Cylindrical Surface Pasting

Computer Science Department
University of Waterloo
Research Report CS-99-13

Stephen Mann and Teresa Pui-Yee Yeung

June 1, 1999

Abstract

The idea of hierarchical modeling is that many surfaces have different levels of detail and for modeling purposes it is useful to interactively edit these surfaces at any level of detail. Detail can be added to tensor-product B-spline via knot insertion, but once the detail has been added the the surface can not be edited at a lower level of detail. Hierarchical B-splines were designed to allow for a type of knot insertion method that still admits editing the surface at any level of detail. However, the details can not be rotated, nor is it convenient to create a library of such details. Surface pasting is a generalization of hierarchical B-splines that allows for the creation of a feature library, and where the pasted features can be rotated and scaled, and translated across the base surface.

Cylindrical pasting is a parametric-blending method that creates a smooth transition surface between a pair of B-spline surfaces that do not originally intersect. This blending surface is a deformed cylinder, and its creation is based on the surface pasting composition method, which adds detailed features to base surfaces by means of an efficient displacement method. In cylindrical pasting, a transition cylinder can be pasted on a NUBS surface or onto a NUBS cylinder. A displacement scheme is used to locate the control points of the blending cylinder to achieve approximate C^1 continuity between the boundaries of the base surfaces and the edges of the cylinders.

Acknowledgments

Many thanks to Richard Bartels and Kirk Haller for their discussions and input on the ideas that appear in this paper. Thanks to Blair Conrad for creating some of the images used in the appendix.

This work was funded by NSERC.

Trademarks

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Computer aided surface modeling and animation are two research areas in computer graphics. Designers continue to find new methods to computerize sophisticated smooth surfaces with fine surface control and adjustment, while animators are looking for new techniques to generate faster and smoother animation. Computer graphics has reached the state at which it is possible to integrate the results of a large amount of research and varieties of techniques into models by generalizing and abstracting the knowledge. This process is called *modeling.*

For character animation, smooth transition between frames and simpler models are recommended for easy control. However, in most cases sophisticated models like human beings are animated. In such models, separate surfaces must coexist as the model undergoes modifications while certain aspect criteria have to be preserved. For example, surfaces defining the joints between the upper and lower arms of a robot must remain smooth throughout any motion or stretching. This can be solved by using tensor product B-spline surfaces, which consist of internally aligned patches guaranteed to meet smoothly. However, enhancing the detail of a tensor product surface by using knot insertion (a common method for adding detail) to split patches results in unneeded control points and higher computational cost. By elaborating Forsey's *hierarchical B-splines surfaces*, *surface pasting* was developed for creating composite B-spline surface models. Surface pasting defines an approximate displacement mechanism that involves only the control points, adding detail only where needed [BF91, BBF95].

In the previous example of modeling a robot's arm, we can model the arm as three pieces, the lower arm, the upper arm, and the shoulder. However, when this arm is animated, gaps will appear at the joints as the arm is being transformed. Blending or filleting operations need to be employed here to assemble those components together. There are many filleting methods that blend two surfaces together in a smooth fashion. With the inspiration of basic pasting, we propose new a blending method in this paper, *Cylindrical Pasting*, that elaborates the domain mapping and displacement schemes of surface pasting, and applies it to cylinders onto NURBS base surfaces. We will focus on the mathematical details behind these operations rather than the user interface for modeling with these blends.

1.1 Overview of Chapters

Chapter 2 describes the background material involved in defining B-spline surfaces and cylinders, the basic surface pasting process. It also presents the ideas involved in blending parametric surfaces.

Chapter 3 discusses the techniques involved in performing cylindrical pasting, the focus of this paper. It also presents a *Sheet-On-Sheet* method that gives the general idea involved in the transition between surface pasting and cylindrical pasting. Mathematical details behind these operations, domain mapping and displacement schemes when a blending cylinder is being pasted on a base surface or cylinder will also be explained.

Chapter 4 introduces some attempts, improvements, and various schemes that have been made in the process of developing the final results of the cylindrical pasting method.

Chapter 5 summarizes the results of this work and suggests some further research.

An overview for the user interface components and a description of the functionality of `CylindersPaster` appear in the Appendix.

Chapter 2

Background

With surface pasting, detailed features can be added to a smooth surface without increasing the complexity of the base surface. The character of the feature surface's shape will be retained after pasting, and the topography of the base surface will also be reflected. An efficient displacement scheme is used to paste the features on the base surface in a multi-layered fashion, called hierarchical pasting [BBF95, Bar94]. Details of hierarchical pasting will not be discussed in this paper.

2.1 Surface Pasting

Tensor product B-spline surfaces play an important role in current surface design [Far94], especially in surface pasting. This section presents the basic concepts of a B-spline surface, the derivation of Greville displacement B-spline surfaces, and the definition of the basic pasting procedure.

2.1.1 An Overview of B-Spline Surfaces

In general, a three-dimensional tensor product B-spline surface is defined by a polygonal mesh of three-dimensional control points and two basis functions over a two-dimensional tensor product domain D. There are two kinds of spline surfaces: *point splines* and *vector splines*, which are explained in the following section.

2.1.2 Point and Vector B-Splines

A parametric tensor product B-spline surface (point B-spline surface) $S(u,v)$ is a linear combination of control points $P_{i,j}$ and two basis functions $B_i^k(u)$ and $B_j^\ell(v)$ defined in each parametric direction u and v respectively:

$$S(u,v) = \sum_i \sum_j P_{i,j} B_i^k(u) B_j^\ell(v), \tag{2.1}$$

where k and ℓ are the orders of the first and second basis function respectively. Each basis function, B_k^d, is a piecewise polynomial of order d (degree $d-1$) defined as

$$\begin{aligned} B_k^1(u) &= \begin{cases} 1, & \text{if } u_k \le u \le u_{k+1} \\ 0, & \text{otherwise} \end{cases} \\ B_k^d(u) &= \frac{u-u_k}{u_{k+d-1}-u_k} B_k^{d-1}(u) + \frac{u_{k+d}-u}{u_{k+d}-u_{k+1}} B_{k+1}^{d-1}(u), \end{aligned} \tag{2.2}$$

where $\{u_0, \ldots, u_N\}$ is a non-decreasing *knot sequence.*

The number of times a knot appears in the knot sequence is known as its *knot multiplicity.* If a knot has multiplicity equal to its degree, then the knot is said to have *full knot multiplicity*, and the B-spline curve will interpolate one of its control points. It is common to have full multiplicity of the first and last knots, with the resulting curve interpolating the first and last control points. In this paper, unless otherwise noted (or apparent in a diagram), we will use knot vectors where the first and last knots have full knot multiplicity.

For hierarchical surface modeling, it can be advantageous if the control points $P_{i,j}$ are expressed as vector offsets (displacements) $\mathbf{d}_{i,j}$ from an origin, giving *control vectors* and a *vector B-spline*, $\mathbf{s}(u,v)$ independent of the any coordinate frame:

$$\mathbf{s}(u,v) = \sum_i \sum_j \mathbf{d}_{i,j} B_i^k(u) B_j^\ell(v), \tag{2.3}$$

A *point spline* (Equation 2.1) defined in a frame with origin $O \in \mathcal{R}^3$, can be expressed as a *vector spline* relative to O:

$$\begin{aligned} \mathbf{s}(u,v) &= O + \sum_i \sum_j \mathbf{d}_{i,j} B_i^k(u) B_j^\ell(v) \\ &= \sum_i \sum_j (O + \mathbf{d}_{i,j}) B_i^k(u) B_j^\ell(v) \end{aligned} \tag{2.4}$$

As the B-spline surface $S(u,v)$ is rooted at origin O, each of the control points $P_{i,j}$ can be expressed relative to O and its offset $\mathbf{d}_{i,j}$, where

$$P_{i,j} = O + \mathbf{p}_{i,j} \tag{2.5}$$

Hence, a *point spline* (Equation 2.1) can be rewritten as a point-vector sum of the origin and the *displacement spline* (Equation 2.3) relative to an origin O:

$$S(u,v) \;=\; O+\mathbf{s}(u,v) \tag{2.6}$$

2.1.3 The Diffuse Coordinate System

Since the B-spline surface needs to be transformed from space to space without changing its topology in surface pasting, and the pasting operations are largely based on a displacement scheme, it will be efficient to take the advantage of the vector spline form, and to express the surface relative to a local coordinate system, called a *diffuse coordinate system* (*DCS*).

Under a DCS, each control vector is associated with a local frame $\mathcal{F}_{i,j} \in DCS$ at origin $O_{i,j}$ with an offset vector $\mathbf{d}_{i,j}$ assigned to it. Using this frame, the *point B-spline* defined in Equation 2.1 can be rewritten as:

$$S(u,v) \;=\; \sum_i \sum_j (O_{i,j}+\mathbf{d}_{i,j})B_i^k(u)B_j^\ell(v), \tag{2.7}$$

Choosing the appropriate $O_{i,j}$ and $\mathbf{d}_{i,j}$ improves the quality of pasting process. The way of finding these appropriate values for constructing DCS is based on the underlying domain of the B-spline surface which will be discussed in the next section.

2.1.4 Greville Displacement B-Splines

The domain of a B-spline surface is two-dimensional and is determined by two nondecreasing knot sequences, $u_0, \cdots, u_{\mu-1}$ and $v_0, \cdots, v_{\nu-1}$, where μ and ν are the number of knots in each parametric direction. A non-uniform B-spline surface (NUBS) has two basis functions $B_k(u)$ and $B_l(v)$ as defined in Equations 2.2 over domain the real plane.

If the B-spline surface has $m \times n$ control points ($m < \mu$ and $n < \nu$), the surface is non-zero over a rectangular region D:

$$D \;=\; \{(u,v) \mid u_{k-1} \le u \le u_{\mu-k},\;\; v_{\ell-1} \le v \le v_{\nu-k}\} \tag{2.8}$$

For a NUBS-spline surface, the basis functions in one variable are defined so that each control point $P_{i,j}$ is influenced by a number of knots that is the same as the degree of the basis. For each control point $P_{i,j}$, there is an associated *Greville domain point* $\gamma_{i,j}$ defined as a pair of *Greville abscissae*, one for each parametric direction:

$$\gamma_{i,j} \;=\; (\gamma_i, \gamma_j)$$

where the *Greville abscissa* is the average of 'degree' d knots:

$$\gamma_i = \frac{1}{d}(u_{i+1} + \cdots + u_{i+d}), \quad i = 0, \cdots, m-1 \tag{2.9}$$

$$\gamma_j = \frac{1}{d}(v_{j+1} + \cdots + v_{j+d}), \quad j = 0, \cdots, n-1 \tag{2.10}$$

The image of the Greville domain point on the surface $S(\gamma_i, \gamma_j)$ is the surface point that is maximally affected by the movement of $P_{i,j}$. Each Greville point is embedded into the x-y plane of a three-dimensional space (3D space) by adding a z-coordinate of 0. The three-dimensional Greville point is denoted as

$$\Gamma_{i,j} = (\gamma_i, \gamma_j, 0)$$

A discrete *Greville domain* is formed from the array of Greville points, which are relative to the same space as the control points. Each control point $P_{i,j}$ can be rewritten as an offset from a local origin $\Gamma_{i,j}$ with $\mathbf{d}_{i,j} = P_{i,j} - \Gamma_{i,j}$ as the *Greville displacement.* Then the local frame $\mathcal{F}_{i,j}$ for $P_{i,j}$ is:

$$\mathcal{F}_{i,j} = \{\Gamma_{i,j}, \mathbf{d}_{i,j}\} \tag{2.11}$$

From Equation 2.3, a point B-spline can be represented in the diffuse form in terms of its Greville point, and is defined as a vectoral *Greville Displacement B-Spline*:

$$S(u,v) = \sum_i \sum_j (\Gamma_{i,j} + \mathbf{d}_{i,j}) B_i^k(u) B_j^\ell(v) \tag{2.12}$$

Equation 2.12 is illustrated in Figure 2.1, which demonstrates that the domain is embedded in 3-dimensional space, and the Greville displacement is drawn from the (i,j)th Greville point $\Gamma_{i,j}$ to the corresponding control point $P_{i,j}$.

2.1.5 Pasting Greville Displacement B-Splines

To paste the feature surface S_F onto the base surface S_B, the feature domain D_F has to be transformed with an invertible transformation T to fit inside the base domain D_B. This transformation determines the location and the relative size of the feature surface with respect to the base surface.

Let the feature domain point be $(u,v) \in D_F$, with its mapped domain point as $(u',v') = T(u,v) \in D_S$. For each Greville point $\Gamma_{i,j} = (\gamma_i, \gamma_j)$, its image will be a base surface point $S_B(T(\gamma_i, \gamma_j)) = S_B(\gamma_i{}', \gamma_j{}')$, and is used as the origin of the displacement vector $\mathbf{d}_{i,j}$.

