Multiresolution Editing of Pasted Surfaces

Marryat Ma
Stephen Mann

Abstract. Surface pasting allows the insertion of local detail to a tensor product surface without changing the structure of the underlying surface. It works by applying feature surfaces on top of a base surface to create a composite surface. Previous modelling systems for pasted surfaces only allowed users to translate, rotate, and resize pasted features, and did not support direct manipulation. In this paper, we describe a method for the direct manipulation of pasted surfaces that allows the user to edit a surface at any level in the pasting hierarchy.

§1. Introduction

Hierarchical modelling is currently an active area for research. Many surfaces have varying levels of detail, and modelling techniques that explicitly represent these levels of detail are useful both in terms of reduced storage and in interactive modelling paradigms where users want to interact with their models at different levels of detail.

Tensor product B-spline surfaces are commonly used in the computer industry because they can be represented by little information and have attractive continuity properties. However, it is difficult to add detail to these surfaces without globally increasing the complexity of the surfaces and thus they are poorly suited to multiresolution editing.

Hierarchical B-splines were developed by Forsey and Bartels [7] for adding areas of local detail to a tensor product B-spline surface. A parametrically aligned region of the surface is locally refined to increase its control point density. The control points in the refined region are displaced to create the local detail. Forsey implemented a limited direct manipulation scheme for hierarchical B-splines that allows a user to manipulate a surface at predefined surface points. If the user decides to manipulate the surface at a lower resolution level, then the higher levels of detail are removed from the display. Unfortunately, with the details hidden, the user is unable to see the effect of

Mathematical Methods in CAGD: Oslo 2000
Tom Lyche and Larry L. Schumaker (eds.), pp. 1–10.

ISBN 0-8265-xxxx-x.

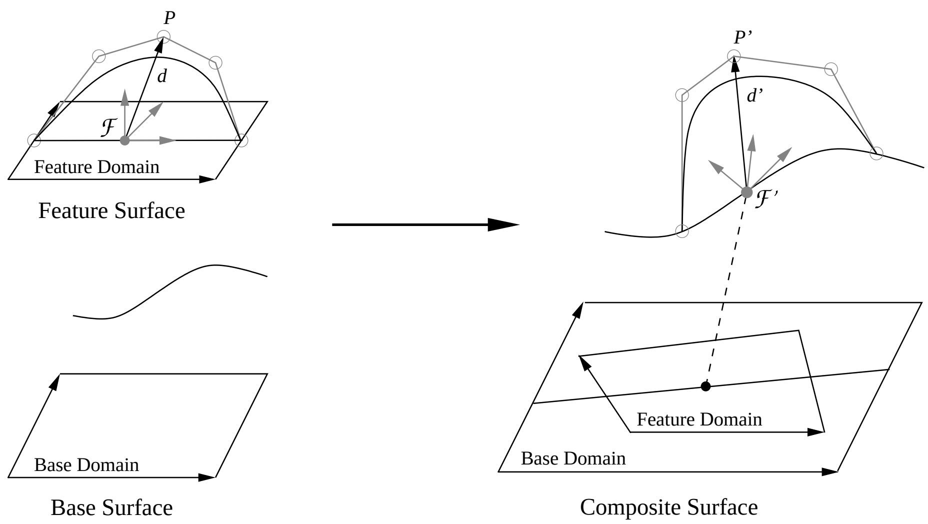

Fig. 1. Pasting a feature on a base surface.

the manipulation on the entire surface. Hierarchical B-splines allow multiresolution editing and maintain a high level of continuity, but the local details can not be translated, rotated, or resized.

Surface pasting, developed by Bartels and Forsey [3], is a generalization of hierarchical B-splines that allows the insertion of local detail to a tensor product surface without changing the structure of the underlying surface. In surface pasting, the area of local detail is represented as a tensor product surface, called a feature. The feature is placed on an existing surface, called the base, to produce a composite surface. Additional features can be pasted hierarchically on the composite surface to create more complex composite surfaces. Surface pasting has been integrated into Side Effects' *Houdini* software, where it has been successfully used for character animation.

