

Feature-based Surface Decomposition for Polyhedral Morphing

Arthur Gregory,

Andrei State, Ming C. Lin, Dinesh Manocha, Mark A. Livingston

Department of Computer Science,
University of North Carolina at Chapel Hill
{gregory, andrei, lin, dm, livingst}@cs.unc.edu
<http://www.cs.unc.edu/~geom/3Dmorphing>

Abstract

We present a new approach for establishing correspondence for morphing between two homeomorphic polyhedral models. The user can specify corresponding feature pairs on the polyhedra with a simple and intuitive interface. Based on these features, our algorithm decomposes the boundary of each polyhedron into the same number of morphing patches. A 2D mapping for each morphing patch is computed in order to merge the topologies of the polyhedra one patch at a time. We create a morph by defining morphing trajectories between the feature pairs and by interpolating them across the merged polyhedron. The user interface provides high-level control as well as local refinement to improve the morph. The implementation has been applied to several polyhedra composed of thousands of polygons. The system can also handle homeomorphic non-simple polyhedra that have holes.

1 Introduction

Image and object morphing techniques have gained increasing importance in the last few years. Given two objects, metamorphosis involves producing a sequence of intermediate objects that gradually evolve from one object to another. The techniques have been used in a number of applications, including scientific visualization, education, entertainment, and computer animation. Morphing, whether in 2D or 3D, consists of two basic phases, establishing a correspondence between the images or objects and interpolating between them, in conjunction with blending their colors or textures.

Main Contribution: We present a new approach for establishing correspondence for morphing between two homeomorphic polyhedra. Initially the user selects some corresponding elements called *feature pairs*. Although we borrow this term from previous morphing algorithms for images or 3D volumetric models [2, 26], our concept of a feature is closer to the sparse control mesh used in [11]. Our algorithm includes a simple and intuitive user interface for feature specification and automatically generates a *feature net*. Based on the feature nets, the algorithm decomposes the boundary of the polyhedra into morphing patches, computes a mapping for each morphing patch to a 2D polygon, merges them, and constructs a *merged polyhedron* whose topological connectivity contains both of the input polyhedra. In order to create a morph, the merged polyhedron has a *morphing trajectory* for each vertex to move along from one input polyhedron to the other. The overall complexity of the algorithm is $O(K(m+n))$, where K is a user-defined constant and m, n correspond to the number of vertices in the two input polyhedra. Some of the principal attributes of our approach are:

- **Simple User Interface:** The users only need to select a few corresponding pairs of vertices on the two polyhedra to define the feature nets. The trajectories along which these features travel during morphing are currently represented as Bézier curves.
- **Fine User Control:** The algorithm not only provides the user with high-level control in terms of specifying the features and trajectories, it also provides a mechanism to *locally refine* the morph or the animation sequence.
- **Generality:** The algorithm is applicable to all simple or genus zero polyhedra and makes no assumptions about convexity or star-shape. Furthermore, it can also be applied to *non-simple* homeomorphic polyhedra.

Organization: The rest of the paper is organized in the following manner. In Section 2, we survey related work in this area. We give an overview of our approach in Section 4, present the algorithm to compute the correspondence in Section 5, and address the morphing trajectories and other interpolation issues in Section 6. We describe the user interface in Section 7. Then, we discuss the implementation in Section 8 and also highlight its performance on different models. We analyze the algorithm in Section 9 and mention areas for future research in Section 10.

2 Related Work

The problems of establishing correspondence between graphical objects for shape transformation and morphing have been widely investigated in computer animation,

graphics and computational geometry for more than a decade. Most of the earlier work in the area focused on image metamorphosis, though a number of approaches have also been proposed for 3D volumetric models and 3D polyhedral models. An extended abstract of this paper appears in [14].

2.1 2D Morphing

Given two images, the problem of constructing a metamorphosis from one image to the other has been extensively studied in computer graphics and image processing. The set of algorithms can be classified into those that operate on raster images [2, 10, 25, 41, 42] and those that operate on vector-based or geometric representations [15, 35, 36, 40]. A common procedure to establish correspondence is to place a mesh over the entire image, which implies a bilinear warping function. Warping of these features can be controlled by energy minimization [36] or feature-based constraints [35]. Beier and Neely presented an elegant feature-based approach [2]. The features may correspond to points or lines [2, 39] or to snakes [25]. Mappings that have been used include (bi)linear mappings [2, 39] and spline-based mappings or free-form deformations [25] in conjunction with a weighting function which effectively controls the range over which a feature has influence. Ranjan and Fournier [32] presented an approach which uses circles to partition the objects. Other algorithms for transforming one image into another are based on two-dimensional particle systems to map the pixels [34].

It is possible to generate 2D images from a 3D model and apply 2D morphing algorithms to these. In this case, the intermediate stages of the morph are images. For many applications in animation and design, the 3D models themselves should be transformed and not their images [6, 21]. Furthermore, if the viewpoint or the lighting parameters are changed, the 2D morph has to be recomputed. On the other hand, 3D morphing is independent of viewing or lighting parameters. Having a 3D representation also allows the use of computer animation techniques such as keyframing.

2.2 3D Volume Morphing

Given two volumes, 3D volume morphing involves producing a sequence of volumes to transform them. A number of approaches have been published [7, 9, 16, 26, 31]. These include use of Fourier transforms that warp linearly in Fourier space [19]. This is a simple approach and requires minimal user specification, but inhibits intuitive understanding of the morph. Leros et al. [26] have presented a 3D extension of Beier and Neely's [2] approach. It allows the user to specify a set of features and permits fine user control.

3D polyhedral models can be voxelized to allow use of 3D-volume morphing. However, the intermediate stages of the morph are volumes and converting them into

geometric models produces topologically complex objects. Given a geometric surface description of the model, we can use current graphics systems for fast rendering as well as utilize different geometric algorithms for applications such as physically-based simulation or 3D object manipulation. Moreover, for feature-based approaches, it is simpler and more intuitive to design a user interface based on the geometric model as opposed to the volumetric model. For example, a user can pick any vertex, edge, face or contour of the original polyhedron as a feature, which is not easy with a voxel-based representation. Therefore, approaches based on voxelization followed by 3D volume morphing have limitations as well.

2.3 3D Shape Transformations and Metamorphosis

Several approaches related to establishing correspondence between 3D polygonal objects for shape transformation and metamorphosis have been proposed. Physically-based modeling techniques based on deformations [1, 37] and particle systems [33] can be used for object metamorphosis. Hong et al. [18] present an approach for polyhedral models that matches the faces whose centroids are closest. Chen and Parent [6] present an algorithm to transform piecewise linear 2D contours and extend it to 3D cylindrical objects. Bethel and Uzelton [4] add degenerate vertices and faces to two polyhedra until they have a common vertex neighborhood graph. Kaul and Rossignac [21] transform a pair of polyhedra by using their Minkowski sums. Hodgins and Pollard [17] have presented an algorithm for interpolating between control systems of dynamic models. Wyvill [43] has described an approach for implicit surfaces. Parent [30] has presented an approach that splits the surface of the model into pairs of sheets of faces and recursively subdivides them until they have the same topology. Parent [29, 30] has also described a method for deformation of polyhedral objects based on implicit functions. Kent et al. [22, 23] have presented a shape transformation algorithm for genus-0 polyhedra that involves projecting the models onto a sphere. Chen et al. [5] have produced 3D morphs of cylindrical images. Galin and Akkouche [13] have presented an algorithm for blob metamorphosis based on Minkowski sums. Lazarus and Verroust [24] have proposed a method based on skeletal curves. Kanai et al. [20] have presented an algorithm for shape transformation of genus-0 polyhedra using harmonic maps. DeCarlo and Gallier [11] have proposed a morphing technique that establishes correspondence by allowing the user to divide the surface into triangular and quadrilateral patches which can be projected onto a plane. Our overall approach shares their theme. However, we improve upon several restrictions in their technique making it easier for the user to specify correspondence between complicated models. For example, we remove the requirement that the user specified surface patches must be triangular or quadrilateral, and that can each be directly projected onto a plane.

