

Distortion Minimization and Continuity Preservation in Surface Pasting

Rick Leung
ryleung@cgl.uwaterloo.ca

Stephen Mann
smann@uwaterloo.ca

Computer Graphics Laboratory
University of Waterloo

Abstract

Surface pasting is a hierarchical modeling technique capable of adding local details to tensor product B-spline surfaces without incurring significant computational costs. In this paper, we describe how the continuity conditions of this technique can be improved through the use of least squares fitting and the application of some general B-spline continuity properties. More importantly, we address distortion issues inherent to the standard pasting technique by using an alternative mapping of the interior control points.

Key words: B-spline, surface pasting, continuity, least squares, distortion reduction

1 Introduction

Spline curves and surfaces are used in many areas of computer graphics and computer aided geometric design. In particular, tensor product B-spline surfaces are commonly used in modeling and computer animation because of their desirable properties such as compact representation, local control of surface geometry, and adjustable levels of internal and cross-boundary continuity.

Despite its wide-spread acceptance, it is somewhat difficult to add regions of local detail to tensor product B-spline surfaces. Traditional methods call for the insertion of extra knots or the degree raising of the surface. These methods erode the benefits of B-spline surfaces by increasing the storage requirements and/or computation costs. More importantly, they impose restrictions on the orientation and placement of the resulting features.

To address the problem of local details, Barghiel, Bartels, and Forsey [1] introduced *surface pasting* (an extension of hierarchical B-splines [7]), which was further developed by others [4, 9, 12]. This modeling technique combines the performance and storage benefits of tensor product surfaces with the flexibility of arbitrary orientation and placement of the features. These properties have attracted the attention of the modeling industry — support for surface pasting is included in recent versions of *Houdini*, a commercial animation tool by Side Effects

Software. More recently, by adopting a slightly different definition of “surface pasting”, Biermann et al. [2] extended the pasting ideas to multiresolution subdivision surfaces.

Surface pasting is an approximation algorithm. While this is the inherent factor that allows pasted features to be oriented arbitrarily, it means that the surfaces only meet with approximate C^0 and C^1 continuity. As a result, although continuous joints can be reasonably approximated when the feature has a sufficiently dense knot structure or when the pasting region on the base is relatively flat, continuity issues arise when these conditions are violated. A second, more important problem with surface pasting is the distortion that occurs in the feature surface when the base has high curvature.

In this paper, we address these short-comings of the original surface pasting technique by proposing alternate schemes to place the control vertices of the feature surface. In particular, least squares approximation and B-spline continuity principles adapted from cylindrical surface pasting [10] are used to more strongly enforce approximate C^0 and C^1 continuity, respectively. Furthermore, we show a new method for setting the remaining feature control points to minimize distortions in its overall shape.

2 Previous Work

2.1 Standard Surface Pasting

The original surface pasting algorithm is a technique in which a *feature* surface is attached to a *base* surface via a change of basis on each of the feature’s control vertices. Since we are modifying the feature control points, using tensor product B-splines as the feature surface eases the pasting process — the B-splines will maintain continuity within the pasted feature with no effort on our part, and still allow for complex, piecewise polynomial features. The form of the base surface is less important, although if we paste hierarchically then we would also want tensor product B-splines as base surfaces for the same reason.

Typically, a tensor product B-spline surface $S(u, v)$ is

written as

$$S(u, v) = \sum_{i=0}^{m+s-1} \sum_{j=0}^{n+t-1} P_{i,j} N_i^m(u) N_j^n(v),$$

where $P_{i,j}$ specifies the two-dimensional grid of control vertices, $N_i^m(u)$ and $N_j^n(v)$ the m - and n -degree B-spline basis functions defined over the knot vectors $[u_0, \dots, u_{2m+s-2}]$ and $[v_0, \dots, v_{2n+t-2}]$, respectively,¹ and s and t the number of segments the tensor product surface has in the u and v directions, respectively. We will assume that the end knots of both knot vectors have full multiplicity. In standard surface pasting, the feature surface $S_f(u, v)$ is re-expressed in *diffuse representation* as

$$S_f(u, v) = \sum_{i=0}^{m+s-1} \sum_{j=0}^{n+t-1} (\Gamma_{i,j} + \vec{d}_{i,j}) N_i^m(u) N_j^n(v),$$

