surfaces) the resulting calculations will likely be much more efficient than a tensor product construction of the same surface. Additionally, if we are looking at a surface which roughly approximates a pipe surface from the side, it is frequently much easier to discern the precise shape of its spine curve (which is governed by f) than the circle surrounding it (which is governed by g); thus, in many graphical examples, the construction may be useful even if the precise curve desired is not rational.

trefoil.nb

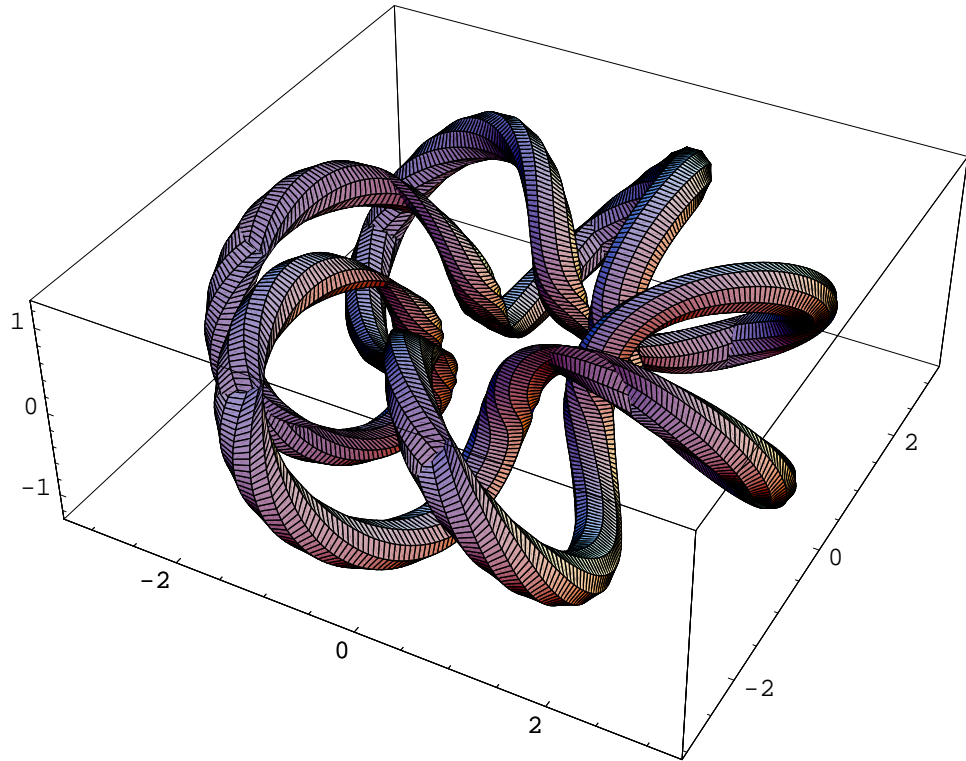


Figure 5.3: A surface surrounding the $(7, 2)$ -torus knot, and useful candidate for the spline-action surface construction

For example, a surface which could be profitably constructed using spline-action surfaces is shown in Figure 5.3. This surface is constructed through action on a circle; however, the circle is being rotated and translated in a fairly complex way, so as to surround a parametrization of the $(7, 2)$ -torus knot.

Chapter 6

Inversion and Generalized Actions

Up until now, we have for the most part been studying the quaternion action of one parametrized curve on another. This has resulted in surfaces of the form

$$S(u, v) = f_{\mathbb{R}}(u)\mathbf{f}_{\text{left}}(u)\mathbf{g}(v)\mathbf{f}_{\text{right}}(u)^c + \mathbf{f}_{\text{offset}}(u). \quad (6.1)$$

While this class of surfaces has many useful properties, as described in previous chapters, it also has its defects. Among these is that it is not closed under the geometric operation of inversion, which may be used to construct many useful surfaces. In this chapter, we describe a generalization of the quaternion action which is naturally closed under inversion, discuss its utility, and show that it is completely general in the sense that any compact differentiable surface may be constructed in this form. In later chapters, we will return to the problem of inverting surfaces resulting from the quaternion action, and discuss special cases in which it is possible.

6.1 A review of inversion

We begin by briefly reminding the reader of the definition of inversion.