Although surface pasting is a hierarchical modelling method, the user interfaces implemented for previous research have concentrated on positioning the features upon the base surfaces, and adjusting—i.e., translating, rotating, and resizing—the features once they have been pasted. In this paper, we present a technique for the direct manipulation of pasted surfaces that allows the user to edit the surfaces at any resolution in the hierarchy.

§2. Background

A tensor product B-spline surface is a piecewise polynomial surface that is defined by a grid of control points $\{P_{i,j}\}$ and a set of basis functions $\{N_{i,j}\}$:

$$S(u,v) = \sum_{i=1}^{M} \sum_{j=1}^{N} N_{i,j}(u,v) P_{i,j}.$$

Here $N_{i,j}(u,v) = N_i^m(u) N_j^n(v)$, and the N_is and N_js are the degree m and n B-splines for the two parametric domain directions. For a more detailed introduction to B-splines, see any introductory splines text, such as Farin's [6].

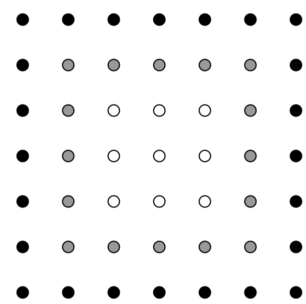

Fig. 2. Outer two layers of feature control points located at Greville points.

The surface pasting process (illustrated in Figure 1) is a computationally inexpensive method for adding local detail to tensor product surfaces. First, the feature's domain is embedded into its range space. Next, a local coordinate frame $\mathcal{F}_{i,j} = \{\mathcal{O}_{i,j}, \vec{r}_{i,j}, \vec{s}_{i,j}, \vec{t}_{i,j}\}$ is constructed for each feature control point $P_{i,j}$. The origin of the coordinate frame $\mathcal{O}_{i,j}$ is the embedded Greville point corresponding to $P_{i,j}$. The vectors $\vec{r}_{i,j}$ and $\vec{s}_{i,j}$ are the two parametric domain directions for the feature and $\vec{t}_{i,j} = \vec{r}_{i,j} \times \vec{s}_{i,j}$. Each feature control point is represented as a displacement vector $\vec{d}_{i,j}$ expressed relative to $\mathcal{F}_{i,j}$ by subtracting $\mathcal{O}_{i,j}$ from $P_{i,j}$:

$$\vec{d}_{i,j} = P_{i,j} - \mathcal{O}_{i,j} = \alpha_{i,j}\vec{r}_{i,j} + \beta_{i,j}\vec{s}_{i,j} + \gamma_{i,j}\vec{t}_{i,j}.$$

Next, the feature's domain is mapped into the base's domain. For each feature control point, the corresponding Greville point is mapped into the base's domain. The base surface is evaluated at this domain point to form a new local coordinate frame $\mathcal{F}'_{i,j} = \{\mathcal{O}'_{i,j}, \vec{r}'_{i,j}, \vec{s}'_{i,j}, \vec{t}'_{i,j}\}$. The control point's new origin on the base surface is $\mathcal{O}'_{i,j}$; the vectors $\vec{r}'_{i,j}$ and $\vec{s}'_{i,j}$ are the two partial derivatives at $\mathcal{O}'_{i,j}$, and $\vec{t}'_{i,j} = \vec{r}'_{i,j} \times \vec{s}'_{i,j}$. The feature control point is placed by expressing its displacement vector relative this new local coordinate frame:

$$P'_{i,j} = \mathcal{O}'_{i,j} + \vec{d}'_{i,j}$$

where

$$\vec{d}'_{i,j} = \alpha_{i,j}\vec{r}'_{i,j} + \beta_{i,j}\vec{s}'_{i,j} + \gamma_{i,j}\vec{t}'_{i,j}. \tag{1}$$

To ensure that the boundary of the feature lies near the base surface, we place the first layer of the feature's control points (the black points of Figure 2) at the Greville points so they have zero displacement vectors. After pasting, these feature control points will lie on the base surface, and the boundary of the feature will lie near the base. By inserting knots into the feature surface, the discontinuity between the feature and the base can be made as small as desired.