where $\Gamma_{i,j}$ is the three-dimensional *Greville point*

$$(\gamma_i, \gamma_j, 0) = \left(\frac{u_i + \dots + u_{i+m-1}}{m}, \frac{v_j + \dots + v_{j+n-1}}{n}, 0 \right)$$

and $\vec{d}_{i,j}$ is the *Greville displacement*

$$\begin{aligned} \vec{d}_{i,j} &= P_{i,j} - \Gamma_{i,j} \\ &= d_{i,j}^x \vec{i} + d_{i,j}^y \vec{j} + d_{i,j}^z \vec{k}, \end{aligned}$$

where $\vec{i}, \vec{j}, \vec{k}$ is a basis for the feature space.




Figure 1: Standard surface pasting

Figure 1 shows a two-dimensional example of standard surface pasting. The feature S_f is pasted onto the base S_b by re-defining the origin of its control vertices and its associated local coordinate frame. Given an invertible transformation \mathbf{T} that maps the feature domain into the base domain, we map the Greville points in the feature domain into the base domain:

$$\begin{aligned} \mathbf{T}(\Gamma_{i,j}) &= \mathbf{T}((\gamma_i, \gamma_j, 0)) \\ &= (\gamma'_i, \gamma'_j, 0). \end{aligned}$$

¹Some definitions of B-splines add an extra knot at both ends of each knot vector.

We now write the pasted feature S'_f as

$$S'_f(u, v) = \sum_{i=0}^{m+s-1} \sum_{j=0}^{n+t-1} \left(S_b(\gamma'_i, \gamma'_j) + d_{i,j}^x \vec{i}' + d_{i,j}^y \vec{j}' + d_{i,j}^z \vec{k}' \right) N_i^m(u) N_j^n(v).$$

The vectors \vec{i}' and \vec{j}' are the partial derivatives at $S_b(\mathbf{T}(\Gamma_{i,j}))$, and $\vec{k}' = \vec{i}' \times \vec{j}'$.

2.2 Limitations of Previous Schemes

One of the primary drawbacks of the standard surface pasting algorithm is the often inadequate approximation of continuity between the feature and the base. When the base surface has high curvature, or when the feature surface has a coarse knot structure, gaps may appear between the base and the pasted feature as illustrated in Figure 2. While this lack of C^0 continuity can be alleviated



Figure 2: C^0 discontinuity in standard surface pasting

by *knot insertion*, this approach is far from ideal. The extra knots introduced by knot insertion add additional control vertices across the entire feature row-wise and/or column-wise. This increases the data storage requirements of the feature, and more importantly, the larger number of control vertices makes it more costly to paste the feature.

In addition to the C^0 continuity issues, standard surface pasting suffers from a lack of C^1 continuity. Although restricting the boundary and second layer control points to have zero displacement vectors can partially remedy this short-coming, the effectiveness of the resulting C^1 approximation is influenced by the curvature of the base surface. Furthermore, enforcing this practical limitation means that the feature must have a rectangular footprint.

By using *quasi-interpolation* to determine the locations of the feature's boundary control vertices, Conrad [5, 6] improved the C^0 and C^1 approximations at the joints between the feature and the base. Unfortunately, the underlying mathematics that drive his algorithm are rather complex, and it is unclear whether the improved continuity can be maintained when the feature

possess a non-rectangular footprint. More importantly, quasi-interpolated pasting does not address the distortion issue that the standard algorithm faces.

In the standard technique, the final shape of the pasted feature is determined by a combination of its original unpasted shape and the shape of the base surface. This allows the feature to adapt to the shape of the base, which is one of the advantages of surface pasting. Unfortunately, if the base surface has a high curvature, the pasted feature may become grossly distorted.

3 Distortion-Minimizing and Continuity-Preserving Surface Pasting

The new surface pasting scheme proposed here operates on the control vertices of the feature surface using three methods, depending into what class each control vertex falls. The three different classes (illustrated in Figure 4) are the boundary (black) control points, the second layer (gray) control points, and the interior (white) control points. The boundary control points form the outer-most layer of control vertices in a tensor product B-spline surface. In our new method, these control points are pasted onto the base surface using least squares approximation to provide an improved C^0 joint. To obtain a more accurate C^1 approximation, the second layer control vertices are pasted using a technique adapted from cylindrical surface pasting [10]. Last but not least, the interior control points consists of the control vertices that do not belong to the previous two groups. They are pasted through an affine transformation \mathbf{T}_{int} that maps the control vertices from the unpasted surface to the pasted region to preserve the shape of the unpasted feature.