Definition 8 Let $p \in \mathbb{R}^k - \{0\}$. The inversion of p about the origin, denoted hereafter $\iota(p)$, is the point $\frac{p}{|p|^2}$. For any $S \subset \mathbb{R} - \{0\}$, the inversion of S about the origin is given by

$$\iota(S) = \{\iota(p) | p \in S\}. \quad (6.2)$$

From a geometric perspective, $\iota(p)$ is the point lying on the same ray emanating from the origin as p , whose distance from the origin is the reciprocal of p 's.

The quaternions provide a particularly natural framework for studying inversions in \mathbb{R}^4 , as is indicated by the following fact:

Fact 5 For any $p \in \mathbb{H}$, $\iota(p) = (p^c)^{-1}$.

Proof : $\iota(p) = \frac{p}{|p|^2} = p(p^c p)^{-1} = (p^c)^{-1}$. ■

6.2 Inverting the quaternion action

Armed with the above fact, we can determine $\iota(A \cdot X)$ for any parametrized curves A and X . Suppose $A \subset \mathbb{R} \times \mathbb{H}^3$ is parametrized as $(f_{\mathbb{R}}(u), \mathbf{f}_{\text{left}}(u), \mathbf{f}_{\text{right}}(u), \mathbf{f}_{\text{offset}}(u))$, and $X \subset \mathbb{H}$ is parametrized by $\mathbf{g}(v)$. We will restrict our attention to the cases where both g and every component of A 's parametrization are everywhere nonzero.

Begin by defining

$$\tilde{\mathbf{f}}_{\text{offset}}(u) = \frac{(f_{\text{right}}(u)^c)^{-1} f_{\text{offset}}(u) f_{\text{left}}(u)}{f_{\mathbb{R}}}. \quad (6.3)$$

Then (6.1) reduces to

$$S(u, v) = f_{\mathbb{R}}(u) \mathbf{f}_{\text{left}}(u) (\mathbf{g}(v) + \tilde{\mathbf{f}}_{\text{offset}}(u)) \mathbf{f}_{\text{right}}(u)^c. \quad (6.4)$$

Thus, applying Fact 5, we have

$$\iota(S(u, v)) = \frac{\mathbf{f}_{\text{right}}(u)^{-1} (\mathbf{g}(v)^c + \tilde{\mathbf{f}}_{\text{offset}}(u)^c)^{-1} \mathbf{f}_{\text{left}}(u)^{-1}}{f_{\mathbb{R}}(u)}. \quad (6.5)$$

But

$$(\mathbf{g}(v)^c + \tilde{\mathbf{f}}_{\text{offset}}(u)^c)^{-1} = \frac{\mathbf{g}(v) + \tilde{\mathbf{f}}_{\text{offset}}(u)}{|\mathbf{g}(v)|^2 + |\tilde{\mathbf{f}}_{\text{offset}}(u)|^2 + 2\text{Re}\{\mathbf{g}(v)\tilde{\mathbf{f}}_{\text{offset}}(u)\}}; \quad (6.6)$$

substituting this into (6.5) yields

$$\iota(S(u, v)) = \frac{\mathbf{f}_{\text{right}}(u)^{-1} (\mathbf{g}(v) + \tilde{\mathbf{f}}_{\text{offset}}(u)) \mathbf{f}_{\text{left}}(u)^{-1}}{f_{\mathbb{R}}(u) (|\mathbf{g}(v)|^2 + |\tilde{\mathbf{f}}_{\text{offset}}(u)|^2 + 2\text{Re}\{\mathbf{g}(v)\tilde{\mathbf{f}}_{\text{offset}}(u)\})}. \quad (6.7)$$

This is similar in form to a quaternion action. However, in quaternion actions of curves, the real scaling factor solely dependent on the acting curve (in this case, it

is a function of u alone). On the other hand, (6.7) has a real scaling factor which depends on both curves (i.e., it is a function of both u and v). This observation motivates the definition of the generalized action, to which we now proceed.