3 Terminology

The term polyhedron refers to an arrangement of polygons such that two and only two polygons meet at each edge. It is possible to traverse the surface of the polyhedron by crossing its edges and moving from one polygonal face to another until all polygons have been traversed by this continuous path [28]. Topology refers to the vertex/edge/face connectivity of a polyhedron. Simple polyhedra are all polyhedra that can be continuously deformed into a sphere. Non-simple polyhedra are topologically equivalent to a solid object with holes in it. In this paper, we assume that each face of a polyhedron is homeomorphic to a closed disk. The genus g of a polyhedron is the maximum number of non-intersecting loops that do not divide its surface into two regions. Moreover, we assume that the given polyhedra satisfy the Euler-Poincaré formula: $v - e + f - 2(1 - g) = 0$, where v, e, f and g are the number of vertices, edges, faces and genus of the polyhedron respectively. The genus of a simple polyhedron is zero.

4 Overview

Given two homeomorphic polyhedra, our goal is to generate a morph that results in a smooth and gradual transition from one polyhedron to the other. One key aspect of our system is to allow the user to identify the important features of each polyhedron and specify a correspondence between them. The rest of the algorithm consists of a combination of techniques that can produce the desired result from the given user input.

Our algorithm decomposes the problem of morphing two polyhedra into morphing corresponding pairs of surface patches. Given the user's specification, the algorithm automatically partitions each polyhedron into a series of morphing patches, each of which is homeomorphic to a closed disk. Based on this decomposition, our approach is applicable to non-simple polyhedra as well. There are many other advantages of this approach. It is simpler to compute a 2D parameterization for one patch at a time. We use such mappings to merge the topologies of the polyhedra. Moreover, it allows us to use a number of algorithms from computational geometry. These include computing arrangements of lines, triangulations of polygons and planar straight-line graphs, and determining point locations in planar subdivisions. All these techniques are used for establishing correspondence between the two polyhedra.

Geometric algorithms are prone to robustness and accuracy problems. These involve dealing with degenerate configurations and inaccuracy problems due to finite precision arithmetic. Since an important component of the morphing algorithm is to merge the topologies of two polyhedra, we need to make sure that the algorithm maintains valid data structures and topology at each stage. In order to develop a robust implementation

of the algorithm, we have at times opted for simpler geometric algorithms, which may not have the best asymptotic performance.

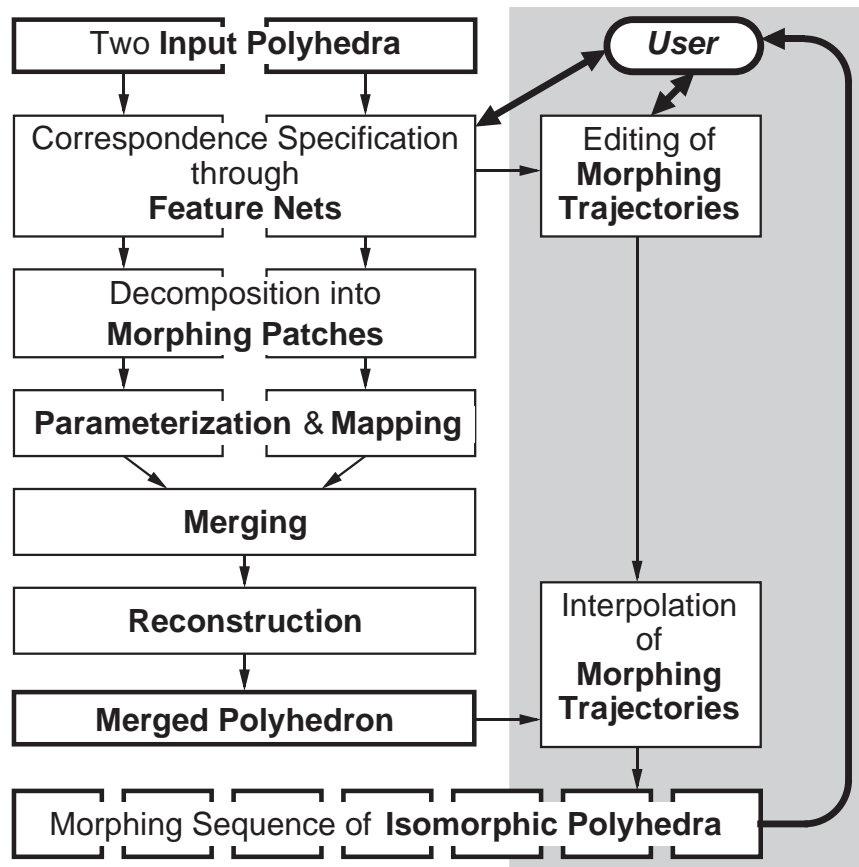


Figure 1. Overview of the polyhedral morphing algorithm.

An overview of our approach is given in figure 1. Given the user input, the algorithm consists of two phases: establishing a correspondence between the two polyhedra and interpolating between corresponding vertex locations.

4.1 Correspondence

- **Feature Net Specification:** The user specifies a network of corresponding chains on the surfaces of the two input polyhedra by specifying the vertices of their endpoints as shown in figures 2 and 3. The interior edges of the chains are then computed as the shortest path between the specified endpoints. The feature net is a sub-graph of the vertex/edge connectivity graph of each polyhedron.
- **Decomposition into Morphing Patches:** Based on the feature nets, the algorithm decomposes the surface of each polyhedron into the same number of morphing patches, each being homeomorphic to a closed disk.

- **Mapping:** A pair of corresponding morphing patches are mapped to a 2D polygon.
- **Merging:** The algorithm merges the topological connectivity of morphing patches in the 2D polygon.
- **Reconstruction:** Using the results from merging, the algorithm reconstructs the facets for the new morphing patch and generates a merged polyhedron with the combined topologies of the original two.
- **Local Refinement:** The user can make local changes to the feature net, such as splitting chains, moving extremal vertices, deleting chains or extremal vertices, or adding new ones, and then re-compute the merged polyhedron.

4.2 Interpolation

- **Trajectory Specification:** The user specifies the trajectories for the vertices of the feature net to follow during the morph. The morphing trajectories for the remaining vertices of the merged polyhedron are computed from these.
- **Morph Generation:** The algorithm makes use of the trajectories and interpolates the surface attributes to generate a morph.
- **Local Control:** The user can modify the trajectories and generate a new morph. This step does not involve re-computation of the merged polyhedron as shown in the shaded "feedback loop" of figure 1.

5 Correspondence between Polyhedra

In this section, we present the algorithm to compute the correspondence between two polyhedra. Given two homeomorphic polyhedra, \mathbf{A} and \mathbf{B} , we represent their vertices as $\mathbf{V}^A = V^A_1, V^A_2, \dots, V^A_m$ and $\mathbf{V}^B = V^B_1, V^B_2, \dots, V^B_n$, respectively. Superscripts represent the corresponding polyhedron. The edges and faces of the polyhedra are represented as $\mathbf{E}^A, \mathbf{E}^B, \mathbf{F}^A$ and \mathbf{F}^B , respectively. The output is a merged polyhedron with the topology of both input models, for which each vertex has a location on the two input models.