3.1 Obtaining Approximate C^0 Continuity

To decrease the C^0 discontinuity between the feature and base surfaces, we apply least squares B-spline curve approximation to each of the feature surface's four boundaries [3, 11]. By sampling the base surface along the boundaries where it meets the feature, we reconstruct B-spline curves that closely approximate the joints between the feature and the base; we will use these B-spline curves as the boundary curves of our pasted feature.

We use a B-spline least squares algorithm that interpolates the first and last sample points; thus, each boundary curve reconstruction may proceed independently of the others as long as the appropriate pairs of corner control points are used as the first and last samples in the reconstructions. No other requirements between the different least squares fits are needed to guarantee that adjacent boundaries meet correctly at the corners of the pasted feature.

Sampling Methodology

Since the least squares reconstruction of each boundary is independent of the others, we describe the reconstruction of a single boundary. Consider the boundary of the tensor product feature surface S_f along the u direction, with $v = v_0$, and assume that the displacement vectors along this boundary have zero lengths as required by the approximate C^0 continuity condition in standard surface pasting. To apply least squares reconstruction, we sample p points from the base surface along the desired boundary between the base and feature surfaces. Note that if the reconstructed curve is a B-spline of degree m with s segments, then the least squares reconstruction procedure requires that $p > m + s - 1$.

To satisfy sampling requirements while adhering to the basic strategy of standard surface pasting, we use the points $S_b(\gamma'_0, \gamma'_0), \dots, S_b(\gamma'_{m+s-1}, \gamma'_0)$, and $S_b(\frac{\gamma'_0 + \gamma'_1}{2}, \gamma'_0), \dots, S_b(\frac{\gamma'_{m+s-2} + \gamma'_{m+s-1}}{2}, \gamma'_0)$ as curve samples. The first set of points are the pasted Greville points along the boundary of the feature. The second set are points on the base surface evaluated at the halfway point between each pair of transformed Greville points. To summarize, the samples to be used are

$$\begin{aligned} Q_0 &= S_b(\gamma'_0, \gamma'_0) \\ Q_1 &= S_b(\frac{\gamma'_0 + \gamma'_1}{2}, \gamma'_0) \\ Q_2 &= S_b(\gamma'_1, \gamma'_0) \\ Q_3 &= S_b(\frac{\gamma'_1 + \gamma'_2}{2}, \gamma'_0) \\ &\vdots \\ Q_{p-1} &= S_b(\frac{\gamma'_{m+s-2} + \gamma'_{m+s-1}}{2}, \gamma'_0) \\ Q_p &= S_b(\gamma'_{m+s-1}, \gamma'_0), \end{aligned}$$

where $p = 2(m + s) - 3$. By this choice of p , we strike a balance between accuracy and efficiency in the least squares reconstruction that produces the desired boundary curve of the feature surface. With the reconstruction, we obtain an improved C^0 approximation at the joint between the feature and the base.

Accommodating Non-Rectangular Footprints

While the previous sampling methodology will produce good C^0 approximations for a feature surface with zero displacement vectors around its boundaries, it cannot accommodate features with non-rectangular footprints. To allow for features with non-rectangular boundaries, we modified the pasting process to allow each boundary control point to lie at a location other than its corresponding Greville point. We still restrict the feature's unpasted boundary control points to lie in the $z = 0$ plane, but rather than insisting on a zero displacement vector, these

control points are allowed to have a non-vertical displacement vector.

Specifically, before computing the sample points on the base surface, we evaluate the unpasted feature surface at its boundary Greville points and the halfway points as shown in Figure 3. Since no vertical displacements




Figure 3: Evaluating the boundary points. Note that each boundary sample point Q_i is the image of a corresponding boundary Greville or halfway point.

are involved up to this point, the images of the boundary Greville and halfway points can still be treated as points that lie on the 2D domain of the base surface. As such, they can be pasted onto the base surface and serve as the sample points for the least squares reconstruction of the boundary curve.