By placing the second layer of the feature's control points (the grey points of Figure 2) at the Greville points, we achieve an approximate C^1 join between the feature and the base. Conrad [5] gives a further discussion of continuity issues of pasted surfaces, and shows how to use quasi-interpolation to further reduce both the C^0 and C^1 discontinuity between the feature and the base.

§3. Direct Manipulation of Tensor Product Surfaces

Traditionally, B-spline curves and surfaces were manipulated by moving their control points. This method is unintuitive and requires that the control points be displayed, thus increasing the clutter on the screen.

Bartels and Beatty [2] developed a technique for the direct manipulation of spline curves in which users could pick any point on a curve, move it to a new location, and have the shape of the curve change appropriately. Their method found a set of control points that had maximal influence over the picked point. The amount that each control point moved was proportional to its influence over the picked point.

A number of researchers have investigated techniques for the direct manipulation of tensor product surfaces. For example, Fowler proposed a method for directly manipulating positions, normal vectors, and partial derivatives at any surface point [8]. He also found that the system of equations that must be solved to perform direct manipulation of tensor product surfaces is underdetermined.

We have chosen to calculate new control point locations using a generalization of Bartels and Beatty's curve manipulation technique [2] where the extra degrees of freedom are used to reduce the overall change in the position of the surface's control points. We have extended and altered their method so that it can be applied to the direct manipulation of tensor product B-spline surfaces.

Given a surface $S(u,v) = \sum_i \sum_j N_{i,j}(u,v)P_{i,j}$, suppose we want to move a surface point $S(\overline{u},\overline{v})$ by a vector $\overrightarrow{\Delta P}$, i.e., we want a surface S' such that

$$S'(\overline{u},\overline{v}) = S(\overline{u},\overline{v}) + \overrightarrow{\Delta P}. \tag{2}$$

Then a block of control points that has the most influence over the picked surface point is found. For each control point $P_{i,j}$ in this block, we calculate a weight $w_{i,j}$ that is proportional to the control point's contribution to the surface point $S(\overline{u},\overline{v})$:

$$w_{i,j} = \frac{N_{i,j}(\overline{u},\overline{v})}{\displaystyle\sum_k \sum_l \big(N_{k,l}(\overline{u},\overline{v})\big)^2}, \tag{3}$$

where the double summation is over the block of control points that we are modifying. Each of these control points is updated as

$$P'_{i,j} = P_{i,j} + w_{i,j}\overrightarrow{\Delta P}. \tag{4}$$

The net result is for the vector $\overrightarrow{\Delta P}$ to be distributed over the control points of S so that when S' is evaluated at $(\overline{u},\overline{v})$, the sum of the $w_{i,j}$s weighted by the basis functions is 1.

Bartels and Beatty showed that in the curve case, moving the single most influential control point produced unstable results. This instability occurs

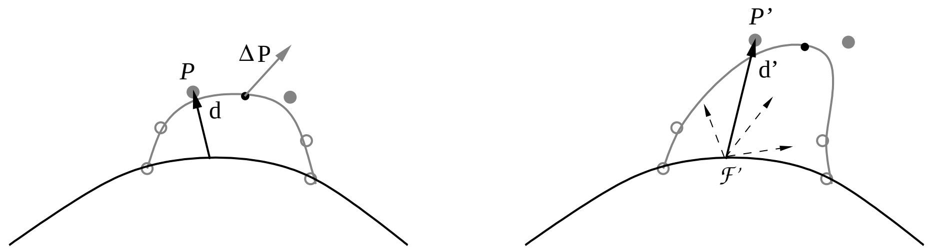

Fig. 3. Updating a control point's displacement vector.

because there will be a point on the curve to the left of which control point P_i will have the most influence, and to the right of which control point P_{i+1} will have the most influence. Picking near this division can have markedly different results depending on which side was picked. A minimum of two control points must be adjusted for their method to be stable.

Just as we need to move two control points in the curve case, we need to move at least two control points in each of the two parametric domain directions to get a stable direct manipulation method for surfaces. We chose to adjust a 2×2 block of control points since this small block size restricts the locality of change. A larger block may be used to modify a larger area of the surface.