The algorithm ensures that each face of the input and output polyhedra is a triangle. Otherwise, the algorithm triangulates the face. The boundary and topology information for each polyhedron is represented using a *winged-edge* representation. Given the polyhedra, the algorithm computes a circularly ordered set of edges for each vertex. For each edge, the algorithm stores the following information: incident vertices, left and right adjacent faces, and preceding and succeeding edge in counter-clockwise order. Along with the winged-edge representation, the algorithm maintains an *adjacency graph*. The vertices and edges of the adjacency graph represent the vertex/edge connectivity information of the polyhedron. We will use the symbols \mathbf{G}^A and \mathbf{G}^B to represent the

adjacency graph of two polyhedra. Furthermore, we assign a weight to each edge of this graph. The weight corresponds to the Euclidean distance between the two vertices defining the edge. Based on the user's specification, the correspondence algorithm marks some of the vertices and edges in these graphs. To start with, each edge and vertex in these graphs is unmarked.

To illustrate the correspondence algorithm described in this section, we will make use of Figs. 2-10. In our illustration, polyhedron **A** corresponds to a model of an igloo and polyhedron **B** corresponds to a model of a house (shown in Fig. 2). Upper case letters denote objects in 3D and lower case letters represent objects in 2D.

5.1 Specifying Corresponding Features

The user selects a pair of unmarked vertices on each of the input polyhedra to be in correspondence denoted by $\{V^A_{i1}, V^A_{i2}\}$ and $\{V^B_{i1}, V^B_{i2}\}$, respectively (shown in Fig. 2). The algorithm computes a shortest path between these vertex pairs in the adjacency graph using only unmarked vertices and edges. Let the shortest paths correspond to: $\{V^A_{i1}, V^A_{j1}, \dots, V^A_{jk}, V^A_{i2}\}$ and $\{V^B_{i1}, V^B_{j1}, \dots, V^B_{jl}, V^B_{i2}\}$, as shown in Fig. 2. All the intermediate vertices and edges on the shortest paths in each graph are marked. We shall call such a shortest path as a *chain*. Moreover, the user-selected vertices are referred to as the *extremal vertices* of a chain.

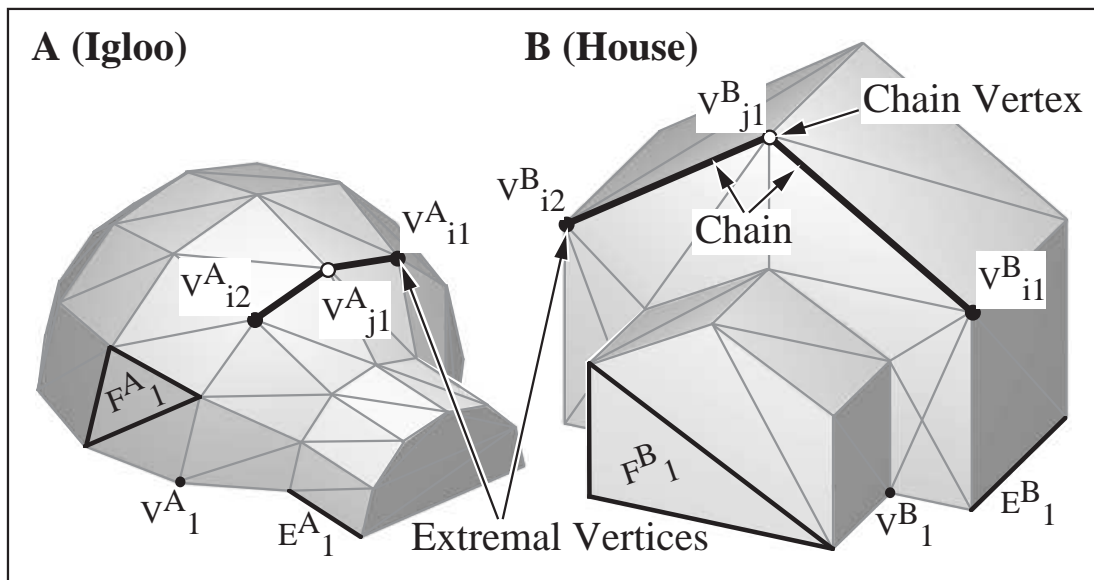


Figure 2. Input polyhedra with user-specified correspondences.

The selected vertices and chains are used to formulate a *feature net* for each polyhedron. We will represent the feature nets as N^A and N^B . They are sub-graphs of G^A and G^B , respectively. The user needs to specify a sufficient number and arrangement of chains in order to partition the boundaries of the polyhedra. The algorithm imposes some

constraints on the user. Each extremal vertex must be adjacent to at least two chains, and each chain must have a connected patch on each side. As a result, N^A and N^B have the same number of chains and extremal vertices and the user has specified a mapping between each extremal vertex and chain. In this way, the two feature nets define a bijection.

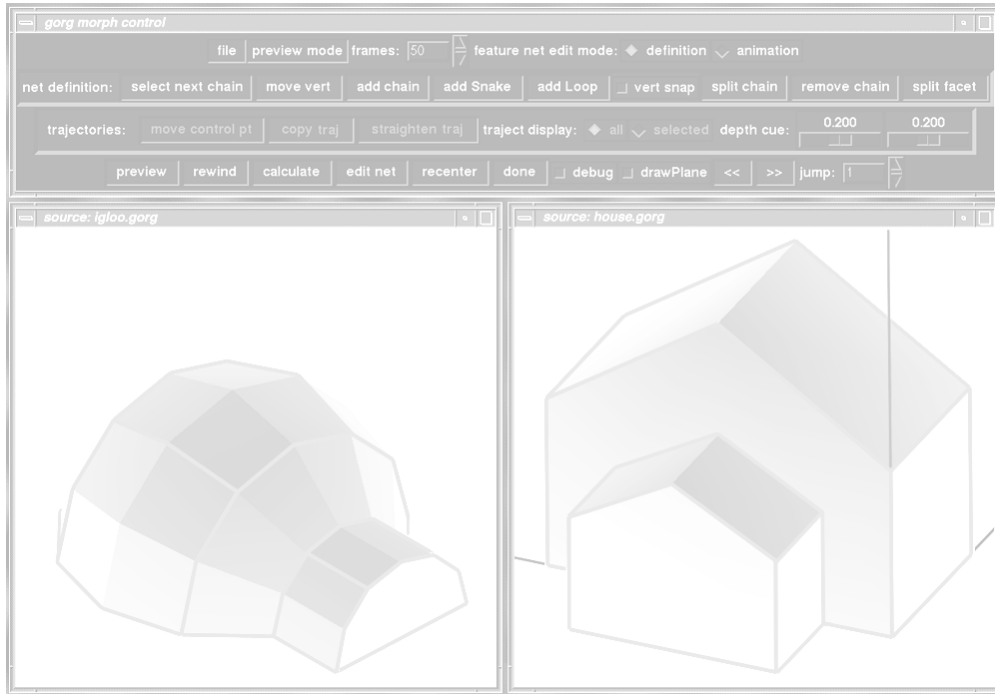


Figure 3. User interface for igloo-house morph showing completed feature net (red) with morphing patches.




Figure 4. Two selected morphing patches.