6.3 The generalized quaternion action

Definition 9 Let $A \subset \mathbb{H}^3$; let $X \subset \mathbb{H}$, and consider a function $f : A \times X \rightarrow \mathbb{R}$. The generalized quaternion action of A on X mediated by f , denoted $A \odot_f X$, is given by

$$A \odot_f X = \{f(\mathbf{a}, x)a_{\text{left}}(x + a_{\text{offset}})a_{\text{right}}^c \mid x \in X, \mathbf{a} = (a_{\text{left}}, a_{\text{right}}, a_{\text{offset}}) \in A\}. \quad (6.8)$$

The definition is analagous to the alternate formulation of the quaternion action suggested by (6.4). Geometrically, the standard formulation corresponds to taking a set, rotating it somehow, and then translating the resulting set elsewhere; the alternate formulation, on the other hand, corresponds to taking a set, translating it, and then rotating the translating set. In the case of the standard action, the two formulations are obviously equivalent. However, it is not at all obvious that, if we were to define a generalized action using the standard definition, it would yield the same class of surfaces. Our use of this alternate formulation is motivated by (6.7), which is of the form given by (6.8).

It should also be noted that it is not quite immediately clear that the generalized quaternion action is in fact a generalization of the quaternion action. In the original quaternion action, it was possible to have any arbitrary set of scaling factors; in the generalized action, though the scaling factors may depend on either set, there is a single well-defined factor for every possible choice of the other parameters. However, since we have not restricted the lengths of the quaternions a_{left} and a_{right}^c , any scaling factor which appears in a standard quaternion action can be absorbed into these quaternions to yield a representation of it as a generalized action.

Geometrically, the generalized action does the following. Given a subset X of \mathbb{H} , it first translates it (in a way determined by a_{offset}), and then rotates and scales it

(in a way determined by a_{left} and a_{right} . Finally, it performs an additional scaling, which may scale different portions of X by different amounts (in a way determined by f). Given this geometric perspective, the following proposition is highly reasonable. Because of this, and because its proof is essentially identical to the argument given in the previous section, we omit its proof.

Proposition 3 *Let $A \otimes_f X$ be formed by generalized quaternion action. Then the set $\iota(A \otimes_f X)$ may also be formed by generalized quaternion action.*

So the class of all surfaces which may be formed under generalized quaternion action is in fact closed under inversion. Moreover, applying the argument from the previous section gives a simple formula for the inversion's parametrization, in terms of the parametrization of the original surface.

We now present a simple example of a generalized action. If we let X be the unit circle in the \mathbf{jk} -plane parametrized by t , let $A = (1, 1, s\mathbf{i})$ (so we are not rotating X at all), and set $f(x, (a_1, a_2, a_3)) = \sin s \sin^2 t + 1$, we obtain the distorted cylinder shown in Figure 6.3.

6.4 The universality of the generalized action

The geometric interpretation of the generalized action described above makes it clear that the generalized action is a very general construction indeed. In fact, as we will show in this section, every compact differentiable surface in \mathbb{R}^3 is essentially a generalized action.

The main part of the argument which establishes the above claim lies in the following lemma:

Lemma 1 *Let $S \subset \text{Im } \mathbb{H} \simeq \mathbb{R}^3$ be a differentiable surface, with $p \in S$. Suppose the line containing p and the origin is not tangent to S . Then there is some open neighborhood $U_p \ni p$ which is a generalized action.*