§4. Direct Manipulation of Pasted Surfaces

When directly manipulating pasted surfaces, we would like to edit any point on the surface, and to edit the surface at any resolution in the hierarchy regardless of where in the hierarchy the selected point lies. In this section, we begin with a discussion of the extension of the direct manipulation technique to the top level of the pasting hierarchy, i.e., the level in which the selected point lies, and then we describe some problems with this simple approach. In the next section, we will extend this method to solve these problems, which will allow us to edit the surface at any resolution at or below the level of the selected point.

The basic technique for directly manipulating tensor product B-spline surfaces carries over to pasted surfaces with only minor modifications. The first step is to update the control points of the surface we wish to modify using direct manipulation of a tensor product surface as described in §3. Then we update the displacement vector $\vec{d}_{i,j}$ for each control point $P_{i,j}$.

To recalculate each control point's displacement vector, the local coordinate frame $\mathcal{F}'_{i,j} = \{\mathcal{O}'_{i,j}, \vec{r}'_{i,j}, \vec{s}'_{i,j}, \vec{t}'_{i,j}\}$ on the base surface must be reconstructed. The difference $\vec{d}'_{i,j}$ between the new control point location $P'_{i,j}$ and the origin of the coordinate frame $\mathcal{O}'_{i,j}$ is found and expressed in terms of $\mathcal{F}'_{i,j}$:

$$\vec{d}'_{i,j} = \alpha'_{i,j}\vec{r}'_{i,j} + \beta'_{i,j}\vec{s}'_{i,j} + \gamma'_{i,j}\vec{t}'_{i,j}.$$

This gives a 3×3 system of equations to determine the new $\alpha'_{i,j}$, $\beta'_{i,j}$, and $\gamma'_{i,j}$ for each updated displacement vector. After the new displacement vectors are

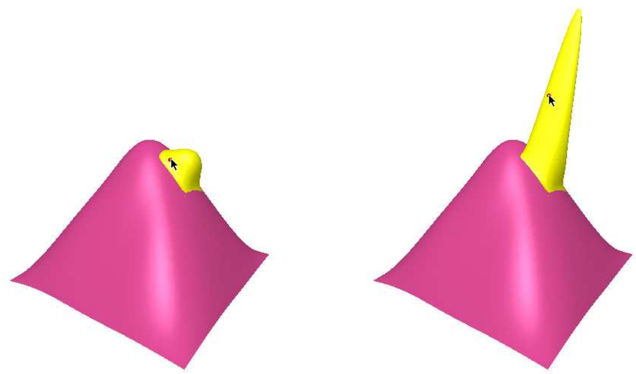

Fig. 4. Example of distortion.

calculated, the feature may be translated or the underlying surface may be changed, and the results of the direct manipulation are preserved.

In the remainder of this paper, we will talk about selecting a surface point $S(\overline{u}, \overline{v})$ and then manipulating points $S_i(\overline{u}, \overline{v})$, i.e., points at other levels in the pasting hierarchy with the same parameter values as the selected point. In effect, we are manipulating S_i at the image of $(\overline{u}, \overline{v})$ mapped into the domain of that surface according to the invertible domain transformations defined by the surface pasting operation.

If we select a point on one surface in a pasting hierarchy, and move the selected surface's control points using this direct manipulation technique, we often get the results we want. However, simply using this basic direct manipulation technique on the component surfaces in a pasting hierarchy results in several problems. The first problem is that if we modify the boundary control points of a surface, then we may lose the desired approximate continuity properties, i.e., the composite surface may stop looking smooth or the surface may detach from its underlying base. Thus, we need to fix the boundary and cross boundary derivatives of the manipulated surface by ensuring that the two outermost rings of the surface's control points (indicated in Figure 2) do not move.

If the user attempts to move a surface point whose most influential 2×2 block of control points intersects the two outermost rings of control points, we are faced with two choices. We can disallow the direct manipulation, or we can find the closest block that does not overlap the two outermost rings. In the latter case, the control points we would change have less influence on the selected surface point, and thus they must be displaced farther to move the picked surface point to its new location. This can cause unsightly distortions, similar to those described by Bartels and Beatty [2], in the area of the surface over which these alternative control points have a higher influence. These distortions would likely confuse the user since the maximal change in the surface would not occur at the picked point. An example of such a distortion can be seen in Figure 4. We chose to disallow direct manipulation of a pasted surface when the most influential 2×2 block of control points intersects the two outermost rings of the surface's control points.