5.2 Decomposition into Morphing Patches

The feature nets are used to decompose the boundary of each polyhedron into the same number of *morphing patches*. A morphing patch (shown in Fig. 4) is simply a subset of a polyhedron that is homeomorphic to a closed disc, thus simplifying the geometric computations necessary to compute a morph. The vertices and edges are partitioned into exterior and interior vertices and edges. The exterior vertices of the morphing patches are those on the specified feature net (Fig. 4).

The decomposition algorithm has two steps. First, the perimeters of the morphing patches are computed by traversing the feature nets. Second, the interiors of the morphing patches are computed. Below is an overview of the partitioning algorithm.

```
Partition polyhedron into morphing patch perimeters(){
  For each extremal vertex in the feature net VI {
    For each chain Cj adjacent to VI that has no adjacent facet clockwise from VI {
      1. OppVert = the extremal vertex at the other end of Cj from VI
      CurrentChain = Cj
      While (OppVert != VI ){
        2. CurrentChain = the next clockwise chain adjacent to OppVert from
           CurrentChain
        3. OppVert = the extremal vertex at the other end of NextChain from OppVert
      }
    }
  }
}
```

Algorithm to partition Feature Net into Morphing Patch Perimeters

The computation of the perimeter, as detailed in the above pseudo-code, uses the circular ordering of edges at each vertex. Beginning at a vertex and chain of the feature net, the algorithm walks through the tightest clockwise loop of extremal feature net vertices it can find. From the first extremal vertex and chain, it moves to the extremal vertex at the other end of the chain (step 1). Next it uses the circular ordering of edges at that vertex to proceed to the closest clockwise chain (step 2). Then it follows that chain to the vertex at the other end (step 3). This process continues until it comes back to the original vertex, hence traversing the perimeter of a morphing patch. Note that this is possible because the chains may not cross. Since there is a bijection between the feature nets, this process is performed simultaneously on both of the input polyhedra. For example, in step 2, the next corresponding pair of clockwise chains can be determined by examining the underlying graph of only one of the input models.

The interior of a morphing patch is computed with a *depth-first search* algorithm modified to traverse through faces of the graph instead of vertices. It starts with an arbitrary edge on the perimeter of the morphing patch, and determines which adjacent face is interior to the patch by choosing the one whose ordering is the same as the exterior vertices. Then it crosses the face that is on the interior of the morphing patch and recursively branches out to the faces on its other two edges. The recursion stops at an edge that has already been traversed or is part of the perimeter.

This process is repeated until all vertices, edges, and faces of the original polyhedra have been partitioned into morphing patches. These morphing patches are represented as $\mathbf{P}^A_1, \dots, \mathbf{P}^A_K$ and $\mathbf{P}^B_1, \dots, \mathbf{P}^B_K$. The algorithm ensures that each morphing patch has no holes in the interior and is *homeomorphic* to a closed disk. In other words, it is bounded by a single loop of adjacent vertices.

Non-simple Polyhedron: For a genus zero polyhedron, no morphing patch could contain a hole. If the input polyhedron has genus greater than zero, a user needs to specify the features in such a manner that each morphing patch is homeomorphic to a disk. Such a decomposition is always possible, as shown in Fig. 11 for a cup and torus.

5.3 Mapping

Given a morphing patch, our goal is to compute a parameterization over a convex polygonal region in 2D. Construction of a parameterization for complex shapes over simple domain is an important problem that occurs in different applications.

5.3.1 Desiderata

A mapping algorithm requires a parameterization that is intuitive from the user's point of view. When the user describes a pair of corresponding morphing patches on the input polyhedra, s/he should be able to intuitively imagine how the interiors will be mapped without any detailed understanding of the system. After some experimentation it appears that we want the mapping to have the following property: Given any two triangular faces F^A_i and F^A_j of a morphing patch, let them map to the triangles f^A_i and f^A_j , respectively, in the 2D polygon. The ratio of the areas between F^A_i and F^A_j should be close to that of f^A_i and f^A_j . Individual triangles are not as important here as the fact that the mapping minimizes the distortion of the area across the patch as a whole. This leads to a more predictable morph for the interiors of corresponding morphing patches. For example, in figure 5, one would not expect that the corresponding morphing patches to split apart during the morph, but instead one should bend into the other. The solution is the equivalent of taking a uniform coordinate system on the 2D polygon as if it were composed of rubber sheets with no potential energy, and placing it on the surface of the 3D morphing patch such that the energy is minimized.




Figure 5. Comparison between the area preservation and harmonic mappings. A corresponding pair of morphing patches A and B are mapped to a 2D polygon using both a harmonic mapping, and an area preservation mapping. The two mappings are then merged and reconstructed, and the resulting fifty percent morphs are shown for the patch, HM and AM, respectively. Note that from a user standpoint one would expect the intermediate patch to look more like AM than HM. During the morph of AM, patch A appears to bend over into patch B, whereas during the morph of HM the tip of patch A shrinks into the side of patch B while the tip of patch B grows out of patch A's side.

5.3.2 Previous Approaches

A number of algorithms have been proposed by Kent et al. [22], Maillot et al. [27], and Eck et al. [12]. One possible solution is to use harmonic maps. They minimize the metric distortion and preserve the aspect ratios of the triangle, but can introduce area compression [12]. In figure 5, it is area compression that causes the intermediate model, HP, to have two peaks instead of one. A harmonic map may not produce a desirable mapping. It would treat the morphing patch as if it were composed of triangular rubber sheets with no potential energy in 3D, and minimize their total energy after being placed into 2D.

5.3.3 Our Approach

We currently use a divide and conquer approach with an area preservation heuristic. An example is shown in figure 6. Note that if the ear had been mapped using a harmonic map all the triangles would be very small in the middle, instead the areas are more uniform at the expense of distorting the shape.




Figure 6. A morphing patch containing an ear mapped to a 2D polygon with our algorithm

Given a pair of morphing patches, \mathbf{P}_i^A and \mathbf{P}_i^B , we compute a mapping from the surface of the patch to a regular 2D polygon. Let these morphing patches consist of m_i

extremal vertices. We map the morphing patch into a regular 2D polygon, inscribed in the unit circle, with m_i edges. We represent the regular 2D polygon as p_i .




Figure 7. The morphing patches are mapped to regular 2D polygons.

To compute a mapping, our algorithm first establishes a bijection between the chains of the two feature nets by splitting edges on the respective chains on the two models. The splitting criterion is based on edge lengths. The splitting of edges does not require the splitting of faces into triangles. After splitting, the algorithm has ensured that the corresponding patches have the same number of exterior edges and vertices.

The extremal vertices of the morphing patches are mapped to the vertices of p_i . Each chain of the morphing patch is mapped to an edge of p_i . All the external vertices lying in the interior of a chain are mapped onto the edges of p_i . The 2D coordinates of the vertices along the chains are interpolated based on the arc length of the chain.