posterpictures.nb

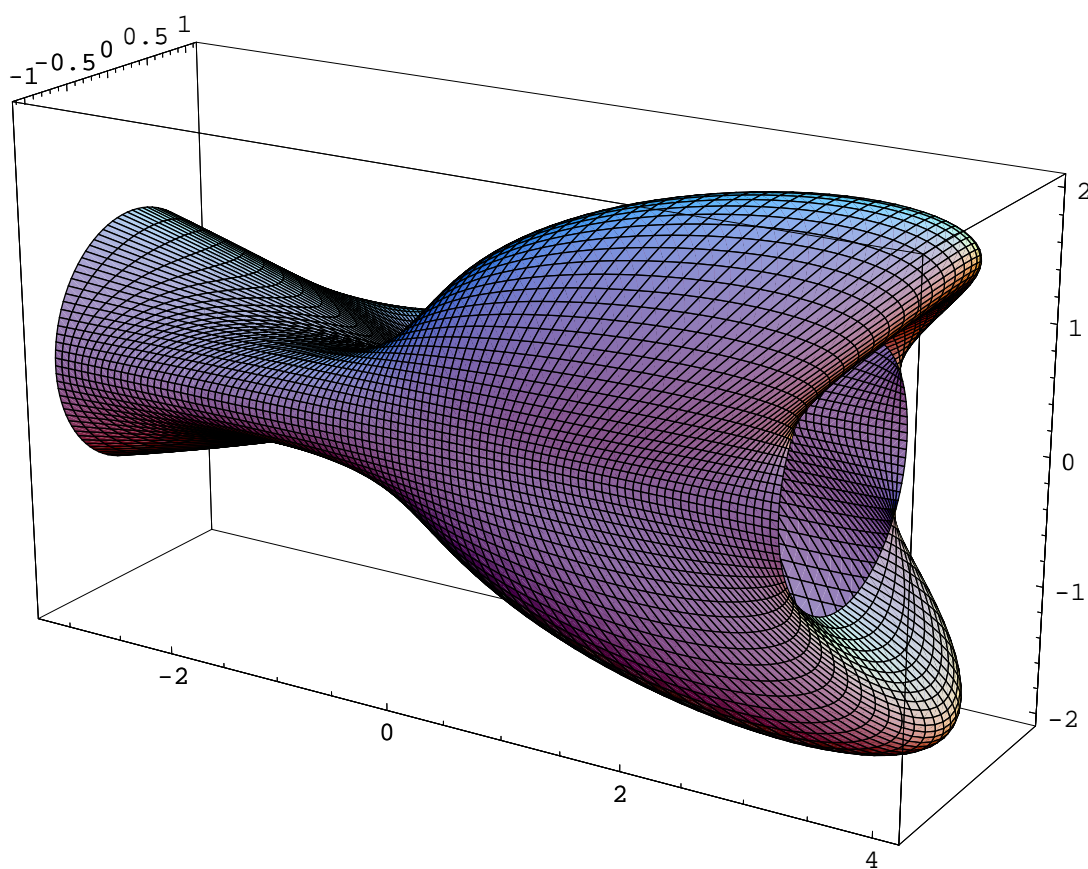


Figure 6.1: The generalized action $A \odot_f X$, where A , X , and f are as defined in the section.

Proof : Let n be the normal vector to S at p ; consider a vector $v \in T_p S$. Then there is some open neighborhood $V \subset S$ of p in which the plane spanned by n and v intersects S in a differentiable curve $\gamma_v(t)$. Moreover, since the line containing p and the origin is not tangent to S , we can suppose without loss of generality (possibly by shrinking V) that the line containing $\gamma_v(t)$ and the origin is not tangent to S , and that $\gamma_v(t)$ lies entirely on one side of the origin (i.e., there is some line in the plane containing $\gamma_v(t)$, passing through the origin, which does not intersect $\gamma_v(t)$). It follows that, for any t where it is defined, the line containing $\gamma_v(t)$ and the origin intersects S only at $\gamma_v(t)$. Thus, since $\gamma_v(t)$ lies in a plane, there exists some real-valued function $f_v(t)$ such that the image of $\frac{\gamma_v(t)}{f_v(t)}$ is a line segment. Moreover, by the compactness of the unit circle, there is some v such that the length of that line segment is minimized; thus we can suppose without loss of generality (again, possibly by shrinking V) that this length is independent of v . Moreover, we can reparametrize $\gamma_v(t)$, if necessary, in terms of distance along that line segment.

Let L be a line segment of this length. Then I claim that V is a generalized action on L , mediated by $f(t, v) = f_v(t)$. It is clear that, for any fixed v , there exists a unit quaternion $q(v)$ such that qLq^c is parallel to and has the same length as $f_v(t)$; thus there also exists a unit quaternion $r(v)$ such that $qLq^c + r$ coincides with $f_v(t)$. It follows that $V = X \odot_f L$, where f and L are as above, and $X = (q(v), q(v), r(v))$. Setting $U_p = V$ yields the desired result. ■

Finally, then, we can prove the theorem alluded to at the beginning of the section.

Theorem 10 *Let S be a compact differentiable surface in \mathbb{R}^3 . Then, up to piecewise constant translation, S is a generalized action.*

Proof : For any point $p \in S$, there is some $c_p \in \mathbb{R}^3$ such that the ray from $p + c_p$ to the origin is not tangent to $S + c_p$. Applying Lemma 1 to $S + c_p$, we obtain a neighborhood U_p of $p + c_p$ which is a generalized action. Set $V_p = U_p - c_p$; then V_p is

an open neighborhood of S . Since it is defined for each point $p \in S$, the collection of all the V_p forms an open cover of S ; by compactness, it has a finite subcover $\{V_{p_1}, \dots, V_{p_n}\}$. We know that each of the V_{p_i} is a generalized action up to constant translation; thus S is a generalized action up to piecewise constant translation. ■

This theorem then leads to two natural questions. First, we would like to know when a compact differentiable surface is a true generalized action. Second, we would like to know when it is a *continuous* generalized action (in the sense that f is continuous, and both A and X are connected).