Using the same example as before, the points of the least squares reconstruction are now

$$\begin{aligned} Q_0 &= S_b(S_f^x(\gamma'_0, \gamma'_0), S_f^y(\gamma'_0, \gamma'_0)) \\ Q_1 &= S_b(S_f^x(\frac{\gamma'_0 + \gamma'_1}{2}, \gamma'_0), S_f^y(\frac{\gamma'_0 + \gamma'_1}{2}, \gamma'_0)) \\ &\vdots \\ Q_p &= S_b(S_f^x(\gamma'_{m+s-1}, \gamma'_0), S_f^y(\gamma'_{m+s-1}, \gamma'_0)). \end{aligned}$$

3.2 Obtaining Approximate C^1 Continuity

Earlier surface pasting methods achieved approximate C^1 continuity between the feature and base surfaces by placing the first two (and sometimes three) layers of control points at their Greville points in the embedded domain. This method gives satisfactory results if one performs a large amount of knot insertion in the feature surface. To reduce the amount of knot insertion and still have satisfactory continuity, we adapted a C^1 continuity condition from cylindrical surface pasting [10]. This latter approach uses the fact that the partial derivatives at the boundary control points of the feature are completely determined by the location of the second layer control points. By pasting the second layer control points based on the cross-boundary derivatives of the base surface, we

obtained improved approximate C^1 continuity between the feature and the base.

Roughly speaking, for each boundary point Q of the feature with Greville point Γ_Q , we compute the cross-boundary (directional) derivative of the base surface at $\Gamma'_Q = \mathbf{T}(\Gamma_Q)$, giving $\vec{c} = D_{\vec{v}} S_b(\Gamma'_Q)$, where \vec{v} is the cross-boundary direction of the feature domain as embedded in the base domain. We then set the second layer control point P corresponding the Q to lie on the ray Q, \vec{c} . The following discusses how we locate the control point along this ray.




Figure 4: Second layer control point calculation with approximate C^1 continuity

We consider three cases, as illustrated in Figure 4, based on whether the second layer control point is influenced by the \vec{v} directional derivative (case 1), the \vec{u} directional derivative (case 2), or both (case 3). In the first case, we set the second layer control point to

$$P = Q + \alpha_P \cdot \left\| D_{\vec{v}} S_b(\Gamma'_Q) \right\|,$$

where α_P is a scaling factor, and $\|\vec{w}\|$ is the unit vector parallel to \vec{w} . In the second case, we set the second layer control point to

$$P = R + \beta_P \cdot \left\| D_{\vec{u}} S_b(\Gamma'_R) \right\|,$$

where β_P is a scaling factor. In the third case, we set the second layer control point to

$$\begin{aligned} P &= \sigma \cdot \left(Q + \alpha_P \cdot \left\| D_{\vec{v}} S_b(\Gamma'_Q) \right\| \right) + \\ &\quad \tau \cdot \left(R + \beta_P \cdot \left\| D_{\vec{u}} S_b(\Gamma'_R) \right\| \right), \end{aligned}$$

where σ is defined to be

$$\sigma = \frac{|R - \Gamma_P|}{|R - \Gamma_P| + |\Gamma_P - Q|},$$

with $||$ being the distance function, and Γ_P being the unpasted Greville point of P . The scaling factor τ is defined in a manner similar to σ , except that $|Q - \Gamma_P|$ is used as the numerator. Together, σ and τ serve as weights to control the amount of influence Q and R have on the final placement of P .

The value of α_P is determined based on the distance between the boundary point Q and Γ_P , and the overall density of the corresponding knot vector along v . Specifically, we define α_P to be

$$\alpha_P = |Q - \Gamma_P| \cdot \frac{n}{n+t-1}, \quad (1)$$

where n is the degree of the feature surface in the v direction, and t is the number of B-spline segments along v . The first term in equation (1) reflects the spatial relationship between the second layer control point P and its adjacent boundary point Q . By taking this spatial relationship into account, the shape of the pasted feature in the region between the boundary and second layer control points will be more appropriately defined. The fractional term of equation (1) scales the distance term to progressively reduce the value of α_P as the number of B-spline segments increases. This pulls P closer to Q and prevents P from being placed too close to the interior control points as the overall number of control vertices increases.