A second problem with this simple method for directly manipulating pasted surfaces occurs when we try to edit at a different level in the pasting hierarchy than at the level containing the picked point. To truly edit the pasted surface at any resolution, we need to be able to select a point in a region of high resolution, and have the direct manipulation changes occur at a lower level of resolution. One simple way to implement this type of manipulation is to find the corresponding point $S_i(\overline{u},\overline{v})$ on a lower level surface when the user selects the surface point $S(\overline{u},\overline{v})$, and to use the $\overrightarrow{\Delta P}$ vector to directly manipulate $S_i(\overline{u},\overline{v})$. We would then reapply Equation (1) to each surface above S_i to update the composite surface.

Unfortunately, Equation (1) is not accounted for in the direct manipulation equations, and this method does not result in direct manipulation. To achieve direct manipulation of the composite surface, we will have to make additional adjustments to the control points of the surfaces between S_i and S, as described in the next section.

§5. Direct Manipulation of Hierarchical Pasted Surfaces

To manipulate a pasted surface at a lower resolution than the level at which the selected point lies, we need to modify the direct manipulation method. A first idea is to derive for the pasting hierarchy formulas similar to Equations (2), (3), and (4). When we expand Equation (2) for a pasted surface, we get

$$\sum_{i,j} N_{i,j}(\overline{u},\overline{v})\Big(\vec{d}\,''_{i,j} + \mathcal{O}''_{i,j}\Big) = \sum_{i,j} N_{i,j}(\overline{u},\overline{v})\Big(\vec{d}\,'_{i,j} + \mathcal{O}'_{i,j}\Big) + \overrightarrow{\Delta P}.$$

To derive the direct manipulation equations, we need to expand $\vec{d}\,''_{i,j}$ (Equation (1)) and $\mathcal{O}''_{i,j}$ (an evaluation of the base surface). Unfortunately, $\vec{d}\,''_{i,j}$ depends upon the control points of the base surface in a non-linear manner as we see from Equation (1) and the formula for $\vec{t}\,'_{i,j}$. Thus, this method of direct manipulation of a hierarchical pasted surface is more expensive than we would like, and as we increase the depth of the hierarchy, the equations become more complicated.

As an alternative method, we chose to modify the control points of multiple surfaces in the hierarchy. While this method is not truly hierarchical, it applies most of the change to a single surface, with other surfaces receiving only minor updates.

The idea behind our method is to push the work down the pasting hierarchy, make a large change at the desired level, and then ascend the hierarchy making small adjustments as needed. Suppose we have a hierarchy of pasted surfaces, $S_0,\ldots,S_H$, with S_0 being the coarsest resolution. Given a point $S(\overline{u},\overline{v}) = S_h(\overline{u},\overline{v})$ on a pasted surface at resolution h, suppose we wish to edit at resolution r, with $h \geq r$. Our method descends the hierarchy of surfaces under the picked point until we reach surface S_r. We adjust S_r so that $S_r(\overline{u},\overline{v})$ is moved by $\overrightarrow{\Delta P}$. This causes S_{r+1} to change (giving a new surface, S'_{r+1}), although $S_{r+1}(\overline{u},\overline{v})$ will not necessarily move by $\overrightarrow{\Delta P}$. We now

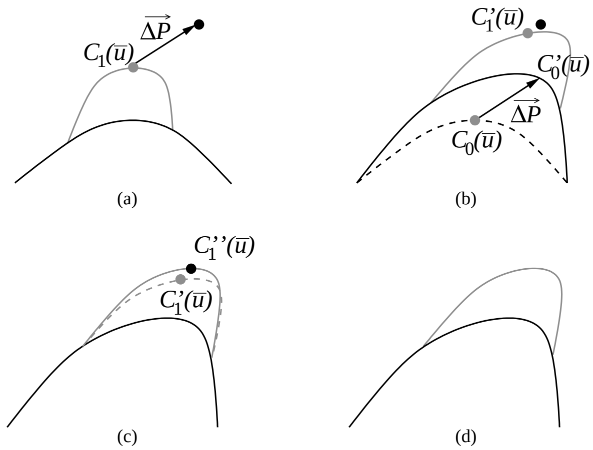

Fig. 5. Example of hierarchical direct manipulation.