Chapter 7

Quaternion Action and Inversion of Canal Surfaces

As promised in the previous chapter, we now return to the (non-generalized) quaternion action. Specifically, we will be examining surfaces formed by quaternion action on the unit circle. In particular, such surfaces include *canal surfaces* (which can be thought of as being surfaces formed by quaternion action on the unit circle, where each circle formed is normal to the *spine curve* formed by the centers of the circles). As we shall see, the inversion of any such surface remains a quaternion action; we will derive the formula for the resulting quaternion action. We then give an application of this material to the theory of blending by use of Dupin cyclids, which is useful in computer-aided geometric design.

The reason that the inversions of canal surfaces are still quaternion actions is simple, and lies in the following well-known fact about inversion, which we shall not prove:

Fact 6 *The inversion of a k -sphere in \mathbb{R}^n is another k -sphere, as long as the original k -sphere does not intersect the origin.*

It follows that inverting a family of k -spheres gives another such family. Since quaternion actions on the unit circle correspond precisely to the unions of collections of circles (1-spheres) in \mathbb{H} , it follows that they are preserved under inversion.

7.1 Inversion of a particular circle

We wish to determine, in general, how to write the inversion of a canal surface as a quaternion action, given a representation of the surface itself as a quaternion

action. In order to do so, we must determine what is obtained by inverting a particular circle. We will assume the circle does not intersect the origin, and thus that its inversion is in fact a circle.

We will in fact only solve a restricted version of the problem, in which the circle lies entirely within a 3-dimensional vector subspace of \mathbb{H} . In particular, note that this allows us to solve the problem for circles lying in $\text{Im } \mathbb{H}$, which we can identify with \mathbb{R}^3 .

Let C be such a circle, lying in a three-dimensional subspace V of \mathbb{H} . C possesses a 1-parameter family of symmetry planes which are orthogonal to the plane it lies in; all of these planes share a common line (namely, the line passing through the center of C , orthogonal to the plane containing C). Such a 1-parameter family of planes sweeps out all of V ; thus there is at least one such plane P which contains the origin.

Let $\{u, v\}$ be an orthonormal basis of P . Without loss of generality, we can set $u = 1$ (since rotating the circle by left-multiplying it by u^{-1} will not alter the geometric effects of inversion). As a plane passing through the origin, P is mapped to itself by inversion. It follows that P must also be a plane of symmetry of the inversion $\iota(C)$; thus P contains a diameter of $\iota(C)$. Since P only contains two points of C , it only contains two points of $\iota(C)$; those two points are therefore the endpoints of the diameter, and allow us to obtain the center and radius of C . To obtain its orientation, we use the fact that P is an axis of symmetry of $\iota(C)$; thus, in order to get from C to $\iota(C)$, only rotation in P , not the plane orthogonal to it, is required.

Since P is a plane containing the real axis, it is isomorphic to \mathbb{C} as a subalgebra of \mathbb{H} . We can therefore apply complex arithmetic to determine the specific diameter of $\iota(C)$ mentioned above. Suppose the two points of C which lie in P are $z \pm w$. Then it is easy to see that their inversions are the points $\frac{\bar{z} \pm \bar{w}}{z^2 - w^2}$.

Furthermore, we know that $\iota(C)$ is symmetric about the plane P ; this fact and

the specification of the diameter in P are sufficient to determine it. But, since P contains 1, it is invariant under quaternion transformations in which the left and right quaternion multipliers are equal. It follows that

$$\iota(C) = qCq, \quad (7.1)$$

where q satisfies the relation $q^{-2} = \bar{z}^2 - \bar{w}^2$.

We now place the above calculation in the context of the quaternion action. Let C be the unit circle in the \mathbf{ij} -plane, and consider the action $(r, q_l, q_r, q_l q_o q_r) \odot C$ of a point on C , where $q_o \in \text{Im } \mathbb{H}$. Geometrically, we are taking C , translating it by q_o , and then rotating and dilating it in some arbitrary way. It is clear that any circle of the form we have been considering can be written as such an action, and conversely. Moreover, in order to determine the inversion of the resulting circle, it suffices to determine the inversion of the circle $C + q_o$, for the reasons discussed in the last section.