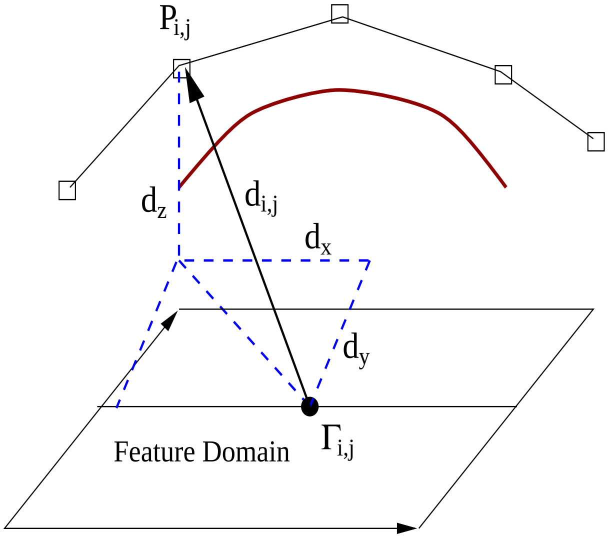

Figure 2.1: Greville displacement B-Spline

To apply the feature displacement vector $\mathbf{d}_{i,j}$ onto the mapped origin $S_B(\gamma_i{}', \gamma_j{}')$, a local coordinate frame $\mathcal{F}$ has to be set up at $S_B(\gamma_i{}', \gamma_j{}')$, where $\mathcal{F} = \{S_B(\gamma_i{}', \gamma_j{}'), \mathbf{v}(\mathbf{i}, \mathbf{j}, \mathbf{k})\}$. To reflect the orientation of the base surface at the mapped origin $S_B(\gamma_i{}', \gamma_j{}')$, the undetermined components of vector $\mathbf{v}$ are the tangent-plane vector components $\mathbf{i}$ and $\mathbf{j}$ and are given by the partial derivatives of base surface S_B in the $\mathbf{u}'$ and $\mathbf{v}'$ directions, $\left(\frac{\partial S_B}{\partial u'}\right)$ and $\left(\frac{\partial S_B}{\partial v'}\right)$, and

$$\mathbf{k} = \begin{cases} \mathbf{i} \times \mathbf{j}, & \text{for positive displacement} \\ \mathbf{j} \times \mathbf{i}, & \text{for negative displacement} \end{cases}$$

As the final result, the pasted image $P'_{i,j}$ was obtained from embedding the feature displacement $\mathbf{d}_{i,j}$ into the local coordinate frame $\mathcal{F}$. The pasting operation is simplified as a point-vector addition in terms of the given frame context:

$$S(u,v) = \sum_i \sum_j (\Gamma_{i,j} + \mathbf{d}_{i,j}) B_i^k(u) B_j^\ell(v)$$

The pasting process can be pictured as Figure 2.2

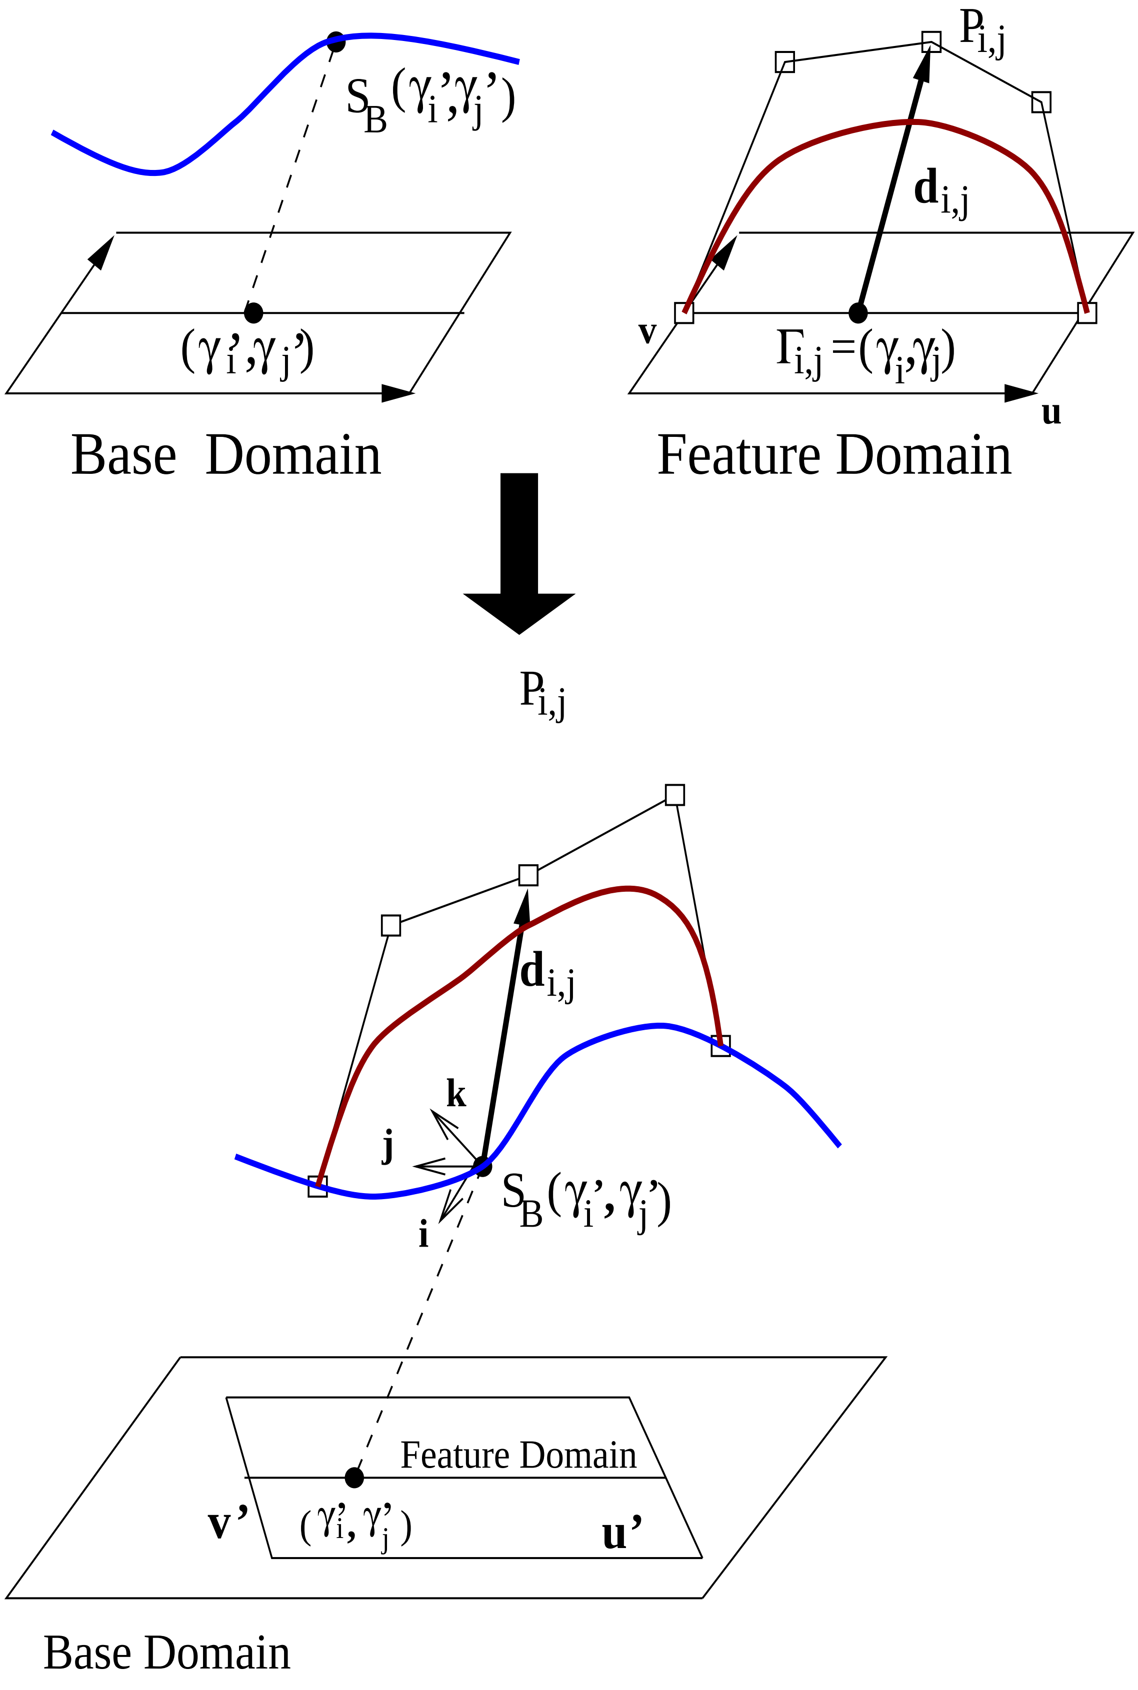

Figure 2.2: Pasting process

2.1.6 Surface Continuity

Splines are piecewise polynomials where the pieces meet with some level of continuity. The advantage of B-splines is that for degree n surfaces, the polynomials automatically meet with C^{n-1} continuity. For surface pasting, we want to match the boundary of the feature with a curve on the base, and match at least first derivatives of the feature and the base along the common curve. For tensor-product B-spline surfaces with full end-knot multiplicity (which is how we represent our features), the boundary row (or column) of control points completely determines the location of the boundary curve of the surface, and can be used to join surfaces with C^0 continuity. Further, the first two rows (columns) completely determine the first derivatives along the boundary curve. Thus, it is these two layers of control points that we will manipulate to achieve a smooth join between the feature and the base.

Surface pasting is only an approximation process as smooth attachment is not guaranteed at the intersection between the feature and base surface. To approximate C^0 continuity at the boundary, the displacement vectors for the control points on the edge are set to zero, which means the edge of the unpaste feature lies on the embedded domain, with the result that the pasted feature lies close to the base. Note that this distance between the boundary curve and the base surface can be made as small as desired by inserting knots in the feature surface.

To achieve approximate C^1 continuity between the feature surface and the base surface, first note that the derivatives of the feature boundary curve are approximately the same as the corresponding derivatives on the base surface. If we can construct the cross-boundary derivatives of the feature to be approximately the same as the cross boundary derivatives of the base, then we will achieve approximate C^1 continuity. In standard surface pasting, approximate C^1 continuity is achieved by setting the displacement vectors of the second inner layer of control points to zero. This results in a second layer of the feature's control points lying on the base surface, and the cross-boundary derivatives of the feature will be approximately the same as those of the base, and thus the feature will meet the base with approximate C^1 continuity.

Thus, with standard pasting the boundary of the feature approximately lies on the base, and the continuity between the pasted feature and the base surface is approximately the same as that between the unpasted feature and the embedded domain (x-y plane), given that the base surface has low curvature. Similar ideas will be used to achieve approximate C^1 continuity for cylindrical pasting, although we will use an improved method for setting the second layer of control points of the feature.

Note that in Figure 2.2, we only have an approximately C^0 join because while the displacement vectors for the end control points are zero, we have non-zero displacement vectors for the control points one in from the ends.

2.2 Blending

Blending is an operation of creating smooth transitions between a pair of adjacent surfaces. Accordingly, the transition surface is simply called a *blend* or a *blending surface*. Blending methods that use parametric surfaces are the most popular techniques. Martin, Vida, and Varady have published a survey of different blending methods using parametric surfaces that clarifies the nature of blending and the relationships between various parametric blending methods [VMV94].

In this paper, the main focus will be on *local parametric-blending* methods. In particular, we use a *trimline-based blend* as the basic idea for *Cylindrical Pasting*. In the following, a brief summary of the most important ideas in parametric blending is given. The different terms used to explain concepts in the blending literature are given and Figure 2.3 can be used as a guide.

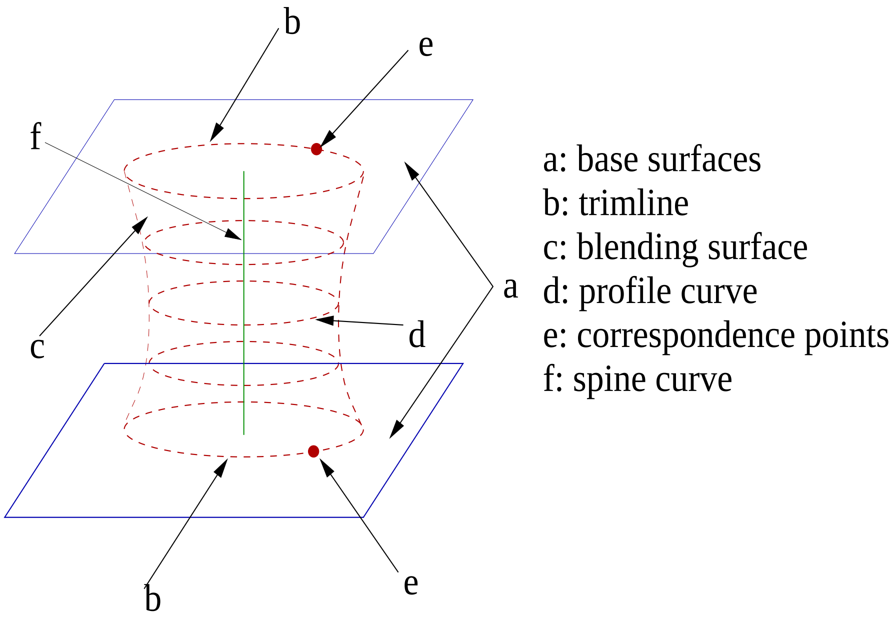

Figure 2.3: Terminology

The surfaces to be joined smoothly (the surfaces being blended) are called *base surfaces.* The curve that forms the common boundary of a base surface and the blend surface is called a *trimline*, at which the base surfaces are trimmed back along these curves. In general, the blending surface is created as a surface or volume swept along a given longitudinal trajectory, which is called the *spine curve.* At each point of the spine, a cross-sectional *profile curve* is associated with it that locally defines the shape of the blend and is usually defined as a planar curve. A profile can be constant or varying along the spine, and can be symmetric or asymmetric, and can be defined as a circular or free-form arc. Having two trimlines, a corresponding point pair, one point from each trimline, can be joined by a profile curve. Correspondences between these pairs of points need to be established by the *assignment* process.