The next step is to compute a mapping for the interior vertices of P_i^A and P_i^B . We use a simple recursive technique that tries to preserve the ratio of areas of the triangles based on a greedy heuristic. The algorithm divides a morphing patch by selecting two exterior vertices V_i^A and V_j^A that do not lie on the same chain. Next it computes a shortest path across the interior of the morphing patch. This path is then adjusted so that the ratio of the surface area on each side is as close as possible to the ratio that it will have once mapped to the 2D polygon. The vertices and edges lying on this path are mapped to the interior of p_i , along the segment connecting V_i^A and V_j^A . Based on this path, the algorithm recursively divides the morphing patch and maps the sub-patches to p_i . This process is then repeated for P_i^B . A pseudo-code description of the algorithm is given below:

```

MapInterior(Perimeter){
  if Perimeter surrounds any vertices (encompasses > 1 facet){

```

- 1 *SplitPath* = the shortest path btw two Perimeter vertices
that approximately divides the Perimeter in half;
 - 2 If there is no subdividing path between Perimeter vertices
vertices then split the Perimeter edges until there is one;
 - 3 Optimize *SplitPath* to preserve the area ratio;
 - 4 Interpolate the 2D coordinates at the endpoints of
SplitPath along its interior;
 - 5 *LPerimeter* = *SplitPath* + part of Perimeter on its left;
RPerimeter = *SplitPath* + part of Perimeter on its right;
 - 6 *MapInterior(LPerimeter)*; *MapInterior(RPerimeter)*;
 - 7 Remove any extra vertices and edges from added in step 2.
- }
}

Algorithm to map the interior of a morphing patch to a 2D polygon.

The shortest path found in steps 1 and 2 tends to preserve the area ratio on each side from the 3D model to the 2D polygon if the portion of the morphing patch enclosed by the perimeter does not have much curvature. Otherwise, the algorithm modifies the path between V_i^A and V_j^A in step 3, until the ratio of the surface area of the morphing patch on either side of the path 3D is as close as possible to that in 2D.

At this point the algorithm has computed a parameterization for each morphing patch such that each triangular face F_j^A and F_k^B has been mapped to a corresponding triangle, f_j^A and f_k^B , respectively, in p_i . The next step of the algorithm is to compute a mapping for each interior vertex of P_i^A to P_i^B (and vice-versa).

Given an interior vertex of P_i^A , the algorithm locates the triangle, say f_k^B that contains the image of that interior vertex. Furthermore, the algorithm computes the corresponding point in f_k^B and represents it using barycentric coordinates in terms of vertices of f_k^B . The barycentric coordinates are then applied to the vertices of F_k^B to compute the corresponding point on P_i^B .

This process is repeated for all the interior vertices of P_i^A and P_i^B . The time complexity of this part of the algorithm depends on the complexity of locating the triangle for each interior vertex of P_i^A . A simple search procedure would be linear in the number of triangles. However, using efficient data structures for planar point location [3], which involves linear time pre-processing, this search time can be reduced to be logarithmic in the number of triangles.

5.4 Merging

The algorithm has so far produced mappings into p_i such that the vertices $V_j^A \Rightarrow v_j^A$, $V_j^B \Rightarrow v_j^B$, and edges $E_j^A \Rightarrow e_j^A$, $E_j^B \Rightarrow e_j^B$. The edges e^A and e^B will in general intersect. We compute the intersections, split the intersecting edges, and create new vertices (Fig. 8).




Figure 8. The 2D polygons are combined into one polygon with merged topology.

Let n_e be the total number of edges and k_e be the number of edge pairs that actually intersect. In the worst case, k_e can be $O(n_e^2)$. Efficient and optimal algorithms of complexity $O(n_e \log n_e + k_e)$ have been proposed by Clarkson and Shor [8] to compute the intersections. However, we are not aware of any robust implementations of these algorithms.

In our application, we encounter many degenerate edge configurations. These include almost co-incident edges and vertices. Motivated by simplicity and robustness, we used an algorithm of complexity $O(n_e^2)$ which checks all edge pairs for overlap. Since the intersection computations can fail on edges that are co-incident, we handle the case where edges lie on the same mapped path separately. Furthermore, to avoid creating an invalid topology, we calculate all the intersections, sort them and then create the edges into which they were split. After intersection computation and splitting, we denote the set of all vertices and edges in p_i by \mathbf{x}^{AB} and \mathbf{g}^{AB} , respectively.

5.5 Reconstruction

After computing the intersection of all the edges, the algorithm produces a planar straight-line graph (PSLG) [3] from those intersections. The PSLG is constructed from the \mathbf{x}^{AB} and \mathbf{g}^{AB} . The next step is to compute a triangulation of these PSLGs. Though

good theoretical algorithms of linear complexity are known [38], it is unclear if they can handle PSLGs (Fig. 7) that have almost collinear edges or have very small angles between them. To robustly handle such cases, we use a simple divide and conquer algorithm. From the counter-clockwise ordering of the winged edge data structure of one of the input models superimposed on the other, we recursively subdivide the connected edges and vertices from the merging step into the smallest counter-clockwise cycles possible. Since the nature of the intersections in the merging step guarantees these cycles are convex, from this point we can triangulate the regions and handle degeneracies. After this step has been performed on all the morphing patches we get a merged polyhedron, as shown in Fig. 10.

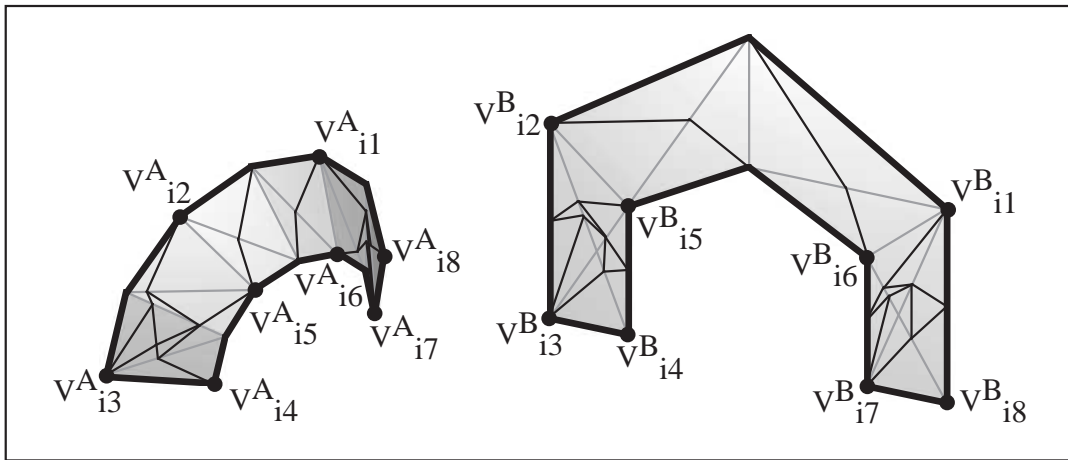


Figure 9. The output after merging and reconstruction steps.




Figure 10. The reconstructed merged polyhedron has the combined topology of the two input polyhedra.

6 Interpolation

At the end of the correspondence process, the merged polyhedron has the combined topological connectivity of polyhedra **A** and **B**. Each vertex on polyhedron **A** has a corresponding vertex on polyhedron **B**. In this section, we discuss issues for interpolating between the two polyhedra to generate a morph.

6.1 Aligning the Input Models

In many cases the location, orientation and size of the input polyhedra are quite different. For example, in Fig. 11, the igloo is much smaller than the house and has been positioned centrally on the ground, against the back wall of the house. Scaling, positioning, and orienting the input polyhedra with respect to each other must be performed by the user in preparation for 3D morphing and directly influences the appearance of the morph. This requirement is quite similar to the preparation required for 2D image morphing[2]. Note that if the user changes the relative scale, position or orientation of the two input polyhedra, only the morphing trajectories must be re-specified and interpolated (shaded feedback loop in Fig. 1); the correspondence specifications remain valid.