The value of β_P is defined analogously.

3.3 Shape Preservation

To preserve the original shape of the unpasted feature as much as possible during the pasting operation, we build an affine transformation \mathbf{T}_{int} to map each interior control point from its unpasted location to a corresponding position in the pasted feature such that the spatial relationships between the interior control points are preserved. Figure 5 illustrates the effect of the \mathbf{T}_{int} transformation.




Figure 5: Affine transformation \mathbf{T}_{int}

Observe that in the figure, the spatial relationships between the three interior control points are the same before and after \mathbf{T}_{int} operates on them. As a result, the overall shape of the feature surface is better maintained after the pasting operation.

Our construction of the affine transformation \mathbf{T}_{int} consists of a number of steps. The first step involves mapping the two dimensional domain plane of the unpasted feature to a three dimensional plane that approximates the pasted boundary Greville points. This 3D plane is known as the *reference plane* (also known as the *datum plane* in many CAD packages). Then, using the four pasted corner boundary Greville points, we construct a local coordinate frame F_T . Using F_T , we map the interior control points to their transformed location in the pasted surface.

The Reference Plane

We create the reference plane to allow the pasted interior control points to be defined in a manner similar to how the unpasted control points are defined relative to the domain plane of the unpasted feature.

To define the plane, we first determine its normal vector by approximating the entire set of pasted boundary Greville points using least squares three dimensional plane approximation. Then, we calculate a point C_T as the centroid of the four pasted corner Greville points. The plane is fixed in three space by substituting C_T into the plane equation. This approach roughly captures the general orientation of the feature boundaries while loosely placing the plane close to the four pasted corner Greville points.

The Local Coordinate Frame F_T

Once the reference plane has been calculated, its normal vector can be used as the frame basis vector \vec{z}_T . We then use the four pasted corner boundary Greville points to determine the remaining frame basis vectors \vec{x}_T and \vec{y}_T . Observe that in the unpasted feature, these Greville points are aligned along the parametric u and v directions, which are typically taken to be the orthonormal frame basis vectors \vec{x} and \vec{y} of the diffuse coordinate system. By a similar measure, the basis vectors \vec{x}_T and \vec{y}_T can be calculated by taking the differences between the appropriate pairs of pasted Greville points.

Given the four pasted corner boundary Greville points $\Gamma_{ul} = S_b(\gamma'_0, \gamma'_{n+t-1})$, $\Gamma_{ur} = S_b(\gamma'_{m+s-1}, \gamma'_{n+t-1})$, $\Gamma_{ul} = S_b(\gamma'_0, \gamma'_0)$, and $\Gamma_{lr} = S_b(\gamma'_{m+s-1}, \gamma'_0)$, the vector \vec{x}'_T can be calculated as follows. Since the reference plane only approximates the four pasted Greville points, it is unlikely for the plane to coincide with any of the pasted Greville points. By projecting the points onto the reference plane along the plane's normal vector, we get the four points Γ_{ul}^p , Γ_{ur}^p , Γ_{ul}^p , and Γ_{lr}^p . Since these points lie directly on the reference plane, the frame basis vector can be calculated as

$$\vec{x}_T = \left\| \frac{\|\Gamma_{ur}^p - \Gamma_{ul}^p\| + \|\Gamma_{lr}^p - \Gamma_{ul}^p\|}{2} \right\|,$$

where $\|\vec{x}\|$ is the unit vector in direction \vec{x} . Note that

\vec{x}_T is calculated as the average of the two vectors that correspond to the parametric u direction because these two vectors are not necessarily parallel. By obtaining an average between the two, the directional error contained in \vec{x}_T is reduced.

In a similar manner, \vec{y}_T is calculated as

$$\vec{y}_T = \left\| \left\| \Gamma_{ul}^p - \Gamma_{ll}^p \right\| + \left\| \Gamma_{ur}^p - \Gamma_{lr}^p \right\| \right\|.$$

Although \vec{x}_T and \vec{y}_T are orthogonal to \vec{z}_T by construction, they are not necessarily orthogonal to each other. Without the \vec{x}_T and \vec{y}_T axes being perpendicular to each other, the \mathbf{T}_{int} transformation defined in the following section cannot preserve the spatial relationships between the interior control points, thereby causing distortions in the pasted feature. To make the axes orthonormal, we rotate \vec{x}_T and \vec{y}_T in opposite directions about \vec{z}_T by the same amount such that all three vectors become orthogonal to each other.