Cylindrical Pasting is based on the idea of the *trimline-based* method, which are a class of techniques where an auxiliary spine is generated from the two trimlines, mainly for the purposes of *assignment* and the creation of *profile curves.* Since we know that blending replaces parts of the base surfaces with blending surfaces, one obvious way of specifying such operation is to decide explicitly which parts are to be substituted by choosing where the trimlines should lie on the base surfaces. Once a pair of trimlines has been chosen, a spine curve is used to choose corresponding points on the trimlines to be assigned together. The final important phase of trimline-based methods is a method of generating profile information that makes it possible to define the profile curves that connect assigned pairs of trimline points and contribute to the blending surface. What pasting provides is a method for joining the blend surface smoothly to the base surfaces.

2.3 Open Inventor and the Splines Library

The implementation for this research idea is based on two software packages: Open Inventor and the Splines Library. Open Inventor is an object oriented library for interactive 3D graphics. It provides a high level rendering mechanism and building blocks for 3D user interfaces. All information about 3D objects is stored in a scene graph that consists of one or more nodes. Each node represents a geometry, property, or grouping object. Hierarchical scenes can also be rendered by adding nodes as children. Open Inventor also provides 3D manipulators, basic geometric objects. With the concept of C++ inheritance, user-defined geometric objects can be created.

The Splines Library is an object oriented library being developed at the Computer Graphics Laboratory, University of Waterloo [VB92]. It contains convenient object classes for developing a prototype spline application quickly, and a pasting class for surface pasting. Using some of the classes provided, a separate package for *Cylindrical Pasting* was created.

Chapter 3

Sheet-On-Sheet Pasting

In surface pasting, all the boundary control points of the feature surface lie on the base surface with zero displacement vectors, while the remaining points in general have non-zero displacement vectors. However, in *Sheet-on-Sheet* pasting, only one edge of feature will be pasted on the base. The remaining control points are pasted relative to the control points on this edge.

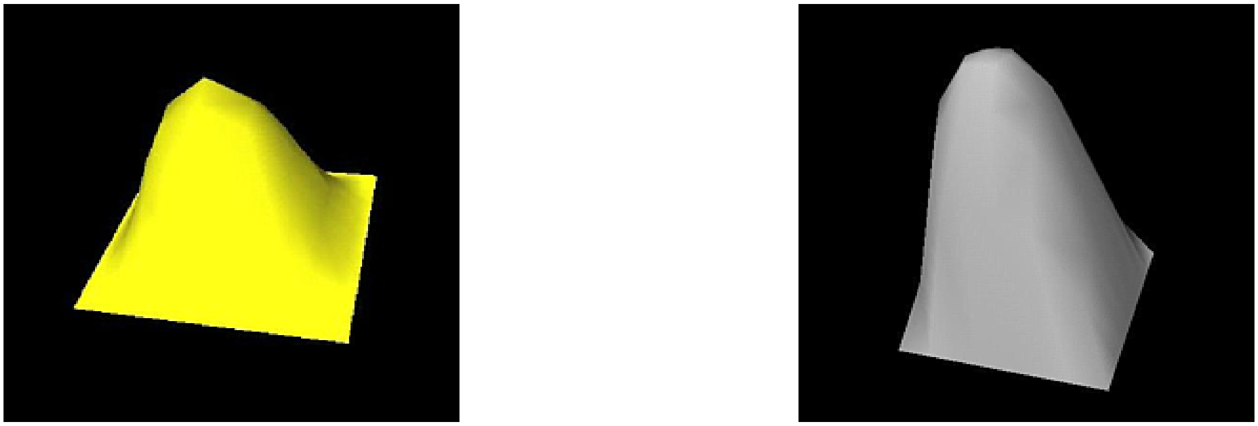

(a) Feature Surface

(b) Base Surface

Figure 3.1: B-spline surfaces for pasting

To show the difference between the standard pasting and the sheet-on-sheet pasting processes, a feature surface (Figure 3.1a) is pasted onto the same base surface (Figure 3.1b). Figure 3.2 shows the resulting image of the surface pasting process with the control points

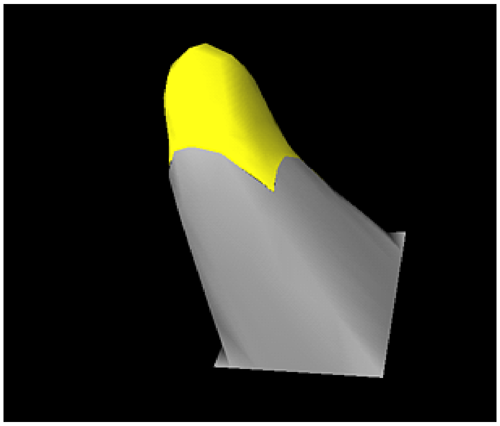

Figure 3.2: Surface pasting

(a) Closer Look

(b) Full View

Figure 3.3: Sheet-On-Sheet pasting

of the feature surface located on the base surface. Figure 3.3 shows two different views of the result for sheet-on-sheet pasting. 3.3(a) shows that only the first column of the feature surface are pasted onto the base surface using the Greville abscissa for locating those control points on the base surface. 3.3(b) shows the rest of the control points are displaced using the new displacement scheme, which will be discussed later in this chapter.

3.1 Domain Mapping

In standard surface pasting, rectangles are used to represent graphically the domains of the surfaces. The feature domain is drawn as a sub-rectangle of the base domain. In sheet-on-sheet pasting, domains are also represented by rectangles. However, only one edge of the feature domain lies in the base domain. We used a line to represent the location of the feature's pasted edge. This line is an oriented line, with the orientation indicating the correspondence between the features cross-boundary derivatives and directional derivatives of the base surface. We used an arrow emanating from the feature line to indicate its orientation.

Figure 3.4 shows the domain mappings for the two different pasting methods.

3.2 Displacement Scheme

In the preceding chapter, the process of finding a Greville Displacement B-spline was introduced. An overview of the displacement scheme in surface pasting is shown in Figure 3.5(a). This scheme will be re-used in sheet-on-sheet pasting, but only applied to the column of control points, $C_{i,j}$ to be pasted on the base. A different displacement scheme is required for the remaining control points to locate the final pasted feature surface.

For simplicity of discussion, we will assume that the edge corresponding to $C_{0,j}$ of the feature is pasted on the base. For sheet-on-sheet pasting, prior to the pasting process, a frame $\mathcal{F}_{0,j}$, with the Greville domain point as the origin and the standard basis as the coordinate vectors, is constructed for each $C_{0,j}$. The respective displacement vectors are set to zero for the control points on the feature edge touching the base. For the remaining points, $P_{i,j}$, their frames are constructed at the corresponding column control points at the same position i, j. And for the displacement vectors, it will be the distance between $P_{i,j}$ and $C_{0,j}$ with respect to the frame $F_{0,j}$. As a result, when the feature is being pasted on the base, the edge will