6.2 Morphing Trajectories

During morphing, the vertices travel from their positions on **A** to their respective positions on **B** along morphing trajectories. [22] suggested using Hermite interpolation between the corresponding vertices with the tangents pointing along normal directions. Similarly we allow the user to represent the trajectory as a Bézier curve for each pair of extremal vertices. Initially, the trajectories are specified by the user for each extremal vertex. The trajectories are represented as cubic Bézier curves and denoted by $B_V(t)$ for each vertex V . The two endpoints lie on **A** and **B**, respectively. The user specifies the tangent directions at each endpoint. Based on the tangents, the algorithm computes the control points for each Bézier curve using Hermite interpolation. The user can modify the tangents for each trajectory belonging to an extremal vertex of the feature net, as shown in Figure 11 (which shows the tangent vectors as green line segments). Starting from the morphing trajectories of the extremal vertices, the system computes trajectories for all other vertices of the merged polyhedron. Different methods are used for interpolating between chain vertices and interior vertices:

1. **Chain Vertices:** For each chain vertex of a morphing patch, only the two adjacent extremal vertices of the chain are used. The algorithm computes weighting factors based on their distances along the chain (or arc length) from the chain vertex.
2. **Interior Vertices:** For each interior vertex of a morphing patch, all extremal vertices lying on the boundary of that morphing patch are used. The algorithm computes

weighting factors for each extremal vertex based on their distances along shortest paths toward the interior vertex. The weighting function is similar to that of Beier and Neely[2].




Figure 11. S-shaped morphing trajectories for the extremal vertices of the feature nets. The highlighted (white) trajectory is shown with its Bézier control points, which are the fixed endpoints of the trajectory. The green control points can be moved by the user.

In both cases above, the weighting factors are applied to the tangent vectors at each endpoint of a trajectory. For trajectory endpoints on polyhedron **A**, the endpoint vectors on **A** are averaged. For trajectory endpoints on **B**, the endpoint vectors on **B** are averaged. This process results in interpolated tangent vectors for trajectories at each vertex, which are then used to compute the two inner Bézier control points (also shown in green in Figure 11). Note that this may result in morphing trajectories that are not straight lines even if the user specifies all extremal vertex trajectories as straight lines.

The speed at which a vertex travels along the morphing trajectory is determined by sampling based on the “frame” number in the morphing sequence. For example, if the morph is to have 100 frames then at frame 30, each point V will be at the position $B_V(0.3)$. Beyond this, the user may want to specify a non-linear mapping between the frame number and the value of t , in order to control the speed at which morphing takes place. The algorithm also allows the user to individually modify this mapping for

individual extremal vertices, in order to make some parts of the polyhedrons morph sooner or later than others. This is analogous to the techniques used for 2D image morphing [2].

6.3 Interpolating Surface Attributes

In addition to the morphing trajectories required by the algorithm, other attributes of the input polyhedra need to be interpolated to generate a good morph. These include vertex colors, lighting coefficients, normal vectors, etc. Interpolation of these surface attributes occurs during the mapping and merging steps. Separate values are computed for the attributes of **A** as well as for the attributes of **B**. During morphing between **A** and **B**, the attributes are linearly interpolated between their values corresponding to each polyhedron.

7 User Interface

The user interface is one of the most important aspects of a morphing system. Although it is easy for the user to conceptualize a morph between two objects, it can be rather difficult to design a system that allows the user to express this easily. Our system achieves this goal by allowing the user to draw the key correspondences on the surfaces of the input models, and to specify the paths that the corresponding features will follow during the morph, taking advantage of graphics hardware to allow real-time interaction.

The user specifies corresponding chains of the feature nets for input models **A** and **B** by selecting the chains' endpoints. In order to enforce a bijection between the two feature nets, the system requires the user to specify the feature net vertices in corresponding pairs. In the case that a path of unmarked vertices and edges between a chain's endpoints is not available, the system creates new vertices and edges by splitting the necessary face(s).

As a simple extension to creating a single pair of corresponding chains, we also allow the user to create "*multi-chains*" and "*loops*". When creating a multi-chain, after specifying the first corresponding feature vertex pair, each additional vertex pair the user specifies makes a pair of chains connecting to the last vertex pair. A loop is simply a multi-chain with the property that the last pair of corresponding vertices is connected to the first pair with an additional pair of chains.

Once the corresponding pairs of chains have been specified, we allow several techniques for *local* refinement. These include splitting a chain into two chains at a selected vertex on one of the input models, removing a chain, and moving extremal vertices of the feature net. Fig. 3 shows the user interface and the feature nets on the two input models.

After a corresponding pair of extremal vertices have been specified on polyhedron **A** and polyhedron **B**, the user can control the *morphing trajectory*. The user can position the tangents of the trajectory, whose endpoints are the location of the vertices on the two input models as shown in Fig. 11. By default, the morphing trajectory is a straight line. After the user specifies the feature net, the rest of the morph is calculated as explained in Sections 5 and 6. Another additional feature of the interface allows the user to easily fix it by *locally* refining the feature net. As already mentioned, it is not necessary to recalculate the merged polyhedron if the user only edits the morphing trajectories shown in Fig. 11.

8 Implementation and Performance

We have implemented the system in C++ using the OpenGL and Tcl/Tk libraries. It features a graphical user interface for specifying features and trajectories and for refining the morph.

The input polyhedra are specified in a shared vertex representation. The adjacency graph of each polyhedron is stored such that each vertex has a list of edges stored in counter-clockwise order, each edge contains the incident vertices and two facets, and each facet contains three vertices and three edges also stored in counter-clockwise order. Furthermore, the system ensures that each polyhedron has valid topology and that it satisfies the *Euler-Poincaré* formula.

When the user specifies the extremal vertices of a chain in the feature net, the system computes the path connecting them using Dijkstra's *shortest path* algorithm. It starts with one of the extremal vertices as the start vertex and incrementally computes shortest paths to other vertices of the polyhedron. It stops when it has computed the path to the other extremal vertex. Since the endpoints of chains are typically close and the shortest path consists of a few edges, the system can compute these paths fast enough for interactive response.

Our implementation also utilizes a number of geometric algorithms for triangulating planar straight-line graphs, edge intersections and for point location. As mentioned in Section 4, we have opted for simplicity and robustness over efficiency or algorithms with optimal asymptotic performance. Our current implementation can handle many degenerate configurations.

8.1 Performance Improvement

The merging algorithm described in Section 5.4 computes all intersections between the mapped edges of each morphing patch. Based on the decomposition algorithm described in Section 5.2, a morphing patch of a large polyhedron may consist of thousands of edges. The number of intersections (and thereby the combinatorial

complexity of the merged polyhedra) grows with the number of edges and in the worst case is a quadratic function of the number of edges. As a result, the merging and reconstruction steps can become a bottleneck in the overall computation. To overcome this problem, we subdivide each morphing patches into smaller sub-patches such that each sub-patch consists of at most Q edges. A typical value for Q in our implementation is 100.

We subdivide the patches using a recursive divide-and-conquer algorithm. It is quite similar to the mapping algorithm presented in Section 5.3 and starts with computing a path between the external vertices of a patch. The system computes a corresponding path on the other patch, and tries to preserve the ratio of the areas on either side of each chain. At the same it maintains the bijection between the feature nets. Hence it divides each patch into sub-patches. This procedure is applied recursively, till each sub-patch has less than Q edges.

The division of morphing patches into sub-patches improves the overall computation time as well as the size of the merged polyhedra, obtained after the merging and reconstruction steps. Furthermore, it also speeds up the interpolation algorithm.