compute the difference between the desired change and the actual change in S_{r+1}, giving a correction factor $\overrightarrow{\Delta P}' = \overrightarrow{\Delta P} - [S'_{r+1}(\overline{u}, \overline{v}) - S_{r+1}(\overline{u}, \overline{v})]$. Then, we directly manipulate S'_{r+1} by $\overrightarrow{\Delta P}'$, giving a new surface S''_{r+1} such that $S''_{r+1}(\overline{u}, \overline{v}) = S_{r+1}(\overline{u}, \overline{v}) + \overrightarrow{\Delta P}$. Note that if $S_i(\overline{u}, \overline{v})$ lies in the unmodifiable region of S_i (area of the surface where every point has a 2×2 block of control points that intersects the surface's two outermost rings of control points), then we skip it. This procedure is repeated for each surface up the hierarchy until S_h is reached.

We illustrate our method for a curve in Figure 5, in which a point on the top surface is selected and the desired resolution level is the base level; the construction for surfaces is analogous. Initially, the user has chosen a point $C_1(\overline{u})$ to manipulate and the amount $\overrightarrow{\Delta P}$ by which to move the point. In diagram (b), the point $C_0(\overline{u})$ is located and moved by $\overrightarrow{\Delta P}$. The shape of the base changes, resulting in the change of the feature's shape and location of the picked point, as shown in (c). Since $C_1'(\overline{u})$ is not at its desired position, a correction factor is applied, giving $C_1''(\overline{u})$. The last diagram shows the resulting composite curve.

Using our method, the user can select any point on the surface, and directly manipulate the surface, with modifications occurring at or below the level in the hierarchy on which the selected point lies. However, if the user selects a point in the unmodifiable region of the top level surface, then we do not allow the user to edit at the level of the selected surface, and only allow the user to edit at lower levels in the hierarchy. This restriction prevents the type of distortions discussed in §4.

Figure 6 illustrates a sequence of events that occurs when performing hierarchical direct manipulation on the composite surface. The selected point is located on the smallest and lightest coloured surface in the first image. The dark coloured surface(s) in the middle three images are affected in each stage of the manipulation. In each case, the bottommost dark surface undergoes the greatest change, while the other dark surfaces only have correction factors applied to them. The more surfaces that are dark coloured, the broader the change in the composite surface. In the second image, only the topmost surface is affected and the selected point was moved to the right and down. In the next image, two surfaces are affected and the selected point was moved up and to the right. Three surfaces are affected and the selected point is moved up and to the right in the fourth image. The last image shows the result of the manipulation.

Our direct manipulation method provides the user with complete feedback no matter what resolution level the user is editing. In contrast, Forsey's hierarchical B-splines editor strips the upper levels of detail out of the display as the user descends the hierarchy, making modelling more difficult.

§6. Conclusions and Future Work

Surface pasting is a method for constructing multiresolution surfaces by hierarchically composing tensor product B-spline surfaces. In this paper, we have shown how to implement direct manipulation of pasted surfaces, allowing the user to edit the composite surface at any level in the hierarchy, either at the selected level, or at a lower level.

Our direct manipulation method is not truly hierarchical since changes are made to multiple resolution levels. However, empirical testing has shown that the correction factors applied to the higher level surfaces tend to be much smaller than the initial $\overrightarrow{\Delta P}$ assigned to the selected resolution level. More research is needed to determine if there are any limits on the size of the correction factors as the pasting hierarchy is ascended.

An improvement to our hierarchical direct manipulation method would be to allow the user to increase the edit resolution for a picked point. For example, suppose there are n levels in the pasting hierarchy and the user wishes to manipulate the picked point at resolution $n+1$. Then a null surface with increased control point density could be pasted under the picked point, and this new surface, rather than the original one, is manipulated. The ability to edit a surface at a higher resolution would allow the user to add local detail to a region without explicitly pasting a feature surface.