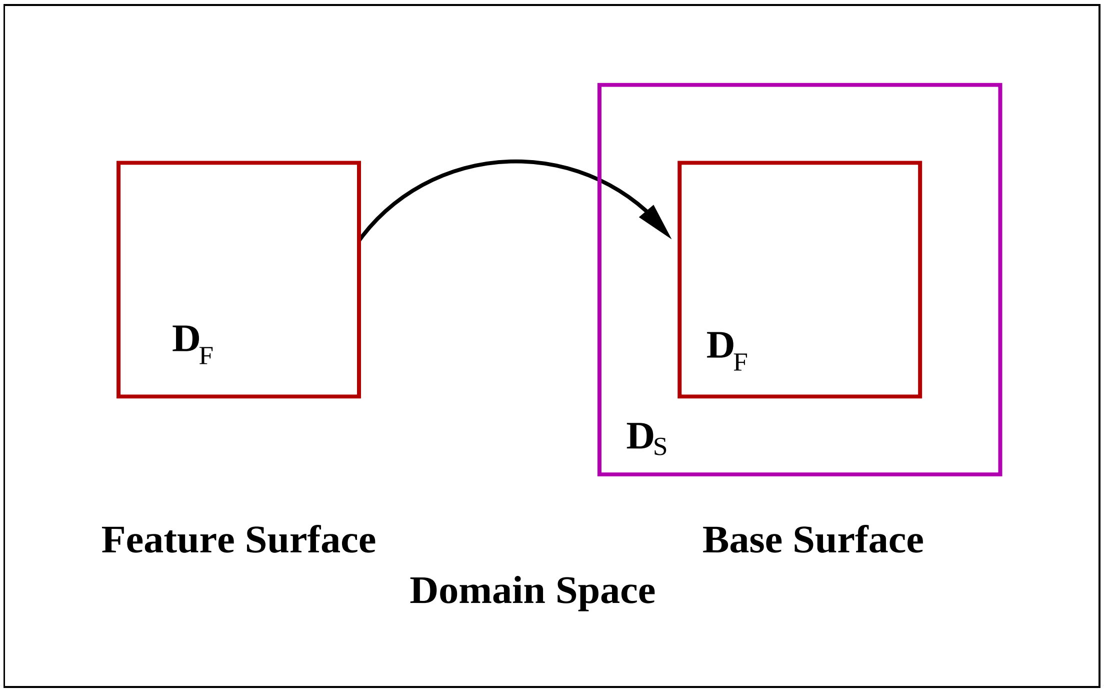

(a) Surface Pasting

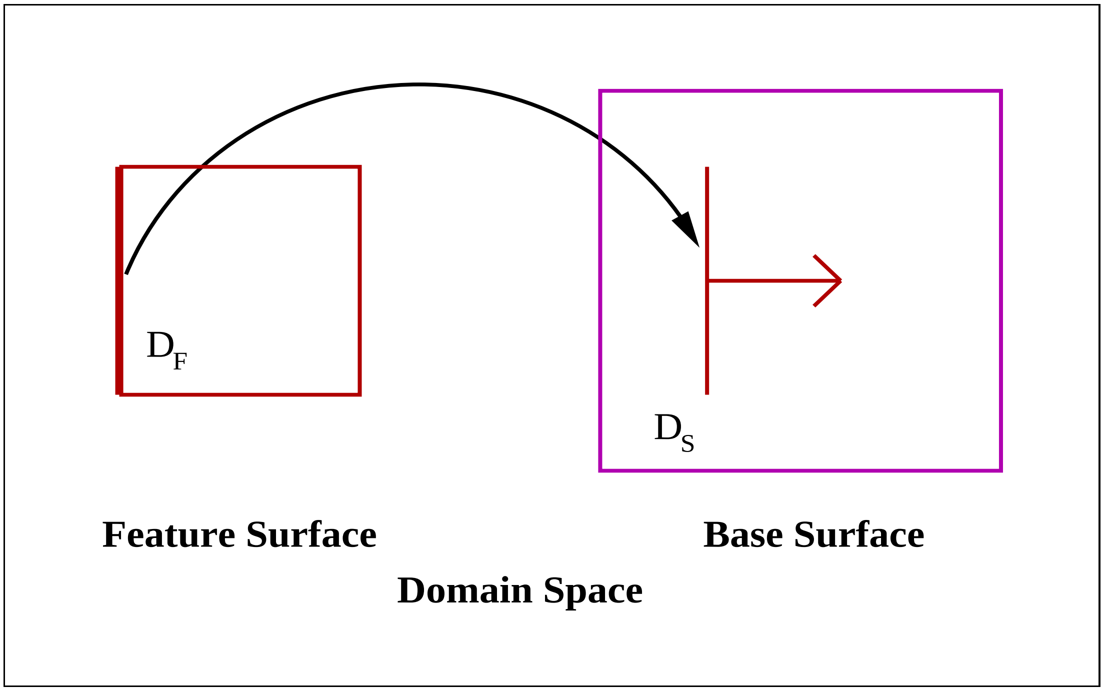

(b) Sheet-On-Sheet Pasting

Figure 3.4: Domain mappings

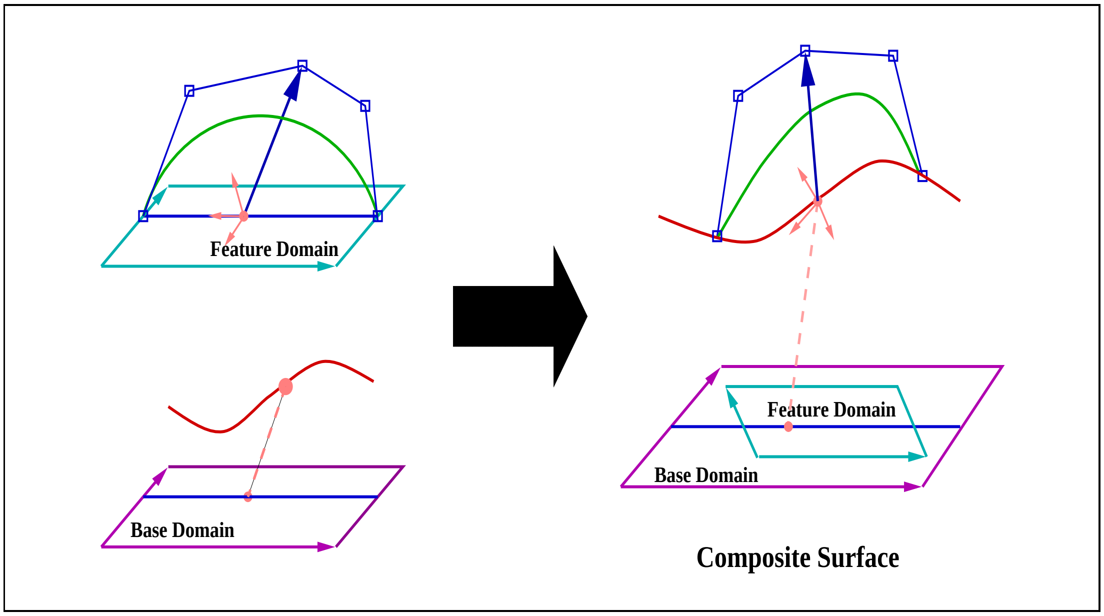

(a) Surface Pasting

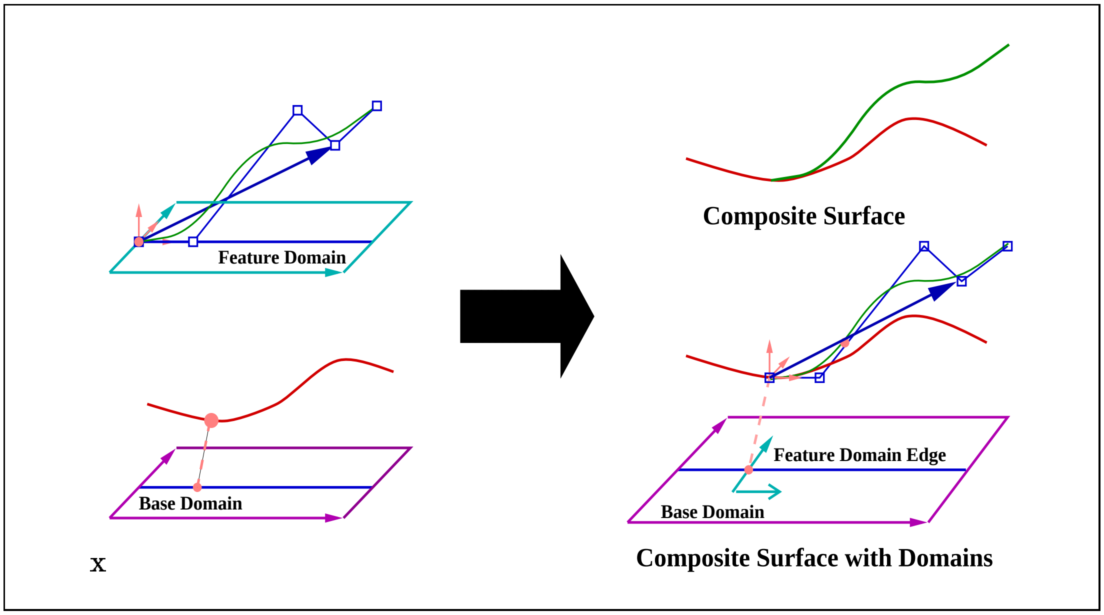

(b) Sheet-On-Sheet Pasting

Figure 3.5: Displacement scheme

lie on the base, and the remaining surface will be displaced from the pasted edge as shown in Figure 3.5(b).

3.3 Results

An example of sheet-on-sheet pasting can be seen in Figure 3.3. The feature surface appears to emanate from the base surface, and the two surfaces meet with approximate C^1 continuity. The result is a bit odd, because it is not a manifold. However, it does show that we can smoothly connect surfaces where only the edge of the feature is strongly associated with the base surface, which is the situation we have with cylindrical pasting. Note further that part of this method could be easily integrated into standard surface pasting. I.e., standard pasting sets a two layers of feature control points to have zero displacements and pastes them using the standard pasting process. In sheet-on-sheet pasting, we also have two layers with zero displacements, but we represent the second layer relative to the coordinate frames of the first layer. When we paste, the new method matches the second layer of control points to the cross-boundary derivatives of the base surface, and gives a better approximation to the cross-boundary vector field. This new construction for the cross-boundary derivatives that we used for sheet-on-sheet pasting would require only a simple change to standard pasting, and should give a reduction in the C^1 discontinuity between the base and feature.

Chapter 4

Cylindrical Pasting

Cylindrical Pasting is a surface modeling tool that integrates the techniques of parametric trimline-based blending into surface pasting, and creates a smooth transition cylinder between two base surfaces. We made modifications to the major pasting techniques, domain mapping and control points displacement, to adapt to the cylindrical pasting environment. All surfaces are parametric surfaces, and each can be considered as a mapping from a rectangular area of u-v parameter space into 3D space. The parametric representation for a cylinder is described below, followed by the progression of research ideas for performing cylindrical pasting.

4.1 Representing a Cylinder Using NUBS

The focus of this paper is to paste a cylinder onto selected B-spline surfaces, which can be a normal NUBS surface or a cylindrical NUBS surface. The mathematical form for representing a general B-spline surface has been previously introduced in Chapter 2 (Equation 2.1). To represent a cylinder using a non-uniform B-spline surface, we use the standard trick of identifying one of the edges of the domain rectangle with the opposite edge of the domain rectangle.

The cross section of a cylinder is a circular curve. Although circles cannot be represented using a NUBS curve, a NUBS curve can represent a closed curve that is a good approximation a circle. To represent a closed curve with a cubic NUBS, we set the last three control points to be the first three control points, with an appropriate setting of the knot vector. Figure 4.1 shows a closed curve of degree 3 with seven control points, which has points P_4, P_5, and P_6

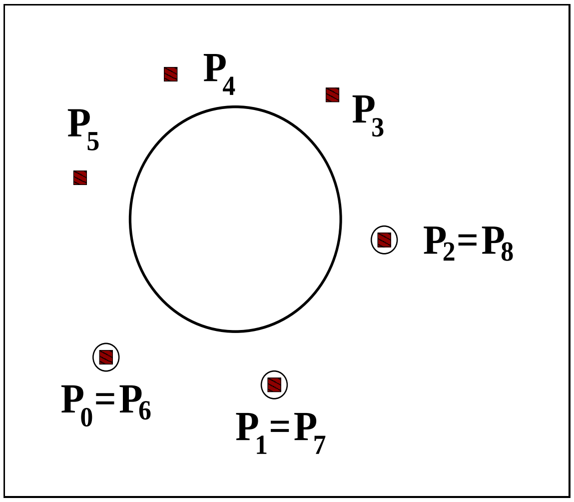

Figure 4.1: Circular base

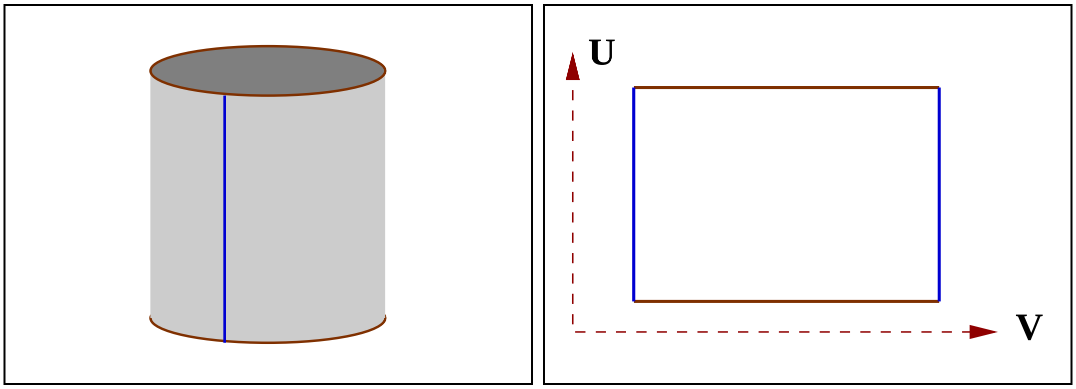

(a) World space

(b) Domain space

Figure 4.2: Representation of a cylinder in world space and domain space. In this figure, the blue (left and right) domain edges are associated with one another.

as the duplicates of P_0, P_1, and P_2, respectively. Here, we are assuming a knot vector whose end knots (five on each end) are all single multiplicity.

Mathematically, if we have a cubic B-spline with a knot vector $\{v_0, \ldots, v_N\}$ and control points $P_0, \ldots, P_{N-2}$, the following conditions must hold to get a closed curve:

$$P_0 = P_{N-5}, \qquad P_1 = P_{N-4}, \qquad P_2 = P_{N-3}, \tag{4.1}$$

and

$$\begin{aligned} v_1 - v_0 = v_{N-3} - v_{N-4}, \quad v_2 - v_1 = v_{N-2} - v_{N-3}, \\ v_3 - v_2 = v_{N-1} - v_{N-2}, \quad v_4 - v_3 = v_N - v_{N-1}. \end{aligned} \tag{4.2}$$

A general B-spline surface has a two-dimensional domain defined in two parametric directions, U and V. A cylinder is also represented by a rectangular domain where the U direction describes the number of circular layers of control points, while the V direction locates the control points on each circular layer. The representation of a cylinder in both world space and domain space is shown in Figure 4.2. Mathematically, a tensor product cylinder will have the form of Equation 2.1, where the V knot vector satisfies Equation 4.2 and for each i, the points $P_{i,j}$ satisfy Equation 4.1.

4.2 Mapping of Domain for Cylinders

The first step in performing basic pasting is to transform the domain of the feature surface to the domain of the base surface. The same procedure has to be done for *Cylindrical Pasting.* The rectangular domain of the blending cylinder has to be first mapped inside the domain of the base surface.

A cylinder can be pasted on two types of NUBS base surfaces: a sheet of normal NUBS surface, or a cylindrical NUBS surface. Depending on the type of the base surface, the rectangular domain of the blending cylinder will be transformed to the base domain in two different ways.

4.2.1 Normal NUBS Base Surface

In the first case, the base surface is a normal NUBS surface with a rectangular domain. Only one of the two edges of the blending cylinder will be pasted on the base in this situation, as shown in Figure 4.3(a). The best way to demonstrate the relative position of the blending cylinder on the base surface in the domain space is to transform the feature domain as a

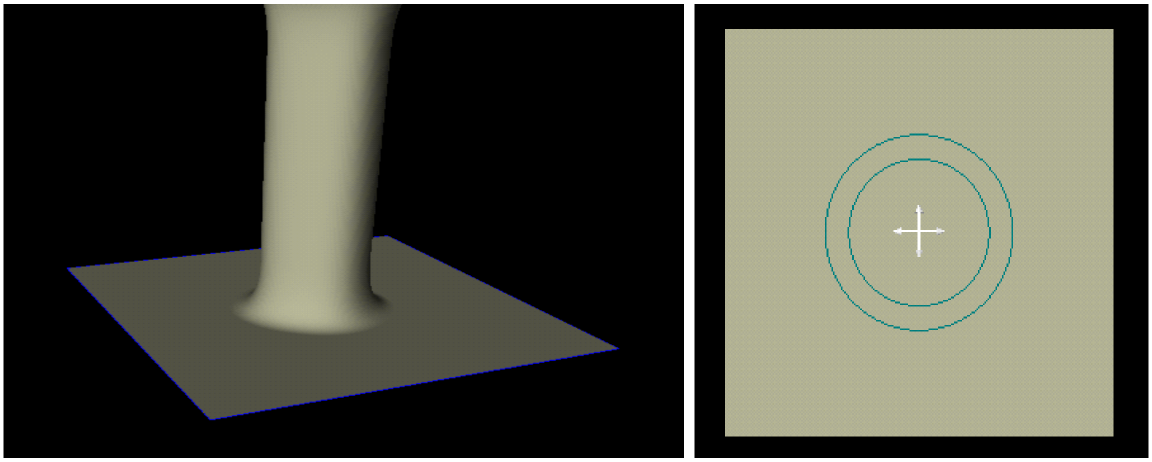

(a) World space (b) Domain space

Figure 4.3: A blending cylinder on a normal NUBS surface.

circle that represents the base of the cylinder, as shown in Figure 4.3(b). As have mentioned earlier, the control points on a single layer are represented in the parametric V direction, so the circular domain is the sole representation of the V parameter for a fixed value of U on the edge. This circular curve is represented as a NUBS curve as defined in Figure 4.1. By default, we initially locate the domain for the blending cylinder at the center of the base domain with a predefined radius.

4.2.2 Cylindrical Base Surface

In the second case, the base is a cylindrical NUBS surface that is the same type as the blending cylinder. Similarly, only one of the blending cylinder's edges is pasted on the base as illustrated in Figure 4.4(a), with the pink (top) cylinder as the base. The feature domain, however, will retain its rectangular shape but will be resized to fit inside the base rectangular domain. As shown in Figure 4.4(b), both of these domains are a 2D representation of the surfaces. With this representation, the height of the rectangle is analogous to the height of the cylinder, while the width relates to the edge of the cylinder. The two sides of the rectangle actually represent the seam of the cylinder when being flatten out. The overlapping distance between the base and blending cylinder, that is, the blue section in the domain picture (light on a black-and-white copy), is predefined, but can be modified by translating the blending

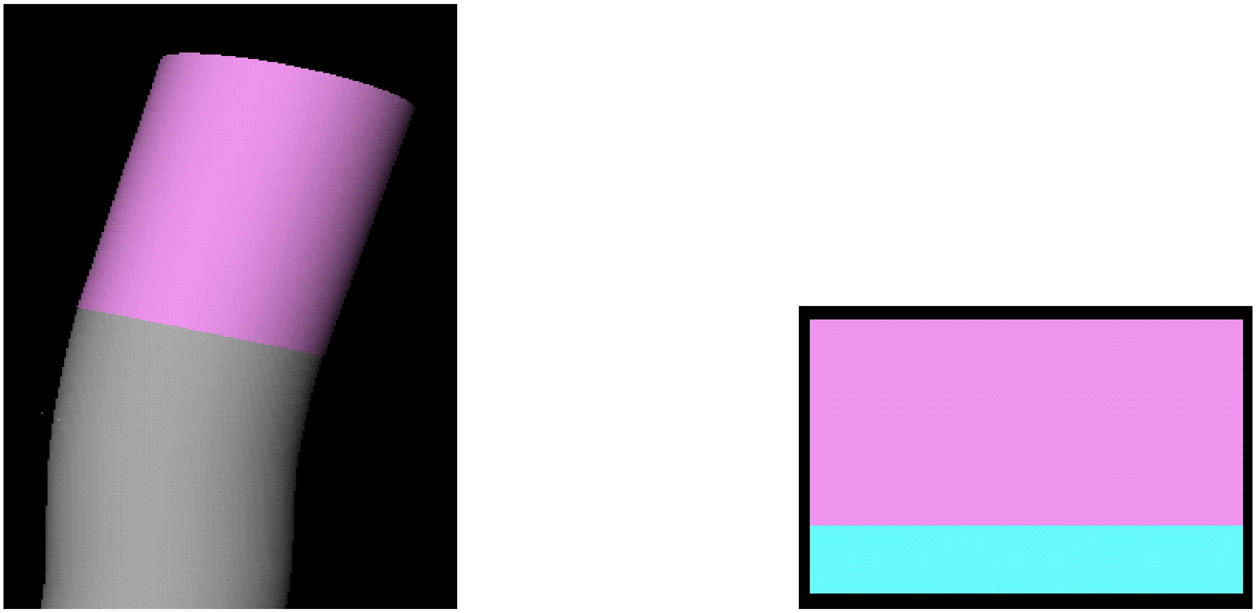

(a) World space

(b) Domain space

Figure 4.4: A blending cylinder on a cylindrical NUBS surface.

cylinder along the base cylinder.