8.2 Results

Our system has been applied to a number of complex polyhedral models and used to successfully create several morphs. These include simple polyhedra (shown in Fig. 10) as well as non-simple polyhedra corresponding to a torus and a cup (shown in Fig. 11). We have presented the results in Table 1. It includes the complexity of input models, their genus, the number of extremal vertex pairs specified and the number of morphing patches. The table also reports the times required by a user to specify the features and time to compute the merged polyhedra.

9 Analysis

In this section, we analyze our algorithm. We first present an asymptotic bound on its running time and after that analyze the results produced by it.

Assume we have two polyhedra with m and n vertices. Most of the algorithms for checking the topology of a polyhedra and constructing the winged-edge data structure and adjacency graphs take $O(m+n)$ time. Let the number of extremal vertex pairs specified by the user be k . For large models, k is much smaller as compared to m and n . The time to compute the feature nets is dominated by shortest path computation algorithm. The worst case time to shortest path algorithm can take $O(km + kn + m \log m + n \log n)$ time, but in practice it is much less than that because the length of shortest paths is typically small. The number of morphing patches can be at most $O(k)$. The

computation of feature nets and morphing patches involves use of depth first search, and its overall time is bounded by $O(k(m+n))$. After the subdivision algorithm, presented in Section 8.1, each morphing sub-patch can have up to Q edges and the number of morphing sub-patches can be $O((m+n)/Q)$. The complexity of merging and reconstructing each pair of morphing sub-patches is $O(Q^2)$. The interpolation algorithm needs to compute the length of the shortest path from each interior vertex of a morphing sub-patch to each extremal vertex of the morphing sub-patch. We make use of single source shortest path algorithms and compute the paths for each extremal vertex on the boundary of the morphing patch. In the worst case it can take $O((m+n) k \log Q)$ time. As result, the overall complexity of the algorithm is $O(K (m+n))$, where $K = \max\{k \log Q, Q\}$.

| Models | Triangles | | Output Triangles | Morphing Patches | User Specification Time | Time to Compute Merged Polyhedron |
|-----------------------------|-----------|--------|------------------|------------------|-------------------------|-----------------------------------|
| | | | | | | |
| House-Igloo (Fig. 3) | 82 | 40 | 214 | 10 | ~5min | < 1s |
| Human-Triceratops (Fig. 14) | 5,660 | 17,528 | 97,900 | 86 | ~6 hours | 2.5min |
| Human Heads (Fig. 12) | 3,426 | 4,020 | 32,520 | 67 | ~3 hours | 30 seconds |
| Donut-Cup (Fig. 13) | 4,096 | 8,452 | 61,701 | 50 | ~4 hours | 1 min |

Table 1: Performance of our algorithm on four different pairs of input polyhedra.

The table shows input and output model complexity, the times to specify the features and trajectories by the user and times to compute the merged polyhedron on a SGI Onyx 2 with 195 MHz R10000.

Our approach does not suffer from the ghosting problems seen in image and volume morphing [2, 26]. However, a similar scenario can occur, self-intersection. We can check for it automatically through collision detection algorithms, but this is prohibitively expensive. Furthermore, self-intersection may be desirable and in some cases even necessary in order to allow some part of a morphing object to reach its target position. Hence, it is currently the responsibility of the user to “re-route” portions of a morphing polyhedron by controlling the morphing trajectories to avoid self-intersections. (For example, we prevented the two sides of the triceratops’ flat tail in Fig. 14 from temporarily passing through each other in this way.)

The visual “quality” of a morph created with our system is quite subjective. As with conventional 2D image morphing, it is influenced primarily by the number of detailed

feature correspondences. Beyond that, controlling the morphing trajectory strongly contributes to a “smoothly flowing” appearance during morphing. Finally, similarly to 2D techniques, morphs between objects that are “similar” in appearance (for which one intuitively notices the corresponding pairs of features by simple visual inspection) result in smoother transitions than morphs between vastly different objects (cf. human heads morph in Fig. 12 vs. cup-donut morph in Fig. 13).

10 Current Limitations & Future Work

The specification of the feature net suffers from two limitations. First, the feature net must be connected. It would be nice to remove this restriction by automatically dividing the feature net into connected components, then adding enough chains to connect them such that each morphing patch is homeomorphic to a disk. Secondly, the chains connecting the feature net vertices are currently restricted to lie on the edges of the source models. It will be beneficial to allow the user to draw on the surface instead of the graph of the polyhedron. It would also be useful to remove the restriction that the input models have to be homeomorphic, one could imagine that the user merely specifies the correspondence, and then the system creates a polyhedron which can “appear” to have the topology of either of the input models.

One of the largest limitations of the current system is the user interface for controlling the morphing trajectories. It can be very difficult to control the shape of the intermediate models by positioning the tangents of a cubic curve. We envision a vast improvement would be to allow the user to actually sculpt the intermediate models, and have the system automatically calculate the morphing trajectories that include the points for the intermediate models sculpted by the user.

The performance of the mapping algorithm highlighted in Section 5.3 varies based on the triangulation of the morphing patch. Since it uses a greedy heuristic, the algorithm may not be able to preserve the area ratios. This can result in some noticeable distortions in the morph, especially when the mapping algorithm introduces area compression for one of the patches. Currently the user can work around this by adjusting the morphing trajectories. There is considerable literature in Graph Drawing on planar embedding of planar graphs. We plan to apply some of those techniques to our problem.

Texture coordinates can be interpolated together with other surface attributes of the input polyhedra. However, this will yield correct results only if both the polyhedra use a common texture map. In such a case, the mapped (unique) texture will seem to smoothly flow across the surface of the polyhedron during morphing. Our algorithm currently handles this case. In many cases, each of the input polyhedra may have different texture maps. One possible solution is to use a weighted blend between the textures of **A** and **B**,

controlled by the morphing interpolation factor t . This would result in an effect similar to an image fade-over between the textures of **A** and **B**. Our algorithm could easily be extended to handle this situation. However, we believe that a more powerful effect could be attained by allowing the user to perform a controlled conventional 2D morph in texture space together with the 3D polyhedral morph specified by our algorithm.

Morphing between animated models. In traditional 2D morphing, an image sequence of a moving actor can be morphed into another sequence showing a different moving actor. This is typically accomplished by (tediously) re-specifying correspondences for each pair (or at least for many pairs) of frames within the image sequences [2]. An extension of our approach could handle this problem in the following way: once correspondences have been specified for a computer-animated character, they can remain attached to the character's topology and carry over throughout the animation sequence. In other words, the correspondence features are animated together with the character. As for the morphing trajectories, they would have to be specified at least for the first and last "frames" of the animation and would have to be interpolated for all other frames.

11 Summary

We have presented a new approach for establishing a correspondence for morphing between two homeomorphic polyhedra, which includes a simple, intuitive user interface. It has been successfully applied to a number of simple and non-simple polyhedral models. We believe it is versatile enough to produce visually pleasing 3D morphs once coupled with an effective method for specifying the interpolation between the two models.

12 Acknowledgements

This research has been funded in part by an Alfred P. Sloan Foundation Fellowship, ARO Contract CAAH04-96-1-0257, NSF Career Award CCR-9625217, ONR young Investigator Award (N0014-97-1-0631), Honda, Intel, and NSF/ARPA Center for Computer Graphics and Scientific Visualization.

Thanks to Robert and Shirley Gregory for additional support, Rhinoceros for making the beta version of their modeling program freely available, Scott Nelson for the triceratops model we found on the web, Todd Gaul for video editing, and Henry Fuchs.