Acknowledgments. We would like to thank Richard Bartels for his contributions to this work. Financial support was provided by NSERC and CITO.

References

1. Barghiel, C., R. Bartels, and D. Forsey, Pasting Spline Surfaces, in *Mathematical Methods for Curves and Surfaces,* Morten Dæhlen, Tom Lyche, Larry L. Schumaker (eds), Vanderbilt University Press, Nashville & London, 1995, 31-40.
2. Bartels, R. and J. Beatty, A Technique for the Direct Manipulation of Spline Curves, in *Proceedings of Graphics Interface '89,* S. MacKay, E. M. Kidd (eds), Canadian Human-Computer Communications Society, London ON, 1989, 33–39.

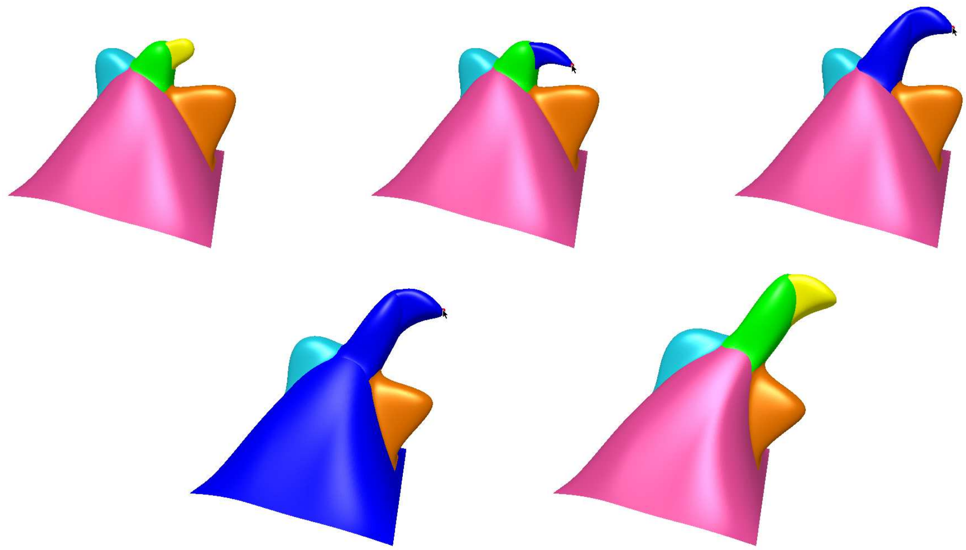

Fig. 6. Example of the direct manipulation of hierarchical pasted surfaces.

3. Bartels, R. and D. Forsey, Spline Overlay Surfaces, University of Waterloo Report CS-92-08, 1991.
4. Chan, L., S. Mann, and R. Bartels, World Space Surface Pasting, in *Proceeding of Graphics Interface '97,* W. A. Davis, M. Mantei, R. V. Klassen (eds), Canadian Human-Computer Communications Society, Kelowna BC, 1997, 146–154.
5. Conrad, B. and S. Mann, Better Pasting via Quasi-Interpolation, in *Curve and Surface Design: Saint-Malo 1999,* P.-J. Laurent, P. Sablonnière, and L. L. Schumaker (eds), Vanderbilt University Press, Nashville TN, 2000, 27–36.
6. Farin, G., *Curves and Surfaces for Computer Aided Geometric Design,* Fourth Edition, Academic Press, NY, 1997.
7. Forsey, D. and R. Bartels, Hierarchical B-spline Refinement, Computer Graphics **22(4)** (1988), 205–212.
8. Fowler, B., Geometric Manipulation of Tensor Product Surfaces, 1992 Symposium on Interactive 3D Graphics **25(2)** (1992), 101–108.
9. Ma, M., The Direct Manipulation of Pasted Surfaces, University of Waterloo Report CS-00-15, 2000.

Marryat Ma and Stephen Mann
Computer Science Department
University of Waterloo
Waterloo, Ontario N2L 3G1
`mma@cgl.uwaterloo.ca`
`smann@cgl.uwaterloo.ca`