4.3 Control Points Displacement Scheme

As the mechanism of *Cylindrical Pasting* is based on the idea of *trimline-based* blending method, the major issue is to determine the trimlines on the base surfaces. The body of the blend is constructed from the spine curve to be defined from the trimlines, and the profile curves along the spine curve.

In accordance with finding the trimlines, obtaining a smooth transition between the base and the feature surfaces is another major concern in surface pasting. As discussed in Chapter 2, pasting is an approximation process. With this concept, the continuity between these surfaces will be approximated at the boundary.

As described in Chapter 2, to paste the feature on the base surface, each control point of the feature has a respective Greville point for constructing a coordinate frame on the base and for calculating the final position relative to the base in the hierarchy. The cylinder is a "wrapped around" NUBS surface, with the last three control points as copies of the first three for each row. Hence, the last three pairs of Greville abscissae for row are duplicated. Unlike standard pasting, not every control point of the blending cylinder has a pair of Greville abscissae.

Only four rows of Greville abscissae are necessary to perform the pasting operation, two for each edge of the cylinder. Those two rows of control points will be named 0^{th} and 1^{st} layers in the world space, denoted as L_0 and L_1, respectively. L_0 is analogous to the *trimline* in the trimline-based blending method. A suitable connection scheme is needed to displace the control points of the blending cylinder from the base for constructing the blend. We have explored several ideas for getting a smooth connection at the boundary. In rest of this section, we will describe each of the attempts and the corresponding results. Moreover, the reasons why each initial method was not chosen, and the improvements made to achieve the final connection scheme will also be explained.

To investigate the visual effects produced by having different continuities at the boundary, we first implemented pasting with C^0 continuity. To obtain this approximation, the Greville displacement vectors for the L_0 layer control points have to be set to zero so that the pasted cylinder will be close to the base surface. The same procedure is done for both of the two edges of the unpasted cylinder. As mentioned in the *Domain Mapping* section, a predefined radius or distance is set for the domain of the blending cylinder. Pasting on a normal

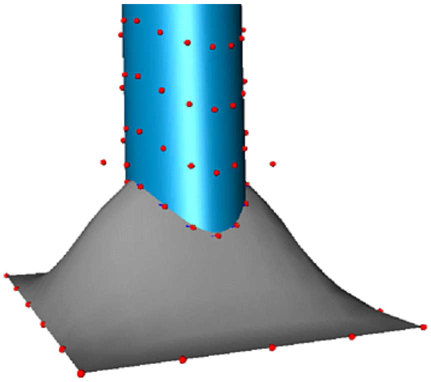

Figure 4.5: Pasting with C^0 continuity.

NUBS base surface is used for demonstration purposes in the following sections. A ring of control points is constructed at the center on the base with their locations derived from their respective Greville abscissae with zero displacements. Since we are interested in the behaviour of the edges, the remaining layers of control points are simply equally spaced between the two L_0 layers. The final image of the pasted cylinder is shown in Figure 4.5; only the lower half of the pasted cylinder appears in this figure.

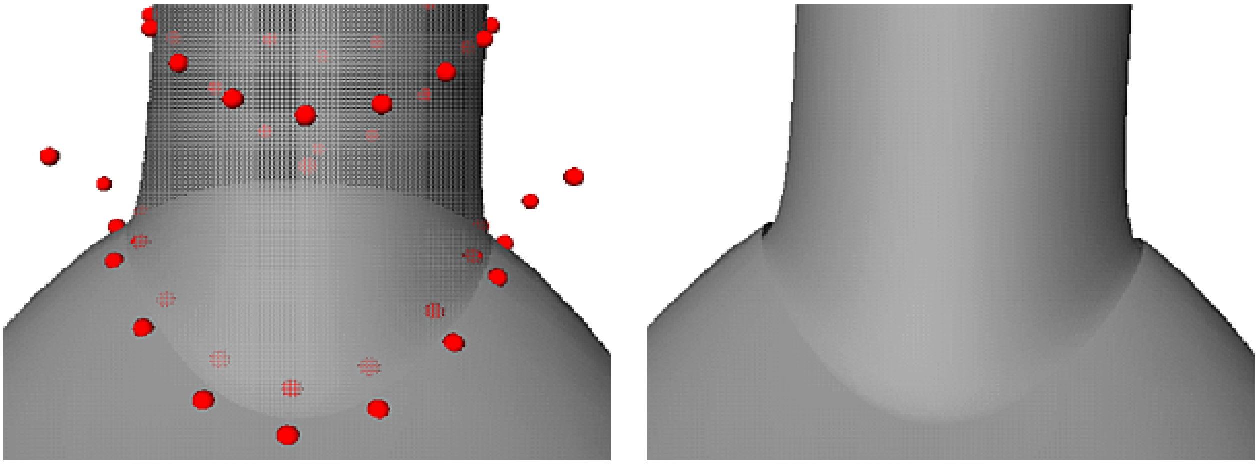

(a) Untrimmed base view

(b) Trimmed base view

Figure 4.6: Pasting with C^1 continuity.

From the picture, we can see that placing the L_0 layer control points on the base results in the cylinder's edge lying close to the base. As expected, it does not produce a smooth transition at the boundaries, since we did nothing to approximate C^1 continuity. Next, we extended the continuity to approximate C^1, which in standard pasting required the displacement vectors for the L_1 layer control points to be set to zero. With cylindrical pasting, we only have the boundaries of the cylinder's domain associated with the base domain, and it was unclear how to paste the L_1 layer of control points. So in our initial attempt, we placed a second circle in the base domain, with the same center but smaller radius than the first circle (Figure 4.3(b)). We used this circle to associate the second layer of the cylinder's control points with the base surface.

Initially, we used an arbitrarily domain distance between the L_0 and l_1 circles for simplicity, although the user is able to adjust the radius of both circles. These two circles gave us mappings for the L_0 and L_1 layers using the method of standard surface pasting. We set the remaining control points to be equally spaced between the two L_1 layers of control points for our initial test. The mapped control points of the L_0 and L_1 layers lie on the base as in Figure 4.6(a). To get a better feel for the C^1 discontinuity, we also trimmed the base surface (Figure 4.6(b)). In looking at these images, we discovered two problems.

The first problem is that while the boundary of the feature lies near the base, it is still unacceptably far from it as seen in Figure 4.6(b). This problem can be handled by inserting more knots in the V parametric direction. The second problem is more serious: the C^1 discontinuity is still unacceptably high (Figure 4.6(a)). Adding knots in the V parametric direction has no effect on this discontinuity, and adding knots in the U parametric direction yields only small improvements.

To reduce the C^1 discontinuity, we changed the method for pasting the L_1 layer of control points by extending the sheet-on-sheet pasting method described in Section 3.2. For cylindrical pasting, we pasted the L_0 layer of control points by using the method of standard pasting using zero length displacement vectors. Then, at each feature control point $C_{0,j}$ of L_0, we construct a coordinate frame $\mathcal{F}_{0,j}$, where $C_{0,j}$ is the origin, the unit derivative vector in the V direction is one coordinate vector, the derivative vector in the U direction is the second coordinate vector, with their cross-product forming the third derivative vector (Figure 4.7(a)). Each L_1 layer control point $P_{1,j}$ is then expressed as a displacement relative to frame $\mathcal{F}_{0,j}$.

Next, we map the L_0 layer to the embedded domain of the base surface, together with the corresponding frames (Figure 4.7(b)). Here, the two frame vectors that are tangent to the feature are mapped to the domain plane so that the basis vector tangent to the circle along the edge of the cylinder maps to be tangent to the circle in the base domain, and other tangent vector that lies in the tangent plane of the cylinder maps to be perpendicular to the circle, pointing inside the circle (the third basis vector is mapped parallel to the z-axis).

Note that with this new method the inner circle in the base domain of Figure 4.7(b) is no longer needed. However, we kept it in our user interface as a means of scaling the lengths of all the cross-boundary derivatives of the pasted cylinder.

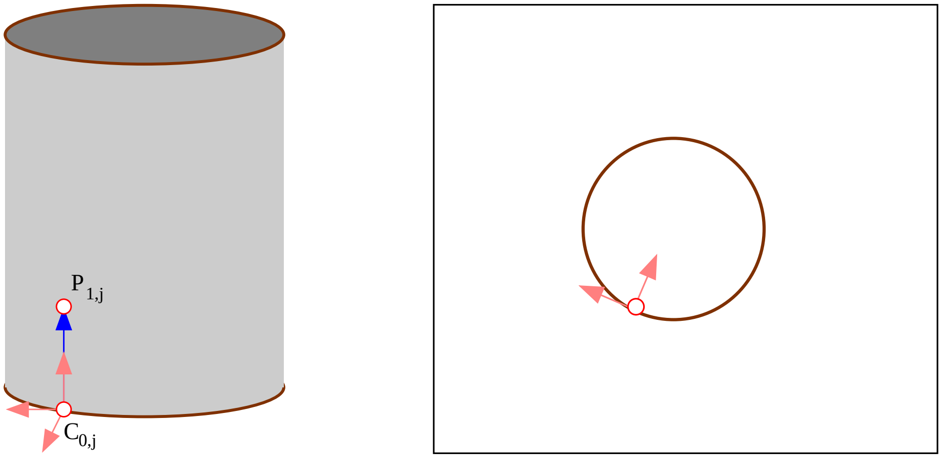

(a) Feature cylinder

(b) Base domain

Figure 4.7: Cylindrical displacement mapping

4.4 Simple Blends

After setting the L_0 and L_1 layers of control points, a method is needed to place the remaining control points for building the body of the cylinder. The method used for standard pasting does not work here, since there is no reasonable surface from which to offset. We tried several methods before settling on the spine method.

Our first approach was to linearly interpolate between the two L_1 layers, where our interpolation parameter was the j index of the control points. This approach has the problem that the blend flattens in the middle in some configurations (in the extreme, the center circle of control points all lie on a line as in Figure 4.8).

Our second approach was to set all the remaining control points as they are set for sheet-on-sheet pasting. Since the blend connects two surfaces, we actually have two settings for

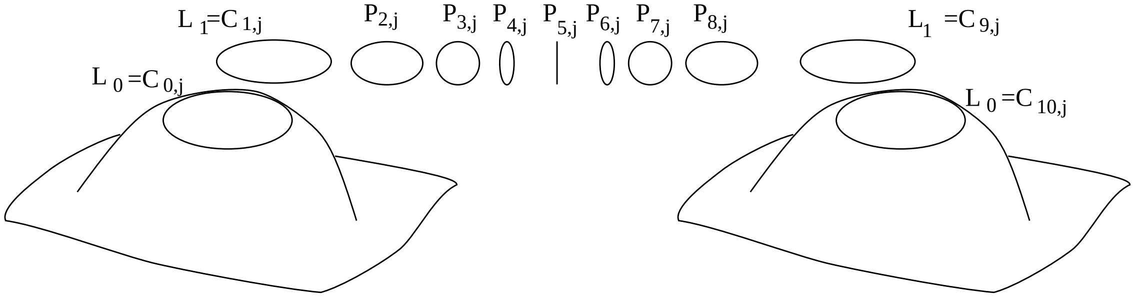

Figure 4.8: A linear blend becomes degenerate

each interior control point of the blending cylinder. Thus, we linearly interpolated between the two choices for each control point. The obvious problems with this approach arose: The interior of the cylinder sometimes became too narrow (indeed, the blend sometimes turned inside out), and sometimes the cylinder had sharp bends with a flat (and skewed) region (Figure 4.9) in the middle.

While these two illustrations are extreme, they show the short comings of using linear blends to find the center rings of control points. Further, they illustrate two different types of blends: The first being a U-shaped pipe, and the other being a simple connecting piece. As seen later, these two types of blends require different approaches when the solving correspondence problem (i.e., deciding how to connect js on the $C_{i,j}$ on the two L_1 layers). As well, we see that simple blends can result in skews in the layers (we obtain better shape if each ring of control points is roughly perpendicular to the curve through the centers of the rings). Thus, we want a method for finding interior points that has two properties: First, the curve through the center (the spine) should be well-shaped, and second, the rings of control points should be roughly perpendicular to the spine.

The other observation we make at this time is that the shape of the blending cylinder depends on the relative location of the two base surfaces. Figure 4.10 shows some scenarios of how the cylinder will bend on different base surfaces.

Finally, we considered using cubic Hermite splines to connect the L_i layers as Kim and Elber did [KE97]. However, we intend to use our method for both blending and for longer connecting pieces, and we found that their method gave us poor shape when used for longer connecting pieces. Note, however, that if our feature cylinder has only four layers (rings) of control points, then our method is constructing cubic Hermite splines almost exactly like their method does. We will discuss their method futher in the last chapter.