Applying the Affine Transformation \mathbf{T}_{int}

Before we apply \mathbf{T}_{int} to the interior control points of the feature, we compute the centroid of the unpasted corner Greville points to serve as the reference point C_{ref} . With the reference point defined, \mathbf{T}_{int} can be applied to the interior control points by determining the spatial relationship between each control vertex and C_{ref} , after which this relationship is re-expressed with C_T as the reference point. Specifically, for each unpasted interior control point $P_{i,j}$ of the feature surface S_f , let

$$\vec{\delta}_{i,j} = P_{i,j} - C_{ref}.$$

Then $P_{i,j}$ can be pasted onto the base surface S_b via the affine transformation \mathbf{T}_{int} in the following manner:

$$\begin{aligned} \mathbf{T}_{int}(P_{i,j}) &= C_T + \varsigma \cdot \vec{\delta}_{i,j} \\ &= C_T + \varsigma \cdot \delta_{i,j}^x \cdot \vec{x}_T + \varsigma \cdot \delta_{i,j}^y \cdot \vec{y}_T + \\ &\quad \varsigma \cdot \delta_{i,j}^z \cdot \vec{z}_T \end{aligned}$$

where ς is a scalar, and $\delta_{i,j}^x$, $\delta_{i,j}^y$, and $\delta_{i,j}^z$ are the individual components of the displacement vector $\vec{\delta}_{i,j}$ relative to $\{C_{ref}, \vec{x}, \vec{y}, \vec{z}\}$.

3.4 Accounting for the Shape of the Base Surface

The affine transformation \mathbf{T}_{int} preserves the shape of a feature surface by using a single local coordinate frame to determine the location of every interior control point. However, \mathbf{T}_{int} is not influenced by the underlying form and orientation of the base surface on which the feature is pasted. To re-introduce some of the influence the base surface has on the pasted feature, the positions of the interior control points can be linearly interpolated between the coordinates given by the standard surface pasting algorithm and those given by the \mathbf{T}_{int} transformation.

4 Results




Figure 6: Modified surface pasting with the application of C^0 and C^1 continuity constraints

Figure 6 shows the result of pasting a parametrically aligned bicubic feature on a bicubic base using least squares curve fitting for the boundaries and the C^1 continuity conditions described in Section 3.2. Compared to Figure 2, which shows the same feature pasted on the same base using the standard technique, notice the reduced C^0 discontinuity. Differences in C^1 continuity between the two figures are minimal because the C^1 improvements were already incorporated in the earlier figure.

Figure 7 illustrates the capability of our boundary-fitting scheme to accommodate feature surfaces with non-rectangular footprints. In the figures, a bicubic feature with three segments in each parametric direction is pasted on a bicubic base, with their parametric directions differing by 45° . While the standard algorithm does not prevent the feature from having non-linear boundaries (Figure 7a), huge gaps are left between the feature and the base. By applying least squares approximation to the boundaries as shown in Figure 7b, our modified scheme maintains approximate C^0 continuity regardless of the shape of the footprint.

As a final example to illustrate the effectiveness of the shape preservation property of the \mathbf{T}_{int} transformation, consider pasting the cabin surface shown in Figure 8a onto the body surface shown in Figure 8b. Figure 8c shows the results of this operation performed using the standard technique. Notice that since the standard pasting algorithm causes the feature surface to follow the shape of the base, the arches on either side of the body surface have effectively caused the cabin surface to cave-in on itself. Furthermore, since the knot structure of the cabin surface is not dense enough relative to the base surface, unsightly gaps are formed around the arches where the feature meets the base.

Figure 8d shows the results of pasting the cabin surface onto the body surface via \mathbf{T}_{int} pasting, least squares-fitted boundaries, and improved C^1 conditions. In addition to the C^0 continuity improvements, notice how the




Figure 7: Feature with non-rectangular footprint pasted using the standard algorithm (a) and the least-squares fitted boundaries scheme (b).