Suppose $q_o = a\mathbf{i} + b\mathbf{j} + c\mathbf{k}$. Then any quaternion multiplication which brings $a\mathbf{i} + b\mathbf{j}$ into the real line while holding \mathbf{k} fixed will cause its symmetry plane containing the origin to also contain the real line (since the symmetry plane of C is the vertical plane containing $a\mathbf{i} + b\mathbf{j}$). In particular, then, consider a quaternion v satisfying $v^2 = -a\mathbf{i} - b\mathbf{j}$. We must left and right multiply by v^{-1} to obtain the inversion of the original circle.

The circle we are inverting has center $vq_o v$. The two points along its diameter which lie in the desired symmetry plane will be $v(q_o \pm (\frac{a\mathbf{i} + b\mathbf{j}}{\sqrt{a^2 + b^2}}))v$. It follows, by the above calculation, that the new circle is obtained from the old one by left and right multiplication by a quaternion q which satisfies the relation

$$q^{-2} = \overline{(vq_o v)^2} - \frac{\overline{(v(a\mathbf{i} + b\mathbf{j})v)^2}}{a^2 + b^2}. \quad (7.2)$$

Thus the inversion of the circle $q_o + C$ is given by

$$\iota(q_o + C) = v^{-1}qv(q_o + C)vqv^{-1}, \quad (7.3)$$

where v and q are as defined above.

7.2 Inversion of canal surfaces and specific examples

The calculations in the previous section motivate the following definition:

Definition 10 Let $p = (r, q_l, q_r, q_l q_o \bar{q}_r)$ be a point in $\mathbb{R} \times \mathbb{H}^3$, with r, q_l, q_r nonzero. We define the virtual inversion of p , denoted $\text{vi}(p)$, to be the point

$$\text{vi}(p) = \left(\frac{1}{r}, \bar{q}_r^{-1} v^{-1} q v, v q v^{-1} \bar{q}_l^{-1}, \bar{q}_r^{-1} v^{-1} q v q_o v q v^{-1} \bar{q}_l^{-1} \right). \quad (7.4)$$

Here v satisfies the relation $v^2 = -a\mathbf{i} - b\mathbf{j}$, where $q_o = a\mathbf{i} + b\mathbf{j} + c\mathbf{k}$, and q satisfies the relation $q^{-2} = \overline{(v q_o v)}^2 - \frac{(v(a\mathbf{i} + b\mathbf{j})v)^2}{a^2 + b^2}$. For a subset A of $\mathbb{R} \times \mathbb{H}^3$, we write $\text{vi}(A)$ for the set $\{\text{vi}(p) | p \in A\}$.

In particular, given this definition, it can be seen that the calculations in the previous section prove the following theorem:

Theorem 11 Let C be the unit circle in the \mathbf{ij} -plane. Consider a quaternion action $A \odot C$ which is purely imaginary. The inversion of the resulting set about the origin in $\text{Im } \mathbb{H}$ is given by the quaternion action $\text{vi}(A) \odot C$.

A particular canal surface which can be naturally thought of in this manner is the Dupin cyclide, which is the inversion of a torus. Dupin cyclides arise in computer-aided geometric design as surfaces which can be used to *blend* two cones: that is, to smoothly interpolate between them. The above argument can be used to construct the quaternion action representation of a Dupin cyclide.

A plot of an entire Dupin cyclide, and an illustration of a pair of cyclide fragments of the sort used when blending cones, are shown in Figures 7.2 and 7.2.