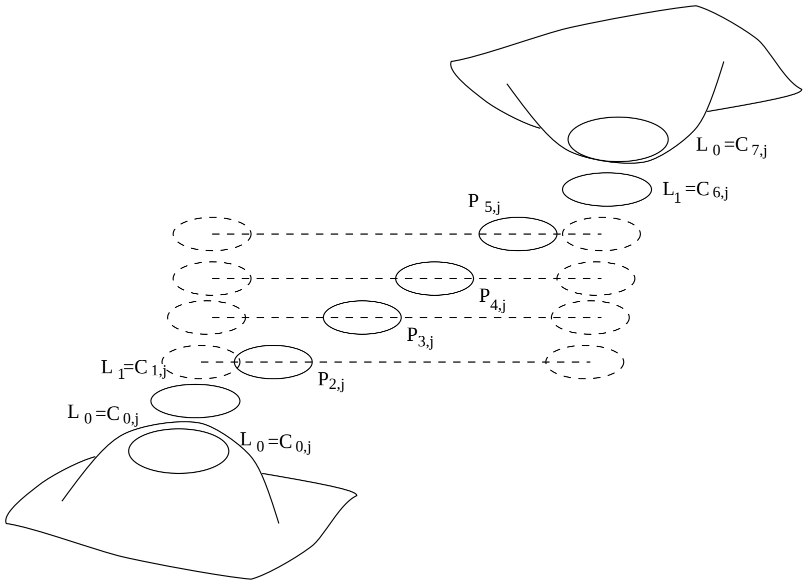

Figure 4.9: A linear blend has sharp bends and skewed, flat regions. The dashed ellipse are layers of control points extended via the sheet-on-sheet method; the connecting dashed lines show which two ellipses are blended, with the solid ellipse being the result of the blend.

(a) Scenario 1

(b) Scenario 2

Figure 4.10: Scenarios of different blending.

4.5 The Blending Spine

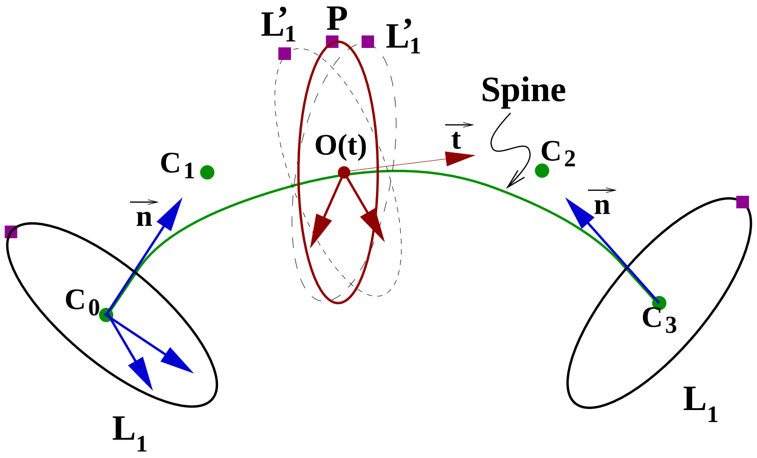

Figure 4.11: Method to determine an inbetween layer of control points.

To get a well-shaped blending cylinder, we constructed the interior control points around a curve, which we call the *spine*. This spine curve plays the role of the skeleton for the cylinder. It is a simple cubic Bézier curve defined by four control points: C_0, C_1, C_2 and C_3. Each of the two end points, C_0 and C_3, is the average point of corresponding L_1 layers of control points. We then construct a vector n at each point by summing the crossproducts with the surrounding points in the layer:

$$\begin{aligned} n_0 &= \sum_j (C_{1,j} - C_0) \times (C_{1,j+1} - C_0) \\ n_3 &= \sum_j (C_{n*,j} - C_3) \times (C_{n*,j+1} - C_0) \end{aligned}$$

where $j+1$ is taken modulo $n*$. The orientations of the n are set to point away from the surface.

Once we have C_0, C_3, n_0 and n_3, we locate the other two control points, C_1 and C_2 by placing the second control point along the normal vector with a certain distance apart from the end point will generate a simple cubic Bézier curve as shown in Figure 4.11. That distance is called the *curvature parameter* that determines the bending of the cylindrical body, and is made available as a shape parameter for the user.

To define the cylindrical body, each remaining layer of control points is built as a ring of points along the spine. We call these rings the *profile curves*, and they are represented as P in Figure 4.11. Basically, each of these profile curves is a linear interpolation of the mapped images of the two L_1 layers at both ends. These mapped images are denoted as L_1' in Figure 4.11.

The effect we want is for the left L_1 layer to sweep along the curve, gradually transforming into the right L_1. If we have $n+4$ layers of control points in our blending cylinder, then four are L_i layers, and we need n profile curves. To construct the L_1' and P layers, we first build a local coordinate frame at C_0 and one at C_3, and represent each of the control points in the L_1 layers relative to the corresponding frame. We then map these coordinate frames along the curve so that they are centered at a point $O(t)$ on the spine obtained by evaluating

$$O(t) \quad = \quad (1-t)^3C_0 + 3(1-t)^2C_1 + 3t^2(1-t)C_2 + t^3C_3 \tag{4.3}$$

at some set of t values.

While the values $t = \frac{i}{n+1}$, for $i = 1 \ldots n$ might seem like appropriate choices for sampling O, the Bézier curve described by Equation 4.3 is not arc-length parameterized. Thus, with uniform samplings of t, we get non-uniform samples on O, resulting in a blend surface with twists. To address that problem, we made an approximate arc-length parameterization of O by sampling O uniformly, computing the distance between these sample points, and using these distances to reparameterize the curve. The result is a close-to-arc-length parameterization, and rings of control points that are uniformly spaced over O.

Once we have the sample points on O, we need to map the L_1 layers to these sample points. We initially considered the idea of rotating L_1 along the spine curve with progressive degrees to get mapped images L_1' has been considered. Unfortunately, it is unclear how to find the appropriate degree variations for how much each L_1 should rotate to give the final profile to best represent the geometry of its base. Instead, we used a geometric transformation of $\vec{n}$, mapping $\vec{n}$ to the vector $\vec{t}$ tangent to the spine curve at $O(t)$. This gives the direction for locating the mapped coordinate frame derived from C_0, hence, the mapped control points can be used to locate L_1'. Applying the same process to L_1 at C_3, two mapped curves L_1' are obtained at $O(t)$. To obtain the final profile curve P that reflects the transition between the base surfaces, we applied linear interpolation on the generated L_1's.

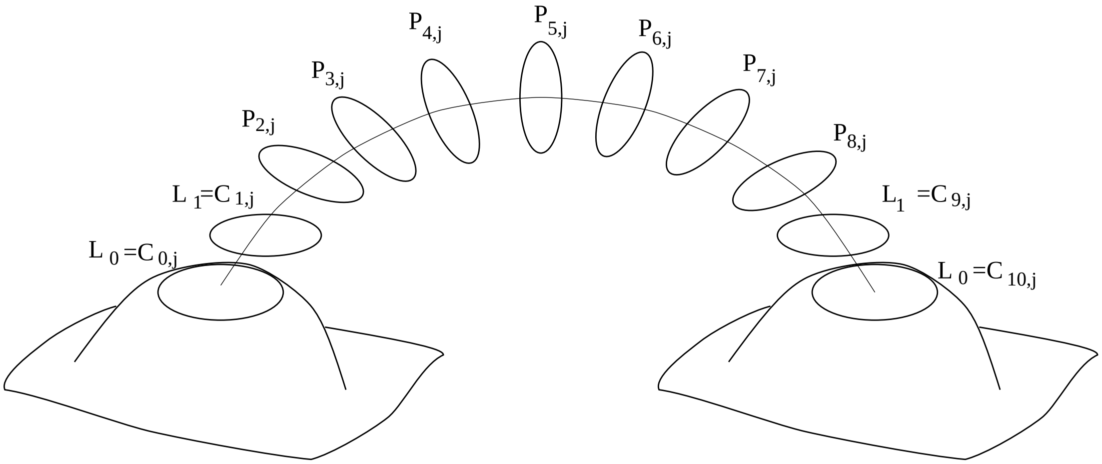

Figure 4.12: Good blend.

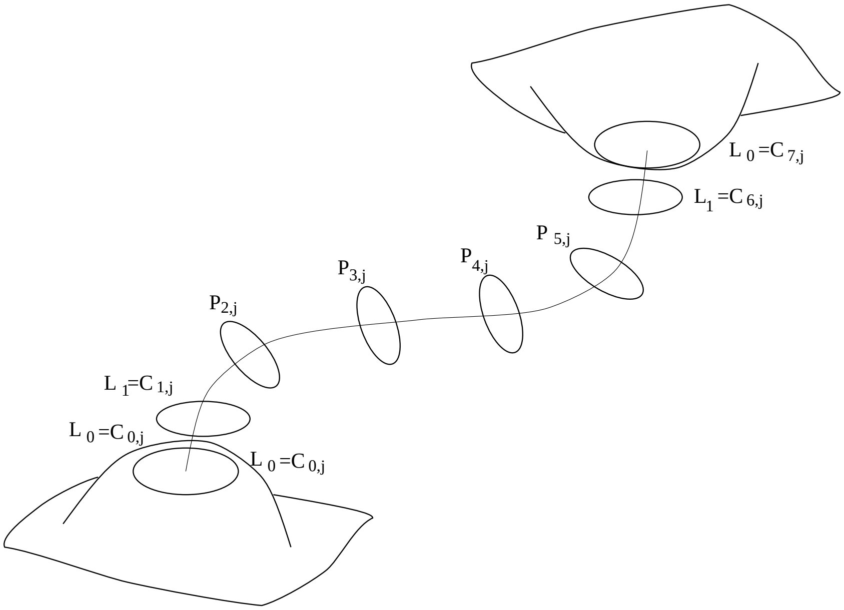

Figure 4.13: Good blend.

4.6 Correspondence Assignment Process

The cylindrical pasting process constructs the first two layers of control points on either end of the pasted cylinder by associating boundary control points with two edges of the cylinder domain, and then mapping these two edges into the base domain. The interior control points are found by transforming the second layer at either end of the cylinder along a curve, and then blending corresponding points, as described earlier in this chapter. However, one question that we have left unanswered is how to set up a correspondence between the two layers of blended control points.

The correspondence process is a non-trivial one. If we make a poor match between the two layers, then we introduce a twist in our blending surface. Because we were developing an initial prototype for cylindrical pasting, we used the following process (which we note is inadequate in general):

1. Using the Cartesian coordinate system for the range space and the normal to the plane approximating the each second layer of pasted control points, find the coordinate direction most perpendicular to each normal.

2. Select the control point within the layer whose coordinate relative to the selected axis is a maximum.

3. Repeat the selection process with the second layer at the other end of the pasted cylinder.

4. Associate the two selected control points, and associate the remaining control points within the layers starting from the selected control points.

As stated, the last step still leaves unspecified the direction around the ring of control points. We chose this direction arbitrarily, and allow the user to change the direction for either ring of control points.

Note that this method works reasonably well if the two layers have relative locations similar to that of Figure 4.13, but is a poor choice if the layers have relative locations similar to that of Figure 4.12. Using a better method for solving the correspondence problem has been left for future work.

4.7 Results

We tested our cylindrical pasting method by blending two surface. The first examples were shown in Figure 4.10. In these figures, the bottom surface is a plane, while the top surface is a curved surface. The plane provides an interesting test case since the pasting method for the boundary control points will result in the boundary of the feature meething the plane with C^1 continuity. Note, however, that once we trim the base, we will not have a C^0 join, since the feature boundary is not the trim curve.

In any case, in these images we see that cylindrical pasting has the desired effect. However, this example is not the expect case for blending surfaces. Commonly, we would expect the surfaces to be closer together, or the trim region on each surface to be closer in size to the distance between the two surfaces.

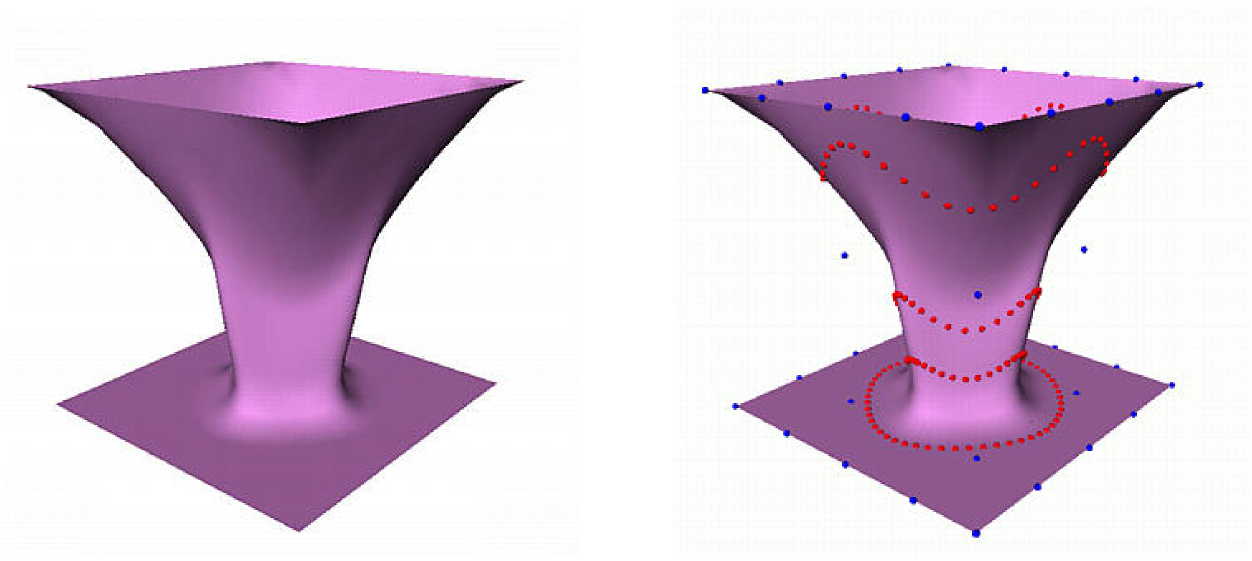

(a) Without control points

(b) With control points

Figure 4.14: Cylindrical pasting example.