Figure 8: The cabin surface (a) is pasted onto the body surface (b) via standard surface pasting (c) and \mathbf{T}_{int} pasting with least squares-fitted boundaries and improved C^1 conditions (d). Pasting via 33% (e) and 66% (f) \mathbf{T}_{int} interpolations with least squares-fitted boundaries and improved C^1 constraints are shown.




Figure 9: Models created via \mathbf{T}_{int} surface pasting.

cabin surface retains most of its characteristic shape after the pasting operation.

Figures 8e and 8f show the results of applying a 33- and 66-percent linear interpolation, respectively, between the standard technique and our modified scheme. Note that the interpolation is applied only to the interior control points of the feature.

5 Conclusion

As one method to provide local details to tensor product surfaces, this paper presented a new and refined procedure based on surface pasting. The new method operates on a given tensor product B-spline feature by dividing the surface's control vertices into three distinct classes. We applied least squares approximation to the boundary control points of the feature to reconstruct the feature boundaries such that the best possible C^0 approximation to the base is produced. To further ensure a smooth and seamless-looking joint, approximate C^1 continuity conditions were borrowed from cylindrical surface pasting and adapted for use in tensor product surfaces to appropriately place the second layer control vertices of the feature. Finally, shape preservation of the pasted feature was achieved by defining a custom affine transformation \mathbf{T}_{int} that operates on the interior control points of the feature. By preserving the spatial relationships between every interior control vertex, little distortion is seen in the resulting pasted feature. If geometric influence from the base surface is required, distortion can be re-introduced in a controlled manner by linearly interpolating the positions of the interior control points between the standard scheme and the shape-preserving technique.

Acknowledgements

We would like to thank everyone in the graphics lab (CGL) at the University of Waterloo for their help and support. This project was funded in part by CITO and NSERC.

References

- [1] Cristin Barghiel, Richard Bartels, and David Forsey. Pasting spline surfaces. In M. Daehlen, T. Lyche, and L. L. Schumaker, editors, *Mathematical Methods for Curves and Surfaces: Ulvik, Norway*, pages 31–40, Nashville, TN., 1994. Vanderbilt University Press.
- [2] Henning Biermann, Ioana Martin, Fausto Bernardini, and Denis Zorin. Cut-and-paste editing of multiresolution surfaces. *ACM Transactions on Graphics*, 21(3):312 – 321, July 2002.
- [3] Richard L. Burden and J. Douglas Faires. *Numerical Analysis*, chapter 8.1. Brooks/Cole Publishing Company, sixth edition, 1997.
- [4] Leith K. Y. Chan. World space user interface for surface pasting. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1996. Available on WWW as mboxftp://cs-archive.uwaterloo.ca/cs-archive/CS-96-32/.
- [5] Blair Conrad. Better pasting through quasi-interpolation. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1999. Available on WWW as ftp://cs-archive.uwaterloo.ca/cs-archive/CS-99-14/.
- [6] Blair Conrad and Stephen Mann. Better pasting via quasi-intepolation. In Schumaker Laurent, Sablonnière, editor, *Curve and Surface Design: Saint-Malo, 1999*. Vanderbilt University Press, 2000.
- [7] David Forsey and Richard Bartels. Hierarchical B-spline refinement. In *Computer Graphics (SIGGRAPH '88 Proceedings)*, pages 205–212, 1988.
- [8] Ricky Ying-Kei Leung. Distortion minimization and continuity preservation in surface pasting. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 2002.
- [9] Marryat Ma. The direct manipulation of pasted surfaces. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 2000. Available on WWW as ftp://cs-archive.uwaterloo.ca/cs-archive/CS-2000-15/.
- [10] Stephen Mann and Teresa Yeung. Cylindrical surface pasting. In G. Farin G. Grunnett, H Bieri, editor, *Geometric Modelling*, pages 233–248. Springer, 2001.
- [11] Les Piegl and Wayne Tiller. *The NURBS Book*, chapter 9. Springer, 1995.
- [12] Clara L. F. Tsang. Animated surface pasting. Master's thesis, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, 1998. Available on WWW as ftp://cs-archive.uwaterloo.ca/cs-archive/CS-98-19/.