torus.nb

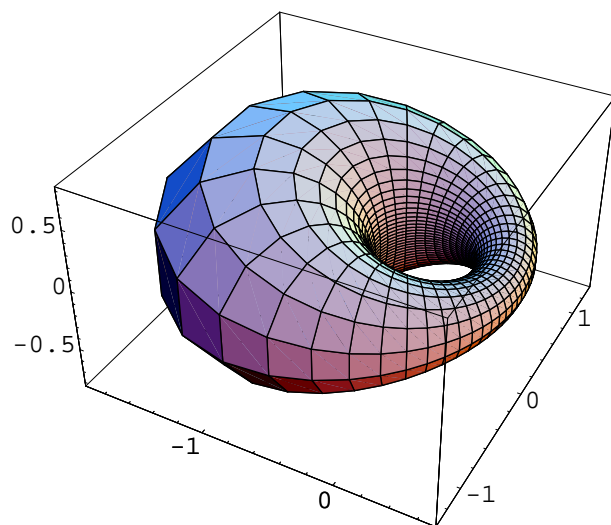


Figure 7.1: A typical Dupin cyclide

torus.nb

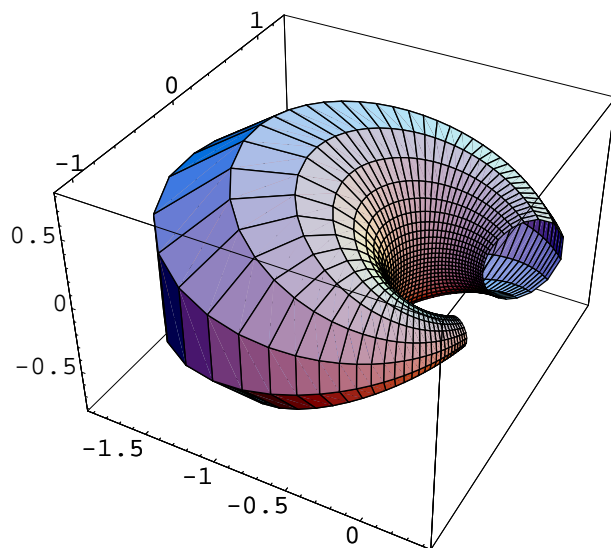


Figure 7.2: A surface which blends two cones

7.3 *The virtual inversion*

So far, we have been using the virtual inversion simply as a tool for calculating the inversion of a canal surface. However, it has other uses. For example, if we are given a quaternion action on a curve in the quaternions which is nearly a unit circle, the virtual inversion of the acting set will result in an approximation of the inversion of the surface.

A specific example of this is given by the virtual inversion used to construct the cyclid. If, instead of acting on a circle, we act with the same curve on arbitrary planar shapes, we achieve a surface which approximates to some degree the inversion of a surface of revolution with that shape as a cross-section; the closer the shape is to a unit circle, the better the approximation.

More importantly, we recall the application of Dupin cyclides mentioned above: they may be used to smoothly blend two (circular) cones. Using a surface constructed by the virtual inversion, we can therefore smoothly blend two cones over any single arbitrary shape.

There is one potential difficulty. Depending on the initial orientations of our two cones, it may be necessary to introduce some twisting into the surfaces, in order to ensure that they have the same orientation (and thus can be smoothly joined). However, this too can be circumvented using the quaternion action. All that is necessary is to introduce an appropriate compensating twist in the surface by rotating the figure in the plane before applying the quaternion action. Since we can apply arbitrary plane rotations to any such plane figure before continuing, it is in fact possible to construct blends of this sort using the quaternion action.

An example surface, which blends two cones over squares, is shown in Figure 7.3.

torus.nb

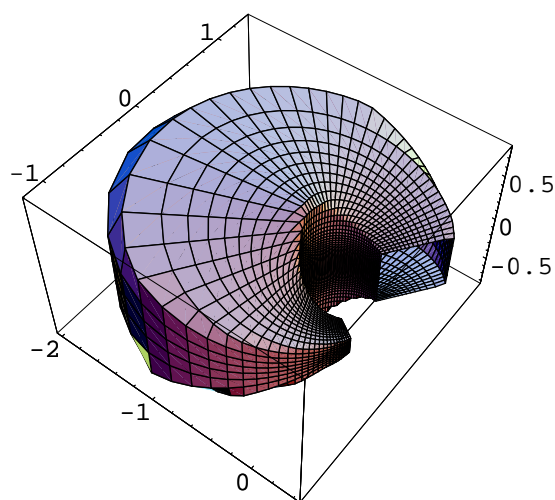


Figure 7.3: A surface which blends two pyramids (cones over squares)

Chapter 8

Future Work

Future work to be done on the quaternion action falls into two categories: the analysis and application of new constructions which use the quaternion action, and further work on refining the already presented applications. We will briefly discuss both.