A second example of cylindrical pasting can be seen in Figure 4.14. In this figure, we have increased the radius of the blend boundary in the two base domains. The result is that each trim curve moves closer to the boundary of its base surfaces, and we see that the pasted cylinder has taken on a "squarish" look. This is a result of the domain curve being a circle, which as seen in Figure 4.14(b) maps to a non-planar curve that also reflects the non-uniform parameterization of the base surface. This non-uniform parameterization results in the displeasing shape of the blending cylinder. This example points to the need for a differently shaped trim curve in the domain.

Chapter 5

Conclusion

Cylindrical Pasting is a connection method to join two surfaces with a cylindrical blend. The process involves selecting domain regions in the two surfaces to connect together, mapping the first two layers of control points on either end of the cylinder to these two surfaces, connecting the two regions with a spline curve, and using this curve to set the remaining control points of the pasted cylinder.

The results in this paper are fairly preliminary. Essentially, this work is a proof-of-concept. Surfaces can be blended using cylindrical pasting, but the images show that further research remains. In addition, this work has shown a method for improving standard pasting. In standard pasting, the construction of the cross-boundary derivatives of the feature is performed by mapping feature control points onto the base surface. If instead we map the cross-boundary derivative vectors of the feature surface onto cross-boundary derivatives of the base, then the feature and base meet with a lower C^1 discontinuity.

5.1 Future Work

The work in this paper has shown many topics for future research. In this section, we will review these topics, and discuss other open topics related to cylindrical pasting.

5.1.1 Polishing

As this work was proof-of-concept, there were several problems for which we used a simple solution. However, many of these solutions are not adequate for use of cylindrical pasting as a blending technique. In particular, the following areas need to be further explored:

- **End directions of spine curve**

  With our method, the direction of the axis of the blend surface is roughly perpendicular to the base surfaces. However, to form a "Y" branch, there is a need for non-perpendicular connections. This should be a straight-forward extension of the existing work.

- **Correspondence problem**

  Our solution to the correspondence problem (i.e., the association of control points in the two end rings of control points) works in some cases, but fails in other cases. A more sophisticated correspondence method should be implemented. The difficulty here is finding a solution that is stable. In particular, as we move one of the base surfaces, we would like the cylindrical blend to move without "popping."

- **Trim curve/domain curve**

  Currently, we use a circle in the domain of each base to determine the location of the trim curve on each base surface. While this works well for some surfaces, it does a poor job on other surfaces (e.g., Figure 4.14). Minimally, we will need to allow for other closed curves in the domains of each base. Ideally, we would be able to select the approximate trim curves on the base surfaces themselves, and find reasonable curves in the domains.

  One approach would be to slice each base surface with a plane, sample the intersection curve, fit a curve to the pre-image of this sample data, and use that as the trim curve. The choice of plane could be automatic, or it could be left to user control.

  Further, we should be able to use quasi-interpolation [LS75, Con99] to get the feature boundary to lie closer to the base surface, and to improve the cross-boundary derivatives. Such a change should have little if any effect on the remainder of our construction, other than to need a less refined cylinder representation.

  A related issue would be to incorporate some of the ideas of Kim and Elber [KE97] in the pasting process. For example, they reparameterize the trim curve to a unit length parameterization, giving a better distribution of the feature boundary points on the base surface. They also adjust the cross-boundary tangent directions to be perpendicular to the trim-curve, which might also improve the blends for our technique.

- **Exact Boundaries**

In some situations, it may turn out that we need an exact representation of the curve on the base surface at which the feature surface joins the base. This curve can be computed exactly using polynomial composition [DGHH93, LM97, KE97].

- **User interface**

  The current user interface is rudimentary. It was designed solely with the intent of testing the pasting ideas. Ideally, we would have a user interface where the user selects the two surfaces to blend, and the computer automatically computes the blend, with some method for allowing the user to adjust the blend as needed.

5.1.2 Hierarchical pasting

Hierarchical pasting allows the base surface to have more than one feature surface pasted on top of it, and one on top of each other. However, we did not implement hierarchical pasting for cylinders. The following discussion gives our ideas on hierarchical cylindrical pasting.

There are several situations and problems for the cylindrical pasting in a hierarchy. Some cases are easy and require no special treatment. For example, if we join two standard pasted surfaces with a cylinder, then the cylindrical pasting method described in this paper works without problem. As a second example, if we paste a surface onto the pasted cylinder, as long as the new feature domain does not overlap the cylinder's domain boundaries, then standard pasting techniques may be used.

The first interesting case occurs when pasting a surface onto a cylinder where the feature domain crosses the boundary of the cylinder's domain, as shown in Figure 5.1. In this figure, the large rectangle is the domain of a cylinder, with the left and right edge identified. If a feature is pasted on top of this domain so that it crosses the right (left) edge, then we need to remap part of the feature domain as shown for the shaded region in this figure. Once this mapping has been performed, the pasting procedure can proceed as normal.

A more complicated problem occurs if the feature domain crosses one of the other two edges of the cylindrical domain as shown in Figure 5.2. In this figure, the large rectangle on the left represents a base domain, the large rectangle on the right is the feature cylinder's domain, the circle is the embedding of the bottom edge of the cylinder's domain in the base domain. A second feature has been pasted on the base so that it overlaps the cylindrical domain. If the second feature were underneath the cylinder domain, then the pasting process works as normal. However, if the second domain is pasted on top of the cylinder, then it is

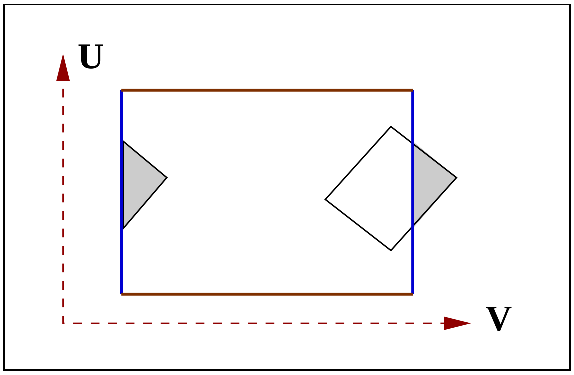

Figure 5.1: Overlapping domains.

unclear how to paste the portion of the part of the new feature that over laps the circle (the shaded region in the right half of the figure).

One idea for mapping the shaded region is to map the edges of the new feature into the cylinder's domain so that they match the first fundamental form, which is fully specified by the tangent to the circle and the line from the edge to the center of the circle (the thin lines in the figure). This mapping will ensure that the boundaries of the feature remain at least C^1.

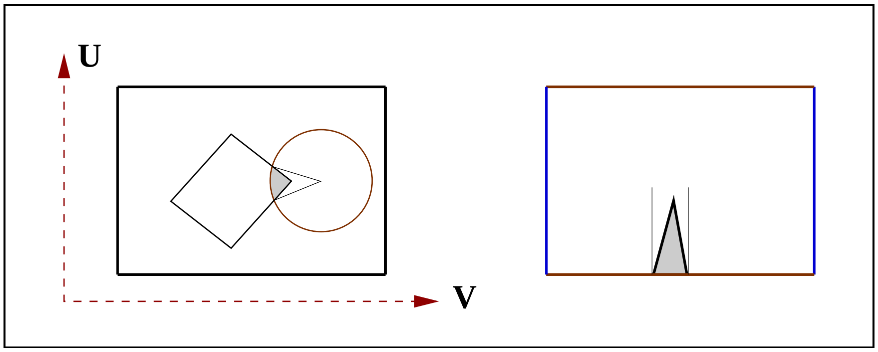

Figure 5.2: Overlapping domains.

While the above method will work in some situations, as we slide the feature further on top of the cylinder, at some point the feature boundaries would map to lines in the cylinder's domain that intersect at a point outside the domain, or worse, intersect on the wrong side of the cylinder's domain boundary line (Figure 5.3). In these cases, we would need to map the edges of the feature to curved lines, which would necessitate a non-linear mapping from the shaded region in Figure 5.4 to a region with curved boundaries in the cylinder's domain.

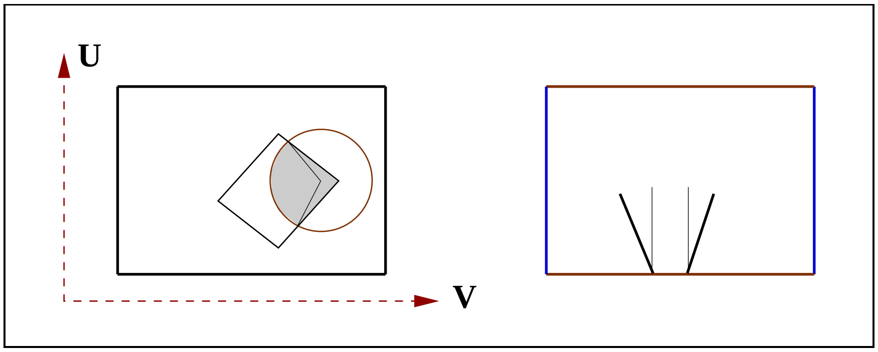

Figure 5.3: Overlapping domains.

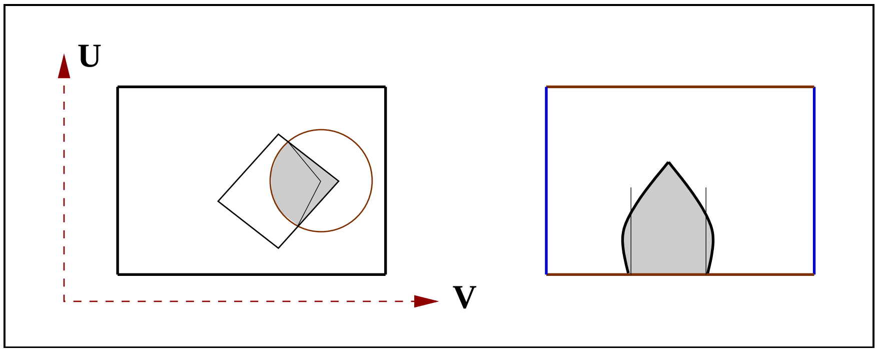

Figure 5.4: Overlapping domains.

Further difficulties occur if we slide the feature even further. If the feature is small enough, then it could slide entirely on top of the cylinder, which would work as expected. If only a small portion of the feature crosses the circle corresponding to the cylinder's domain boundary, then the new feature may just slide on and off of the cylinder. But consider the situation where new feature slides so that it covers antipodal points of the circle, completely covers or passes in and out of the circle as in Figure 5.5. Since the circle is the boundary of the cylinder's domain, the interior of the circle has no meaning. Thus, it is unclear how the domains should map, let alone what the desired pasted surface would look like in such cases.

Perhaps a better approach would be to start shrinking the domain of the new pasted surface so that it is guaranteed to avoid most of the above problems, using a shrinking map that would allow the portion of its edges inside the circle to map to straight lines in the cylinder's domain.

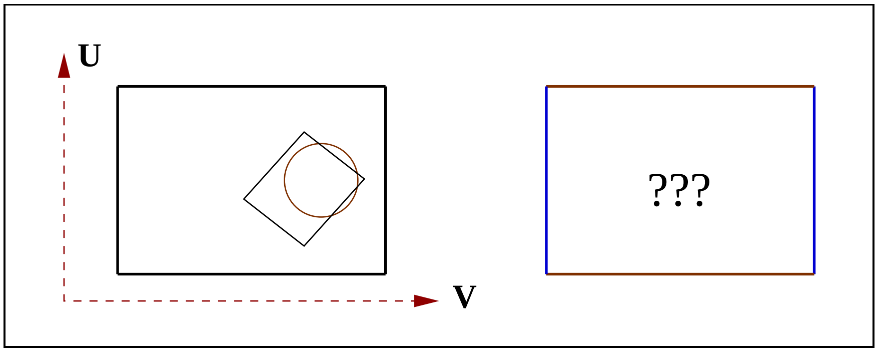

Figure 5.5: Overlapping domains.

As a final problem that might arise, in standard surface pasting, the hierarchy will as be a DAG. If a cylinder is used to create a handle (i.e., both ends of the cylinder are attached to a single surface) then we can create a cycle in our hierarchy.[1] Likewise, it may be possible to get cycles just by adding the capability to paste over the cylinder seam (Figure 5.1) or even have a feature paste onto itself. Something would need to be done to break or otherwise handle such the cycles.

---

[1] Thanks to Jorg Peters for pointing this out.

Appendix A

Graphical User Interface

This chapter gives an overview of all the functionalities the graphical user interface (GUI) provides for performing cylindrical pasting. The GUI was implemented with the Open Inventor building blocks that provide 3D manipulation of the surfaces. When the program starts, the main window of the interface will be shown in the screen. The main window is divided into three parts: the surface window on the left hand-side, two domain windows on the right hand-side with one sitting on top of the other. The top domain window is the domain view of the first base surface being selected, while the bottom one is for the second selected surface. Several scales are provided for the user to adjust parameters for producing different cylindrical pasting results. Error dialogs are also included to notify users of any invalid operations.