13 References

1. Barr, A., *Global and local deformations of solid primitives*. ACM Computer Graphics, 1984. **18**(3): p. 21--30.
2. Beier, T. and S. Neely, *Feature-based image metamorphosis*, in *Computer Graphics (SIGGRAPH '92 Proceedings)*. 1992. p. 35--42.

3. Berg, M.d., *et al.*, *Computational Geometry: Algorithms and Applications*. 1997: Springer-Verlag.
4. Bethel, E. and S. Uselton, *Shape distortion in computer-assisted keyframe animation*, in *State of the Art in Computer Animation*, N.M.-T.a.D. Thalmann, Editor. 1989, Springer-Verlag. p. 215-224.
5. Chen, D., A. State, and D. Banks, *Interactive Shape Metamorphosis*. Proc. of 1995 Symposium on Interactive 3D Graphics, 1995: p. 43--44.
6. Chen, E. and R. Parent, *Shape Averaging and its Applications to Industrial Design*. IEEE Computer Graphics & Applications, 1989. **9**(1): p. 47-54.
7. Chen, M., M.W. Jones, and P. Townsend, *Methods for Volume Metamorphosis*, in *Image Processing for Broadcast and Video Production*, Y.P.a.S. Wilbur, Editor. 1995.
8. Clarkson, K.L. and P.W. Shor, *Applications of random sampling in computational geometry, II*. Discrete Comput. Geom., 1989. **4**: p. 387--421.
9. Cohen-Or, D., D. Levin, and A. Solomovici, *Contour Blending Using Warp-Guided Distance Field Interpolation*, in *IEEE Visualization '96*. 1996, IEEE.
10. Covell, M. and M. Withgott, *Spanning the gap between motion estimation and morphing*. Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing, 1994. **5**: p. 213-216.
11. DeCarlo, D. and J. Gallier, *Topological Evolution of Surfaces*. Prof. of Graphics Interface'96, 1996: p. 229-235.
12. Eck, M., *et al.*, *Multiresolution Analysis of Arbitrary Meshes*, in *SIGGRAPH'95*. 1995. p. 173--182.
13. Galin, E. and S. Akkouche, *Blob Metamorphosis based on Minkowski Sums*. Computer Graphics Forum (Eurographics'96), 1996. **15**: p. 143-153.
14. Gregory, A., *et al.* *Feature-based Surface Decomposition for Correspondence and Morphing between Polyhedra*. in *Proceedings of Computer Animation*. 1998.
15. Guibas, L. and J. Hershberger, *Morphing Simple Polygons*, in *Proc. 10th Annu. ACM Sympos. Comput. Geom.* 1994. p. 267--276.
16. He, T., S. Wang, and A. Kaufmann, *Wavelet-Based Volume Morphing*. Proceedings of IEEE Visualization, 1994: p. 85-91.
17. Hodgins, J. and N. Pollard, *Adapting Simulated Behaviors for New Characters*. SIGGRAPH'97 Conference Proceedings, 1997: p. 153-162.
18. Hong, T., N. Magnenat-Thalmann, and D. Thalmann, *A general algorithm for 3D shape interpolation in a facet-based representation*, in *Proceedings of Graphics Interface '88*. 1988. p. 229--235.
19. Hughes, J., *Scheduled Fourier volume morphing*, in *Computer Graphics (SIGGRAPH '92 Proceedings)*. 1992. p. 43--46.
20. Kanai, T., H. Suzuki, and F. Kimura, *3D Geometric Metamorphosis based on Harmonic Maps*. Pacific Graphics, 1997: p. 97-104.
21. Kaul, A. and J. Rossignac, *Solid-interpolating deformations: construction and animation of PIPs*, in *Proc. Eurographics*. 1991. p. 493--505.
22. Kent, J., W. Carlson, and R. Parent, *Shape transformation for polyhedral objects*. SIGGRAPH '92 Conference Proceedings, 1992: p. 47--54.
23. Kent, J., R. Parent, and W. Carlson, *Establishing Correspondences by Topological Merging: A New Approach to 3-D Shape Transformation*. Proceedings of Graphics Interface '91, 1991: p. 271--278.
24. Lazarus, F. and A. Verroust, *Feature-based Shape Transformation for Polyhedral Objects*. Fifth Eurographics Workshop on Animation and Simulation, 1994.
25. Lee, S., *et al.*, *Image Metamorphosis Using Snakes and Free-Form Deformations*. SIGGRAPH 95 Conference Proceedings, 1995: p. 439--448.
26. Leros, A., C. Garfinkle, and M. Levoy, *Feature-Based Volume Metamorphosis*. SIGGRAPH 95 Conference Proceedings, 1995: p. 449--456.
27. Maillot, J., H. Yahia, and A. Veroust, *Interactive Texture Mapping*. SIGGRAPH'93 Conference Proceedings, 1993: p. 27--34.

28. Mortenson, M.E., *Geometric Modeling*. 1985, New York, NY: Wiley.
29. Parent, R., *Implicit Function Based Deformations of Polyhedral Objects*, in *Implicit Surfaces '95*. 1995.
30. Parent, R., *Shape Transformation by Boundary Representation Interpolation: A Recursive Approach to Establishing Face Correspondences*. *The Journal of Visualization and Computer Animation*, 1992. **3**(4): p. 219-239.
31. Payne, B. and A. Toga, *Distance field manipulation of surface models*. *IEEE Computer Graphics and Applications*, 1992. **12**(1): p. 65--71.
32. Ranjan, V. and A. Fournier, *Matching and Interpolation of Shapes using unions of circles*. *Computer Graphics Forum (Eurographics'96)*, 1996. **15**: p. 129-142.
33. Reeves, W.T., *Particle Systems -- a Technique for Modeling a Class of Fuzzy Objects*. *ACM Trans. Graphics*, April 1983. **2**: p. 91--108.
34. Rosenfeld, M., *Special Effects Production with Computer Graphics and Video Techniques*. SIGGRAPH'87 Course Notes #8, 1987.
35. Sederberg, T., et al., *2D Shape Blending: An Intrinsic Solution to the Vertex Path Problem*. SIGGRAPH '93 Conference Proceedings, 1993: p. 15--18.
36. Sederberg, T. and E. Greenwood, *A physically based approach to 2D shape blending*. SIGGRAPH '92 Conference Proceedings, 1992: p. 25--34.
37. Sederberg, T. and S. Parry, *Free-Form Deformation of Solid Geometric Models*. *Computer Graphics (SIGGRAPH '86 Proceedings)*, 1986. **20**: p. 151--160.
38. Seidel, R., *A simple and Fast Randomized Algorithm for Computing Trapezoidal Decompositions and for Triangulating Polygons*. *Computational Geometry Theory & Applications*, 1991. **1**(1): p. 51--64.
39. Seitz, S. and C. Dyer, *View Morphing: Synthesizing 3D Metamorphoses Using Image Transforms*. SIGGRAPH 96 Conference Proceedings, 1996: p. 21--30.
40. Shapira, M. and A. Rappoport, *Shape Blending Using The Star-Skeleton Representation*. *IEEE CG&A*, March 1995. **15**(2): p. 44--50.
41. Wolberg, G., *Digital Image Warping*. 1990: IEEE Computer Society Press.
42. Wolberg, G., *Skeleton-Based Image Warping*. *The Visual Computer*, 1989. **5**: p. 95--108.
43. Wyvill, W., *Metamorphosis of Implicit Surfaces*. SIGGRAPH'90 Course 23- Modeling and Animation with Implicit Surfaces, 1990.