8.1 New constructions

There are a number of possible uses of the quaternion action which have not yet been studied. Specifically, all the applications presented in this thesis have been to either pure mathematics or computer-aided geometric design. However, given the extreme ease with which we can construct many types of geometric objects using the quaternion action, it is highly likely that it has uses in other fields as well. In particular, we suggest the following two applications of the quaternion action to biological modeling:

- The DNA found in living cells forms complex, possibly highly knotted shapes; however, it maintains its standard double-helical structure. As a result, the geometric shape it describes is locally a double helix, which globally lies around some arbitrary spine curve. The existence of this simple, uniform local structure maintained within an extremely complex global structure suggests that the quaternion action could be used to model such supercoiled DNA. In particular, Hao and Olson [7] have used B-spline curves to model

the supercoiling of DNA; it is plausible that one could extend their results, using the Minkowski action and the techniques in Chapter 5.

- When modeling tumor growth and response to chemotherapy, the tumor is frequently modelled as a sphere of variable radius [8]. It is possible that a more accurate model could be obtained by designing the successive shape of the tumor as a quaternion action or generalized action.

8.2 Refining current constructions

We have presented a number of potentially useful constructions involving the quaternion action. Further research into these constructions include the following:

- We stated in Chapter 7 that, because the quaternion action was capable of accomplishing planar rotations, we could use it to construct blends between arbitrary surfaces. At the moment, however, the only method we know of for ensuring the smoothness of the surface is *ad hoc* trial and error. Developing a general method for performing this operation would be useful.
- The question of which surfaces can be produced as continuous non-self-intersecting generalized quaternion actions remains open. In particular, it seems clear that the topological restriction we placed on them is not sufficient; however, we as yet know of no examples of surfaces which demonstrate this.

Additionally, it would be useful to carry out the topological analysis in Chapter 3 in greater detail for the Minkowski product over the quaternions, as well as for the quaternion action, to the extent that doing so is possible.

Chapter 9

Conclusion

We have presented a very natural construction—the Minkowski product—as well as two of its generalizations, with broad application—the quaternion action and generalized quaternion action. As well as briefly analysing the topological implications of their definition, we also have given various constructive methods that can be used, along with the quaternion action and generalized quaternion action, to produce a number of interesting and potentially desirable surfaces.

We believe the quaternion action to be both theoretically elegant and useful. We hope that we have amply demonstrated it to be worthy of the attention the above questions, and other similar work, would warrant.

Bibliography

- [1] Berger, M. and B. Gostiaux (tr. Silvio Levy), *Differential Geometry: Manifolds, Curves, Surfaces*. Springer-Verlag, 1988.
- [2] Farouki, R. and W. Gu, *Minkowski Roots of Complex Sets*, *Geom. Modeling and Processing*, (2000) pp. 287-300.
- [3] Farouki, R., H. P. Moon, and B. Ravani, *Algorithms for Minkowski Products and Implicitly Defined Sets*, *Advances in Computational Mathematics*, 13 (2000) pp. 199-229.
- [4] Farouki, R., H. P. Moon, and B. Ravani, *Minkowski Geometric Algebra of Complex Sets*, *Geometriae Dedicata*, 85 (2001) pp. 283-315.
- [5] Farouki, R. and H. Pottmann, *Exact Minkowski products of N complex disks*, *Reliable Computing* 8 (2002) pp. 43-66.
- [6] Gallier, J. *Curves and Surfaces in Geometric Modeling*, Morgan Kaufmann Publishers, 2000.
- [7] Hao, M. and W. Olson, *Modeling DNA Supercoils and Knots with B-Spline Functions*, *Biopolymers*, 28 (1989) 873-900.
- [8] Jackson, T. and H. Byrne, *A mathematical model to study the effects of drug resistance on the response of solid tumors to chemotherapy*, *Mathematical Biosciences*, 164 (2000) 17-38.

- [9] Lu, W. and H. Pottmann, *Pipe Surfaces with Rational Spine Curve are Rational*, Comput. Aided Geom. Design, 13 (1996) pp. 621-628.
- [10] Peternell, M. and H. Pottmann, *Computing Rational Parametrizations of Canal Surfaces*, J. Symbolic Computation, 23 (1997) pp. 255-266.
- [11] Smukler, M. and W. Gu, *Topology of Minkowski Products of Plane Curves*, in preparation.
- [12] Ward, J. P. *Quaternions and Cayley Numbers: Algebra and Applications*. Kluwer Academic Publishers, 1997.
- [13] Weiner, I. and W. Gu, *Minkowski Geometric Algebra of Quaternion Sets*, Geometriae Dedicata, in submission.