In the previous basic pasting research, two different user interfaces were implemented, *PasteMaker* [Cha96] and *PasteInterface* [Bar94]. PasteMaker was developed as a domain space user interface that allows user to control pasting through domain manipulations only, while *PasteInterface* provides similar functionality except that user manipulations are done directly on surfaces. *CylindersPasting* was implemented as a domain space interface using Open Inventor as the building blocks in terms of GUI aspects. A summary of the shortcut keys and menu options are listed in the following table:

A.1 Selecting Surfaces

To choose the desired surfaces to connect together, a picking method is required for the user to do selection. Open Inventor provides a built-in picking package for all the objects in the

Menu	**Function**	**Key**	**Menu**	**Function**	**Key**
File	New	Alt n	Edit	Add	Insert
	Open	Alt o		Delete	Delete
	Save	Alt s	Paste	Paste	p
	Save As	Alt a		Unpaste	u
	Quit	Alt q	Options	Trim	T
Surface	Tessellation	t		Frame	F
	Properties	s		Control Points	P
	Paste Adjust	a		Spine Curvature	C

Table A.1: Summary of shortcuts and the corresponding menu options.

rendering scene. With this package, a user can easily select the surfaces using the mouse to click on them in the 3D world space. When a surface is picked, it is highlighted with a blue bounding box. With selected surface(s), the user has an option between changing surface properties or to perform pasting.

At most two surfaces can be picked at a time, when the user tries to select the third one, all surfaces previous picked will get de-selected. When the user tries to pick a blending cylinder, all prior selected surfaces will also be de-selected. The user can either change the surface properties for this blend or to change the pasting properties, but can not paste on blend surface. When a surface is selected, its domain will appear in the corresponding domain window.

A.2 Adjusting Properties For Selected Surface(s)

With this interface, the user can change several properties of the selected surface(s) through adjusting the scaling bars built for the corresponding purposes.

A.2.1 Tessellation of Surface

The surface tessellation technique has a large influence on the rendering performance. In Open Inventor, surfaces are rendered as groups of small triangles. The higher the tessellating level, the more triangles are used, resulting in higher-quality surfaces for visual effect. Unfortunately that also means slowing down the interactive performance. In `CylindersPaster`,

Figure A.1: Surface Tessellation Slider

there is a menu option for setting the tessellation level, which pops up a window with a slider, shown in Figure A.1. Moving the slider to the left decreases the tessellation level. Moving it to the right increases the tessellation level. This tessellation value is passed to Open Inventor, which uses this value to decide how finely to tessellate the surface.

A.2.2 Transparency

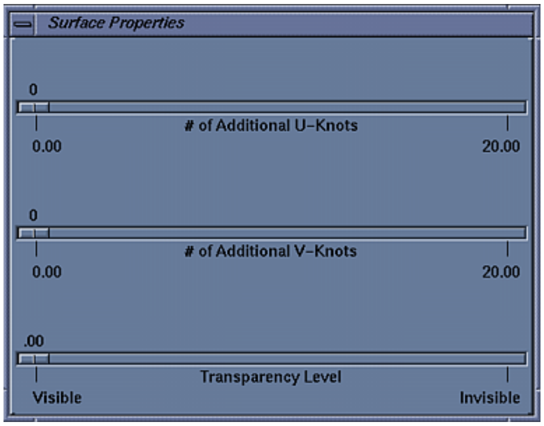

Figure A.2: Adjusting transparency of a surface

With the feature of adjusting transparency level of surfaces, it helps to distinguish between the feature and base surfaces if they have similar color material. Moreover, it can also identify the boundary for showing how well the pasting was performed. Again, there is a menu option for setting surface properties. This menu item pops up a window (Figure A.2

for setting the transparency and for inserting knots. Note that only the selected surfaces have their transparency adjusted.

A.2.3 Knot Insertion

With NUBS surfaces add more control points through knot insertion. Knot insertion on the feature wil also minimize the gap between the pasted cylinder and the base surfaces, giving a smoother connection. This GUI (Figure A.2) provides the user scaling bars for adding control points in both parametric directions. Only the selected surfaces have knots inserted. Note that knots can only be inserted; they can not be deleted. Further, the numbers on the sliders are the number of knots to insert at one time. You may repeatedly set the slider to add additional knots.

A.3 Pasting

Surface 1
● Up
○ Down

Surface 2
○ Up
● Down

Figure A.3: Orienting the blend.

If two surfaces have been selected to paste one on top of each other, a dialog will appear (Figure A.3) for asking orientation for each surface to get the relative position of the blend. There are two choices: *Up* and *Down*. Each surface must have an opposite orientation, that is, if one is oriented upwards, then the other one must be oriented downwards, and the blend will be formed between them.

A.4 Changing Properties of Blending Cylinder

Some aspects of the pasting can be adjusted to give a different look for the blending cylinder via sliders. A menu item brings up the window shown in Figure A.4. The following subsections describe its functionality.

A.4.1 Location of Blending Cylinder

The location of the blend on the base surface can be repositioned. Since this interface is a domain-space GUI, rearranging the blend's position has to be done in the domain space by manipulating the cylinder's domain. The blend has to be selected in the surface window first, and the domains for the two edges will be shown in the two domain windows. With the decision on which edge to be moved, use the left mouse button to click on the corresponding domain and slide the domain within its base domain, with the button being pressed along the motion. As the domain moves, the corresponding mapping location will be refreshed in the surface window. When the desired final position has been reached, release the button and the repositioning is completed. If the base surface is a normal NUBS surface, the circular domain is allowed to translate in both parametric directions provided that the circle is enclosed by the rectangular base domain. On the other hand, if the base is a cylindrical surface, the rectangular domain can only be translated along the height of the base.

A.4.2 Adjusting Radius of the Blending Cylinder

When a cylinder is pasted on a NUBS base, the radius of the cylindrical base (Figure 4.3) can be adjusted through the domain space. There is some restriction on this adjustment. The final circular domain must lie inside the rectangular base domain

A.4.3 Adjusting Distance between 0th and 1st Layers of CPs

The distance between the circles in the base domain for the 0th and 1st layer of the at either end of the blending cylinder (Figure 4.3) can be adjusted using a slider. This gives the user control over the lengths of the cross-boundary derivatives of the blending cylinder.

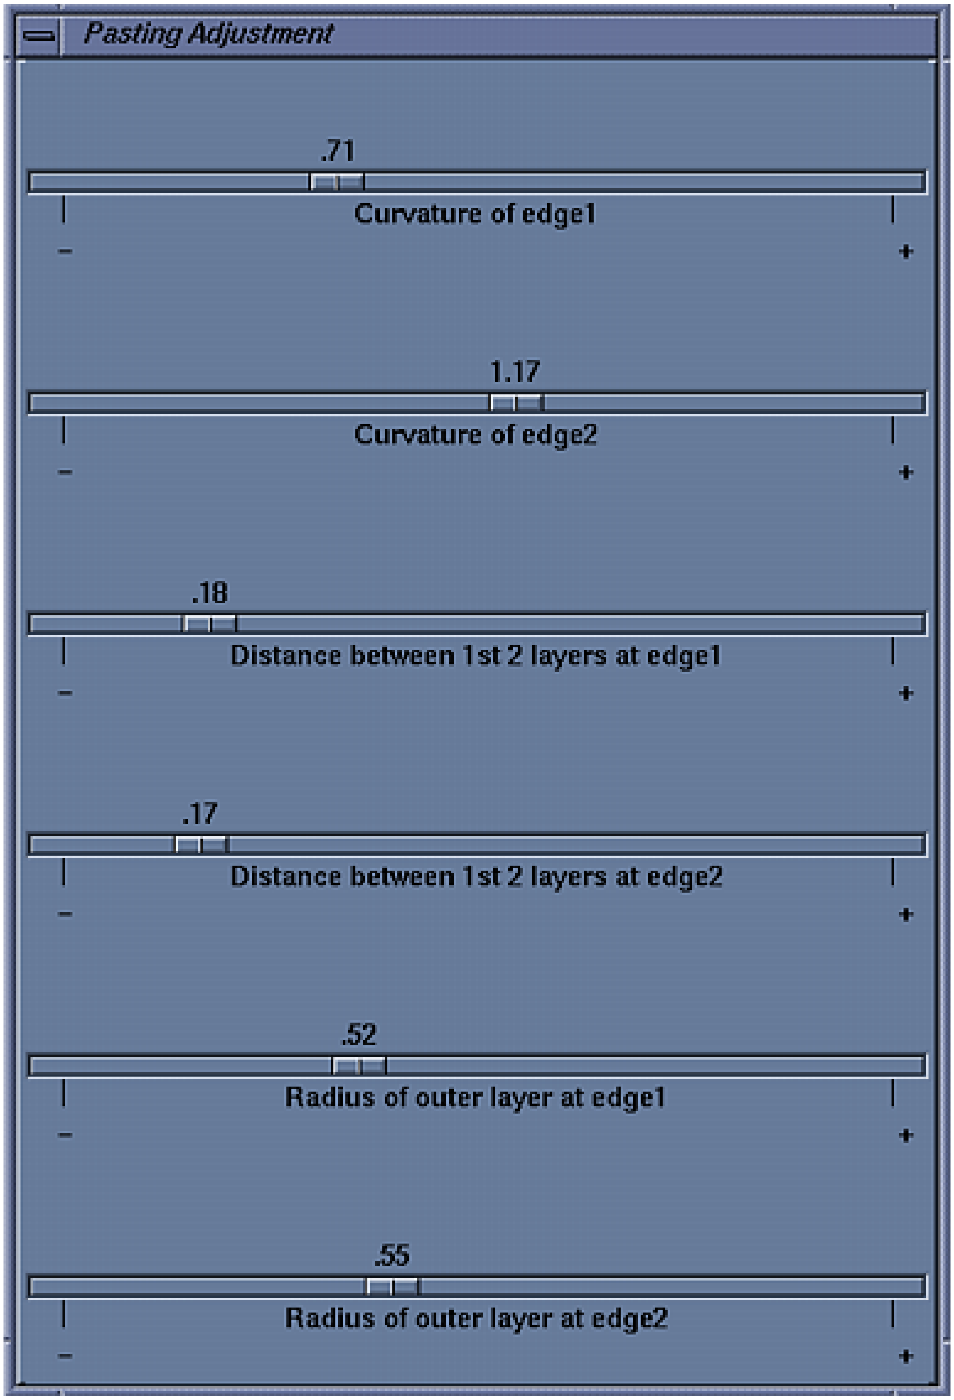

.71
Curvature of edge1

1.17
Curvature of edge2

.18
Distance between 1st 2 layers at edge1

.17
Distance between 1st 2 layers at edge2

.52
Radius of outer layer at edge1

.55
Radius of outer layer at edge2

Figure A.4: Window for adjusting the blending cylinder.

A.4.4 Adjusting Curvature of the Spine Curve

The bending of the blending cylinder is based on the curvature of the spine curve. These sliders adjust the locations of the interior control points of the spine curve.

A.5 Geometric Transformations of Surfaces

Surfaces in the surface window can be translated to a new position or being rotated in a different orientation. The surface has to be selected first. In order to translate the surface, press **z** key and use the mouse to move the surface.

A.6 Miscellaneous

There are also other features that help the user to have some basic back-end understanding of the pasting process and helped using the debugging process.

A.6.1 Control Points Display

The user has the option to display control points for all the surfaces in the world space window. This can give an idea where the points are located and how they can generate the base surfaces and the blending cylinder, in addition to the connection between them. This feature helps in adjusting the distance between the 0th and 1st layers of control points for the blending cylinder.

A.6.2 Spine Curve Display

As described in the previous chapter, the blending cylinder is constructed upon a spine curve. To have a better view of the skeleton of the cylinder, the spine curve can be displayed, and it will assist in adjusting the bending curvature of the cylinder.

A.6.3 Coordinate Frames Display

The method to construct local coordinate frames is critical in a pasting process. The ones constructed at the Greville domain points in the domain space will be mapped over to the world space and displace the Greville B-spline surface. This mapping can be examined by enabling the coordinate frame display option.

Bibliography

[Bar94] Cristin Barghiel. Feature oriented composition of b-spline surfaces. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1994. (Available as Computer Science Department Technical Report CS-94-13).

[BBF95] Cristin Barghiel, Richard Bartels, and David Forsey. Pasting spline surfaces. In L. Schumaker M Daehlen, T. Lyche, editor, *Mathematical Methods for Curves and Surfaces*, pages 31–40. Vanderbilt University Press, 1995.

[BF91] R. Bartels and D. Forsey. Spline overlay surfaces. Technical Report CS-92-08, Univeristy of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1991.

[Cha96] Leith Chan. World space user interface for surface pasting. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1996. Available as Computer Science Department Technical Report CS-96-32, ftp://cs-archive.uwaterloo.ca/cs-archive/CS-96-32/.

[Con99] Blair Conrad. Better pasting through quasi-interpolation. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1999. Available as Computer Science Department Technical Report CS-99-14, ftp://cs-archive.uwaterloo.ca/cs-archive/CS-99-14/.

[DGHH93] Tony DeRose, Ronald Goldman, and and Stephen Mann Hans Hagen. Functional composition via blossoming. *ACM Transactions on Graphics*, 12(2):113–135, April 1993.

[Far94] Gerald Farin. *Curves and Surfaces for Computer Aided Geometric Design*. Academic Press, Inc., 3 edition, 1994.

[KE97] Kwansik Kim and Gershon Elber. New approaches to freeform surface fillets. *The Journal of Visualization and Computer Animation*, 8(2):69–80, 1997.

[LM97] Wayne Liu and Stephen Mann. An optimal algorithm for expanding the composition of polynomials. *ACM Transactions on Graphics*, 16(2):155–178, April 1997.

[LS75] Tom Lyche and Larry Schumaker. Local spline approximation methods. *Journal of Approximation Theoryjat*, 15:294–325, 1975.

[VB92] Allan Vermeulen and Richard Bartels. C++ splines classes for prototyping. In *In Curves and Surfaces in Computer Vision and Graphics II*, pages 1610:121–131, Bellingham, Washington, 1992. SPIE Proceedings.

[VMV94] J. Vida, R.R. Martin, and T. Varady. A survey of blending methods that use parametric surfaces. In *Computer Aided Design*, volume 26(5), pages 341–365, 1994.