Consistent Parametrization by Quinary Subdivision for Remeshing and Mesh Metamorphosis

Jian-Liang Lin Jung-Hong Chuang* Cheng-Chung Lin Chih-Chun Chen

Department of Computer Science and Information Engineering
National Chiao Tung University, Hsinchu, Taiwan, ROC.
{jllin, jhchuang, cclin, chihchun}@csie.nctu.edu.tw

Abstract

The vertex correspondence establishment among multiple objects is a versatile operation in computer graphics and geometry processing. We propose a systematic method called *recursive quinary subdivision* to efficiently find a dissection for a meshed object of genus-zero with little user input. The process can be easily extended to multiple objects, taking into account the alignment of extra feature points for applications such as mesh metamorphosis, to derive a common dissection. Based on the dissection and the parameterization associated with each resulting patch, uniform or adaptive remeshing can be performed to yield a set of semi-regular meshes. Moveover, geometric details can be easily resampled and stored as normal maps. We demonstrate the mesh metamorphosis application between two or more objects based on the vertex correspondence established by the common dissection and parameterization.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations, Hierarchy and geometric transformations.

Keywords: mesh dissection, parameterization, remeshing, metamorphosis, multiresolution modeling

1 Introduction

Recently, *semi-regular* meshes are getting more and more popular as representations of complex objects in computer graphics and geometric modeling. Such meshes are multiresolution representations formed by starting from a coarse irregular base domain and applying recursive regular refinement. Due to their regular structures, parameter and connectivity information can be predicted [Khodakovsky et al. 2000], and efficient tree or array based data structure can be used in such a way that only geometry information needs to be stored. Moreover, signal processing algorithms such as wavelet analysis can be employed [Eck et al. 1995; Lounsbery et al. 1997; Lounsbery 1994].

Remeshing is a process that for a given input irregular mesh, resamples its geometry information and constructs a semi-regular mesh which approximates the original irregular one. The state-of-the-art algorithms of remeshing involve initially dissecting the original irregular meshes into a set of topological disk-like patches, called *base domain*. These patches are later parametrized to compute a bijection between the 3D domain and the parameter domain. Obviously, the patch layout generated by initial dissection is a vital factor that dominates the quality of the remeshing.

A single mesh can be remeshed to yield a semi-regular one. Can we extend the idea to multiple objects? The answer is not true unless the objects have a common initial dissection in which each patch of one object corresponds exactly to one patch in all other objects. In consequence, they will possess the same base domain and their parametrizations will be consistent. The difficulty is that common initial dissections are not easily found, because a good dissection of one object might be bad for another object. Previous works [Praun et al. 2001; Michikawa et al. 2001; Kanai et al. 2000; Gregory et al. 1998; Zöckler et al. 2000] leave this problem to the users, requiring the user to provide a common initial dissection manually and many corresponding feature points. Applications such as metamorphosis (or morphing) and DGP (Digital Geometry Processing) applications [Praun et al. 2001], which require the establishment of correspondences among multiple objects, benefit from consistent parametrizations.

In this paper, we propose a systematic method called *quinary subdivision* to find a common initial dissection for multiple objects of genus-zero with little user intervention. The quinary subdivision scheme is a process that recursively subdivides an undesirable patch into five new patches, that possess better parametrization than their parent patch. The quinary subdivision scheme can be extended to multiple objects and guarantees to yield a common initial dissection. Extra feature correspondences between objects can also be provided by users and aligned during the quinary subdivision by using a foldover-free warping. After the common initial dissection is found, we begin the remeshing process and finally obtain a set of semi-regular meshes. The remeshing process can be either uniform or adaptive. Based on the parametrization, a set of *normal maps*, which capture geometry details, can also be easily resampled to improve the rendering quality of semi-regular meshes with coarse resolution. A multiresolution 3D mesh metamorphosis, with only a set of feature correspondences specified, is demonstrated as an application of the proposed approach.

*Contact author.

2 Related Work

Remeshing process begins with an input irregular connectivity model, dissect the model into a set of patches and then compute a parametrization for each patch. Eck et al. [Eck et al. 1995] proposed a method that produces a semi-regular mesh fully automatically by employing a Voronoi-like algorithm coupled with a Delaunay triangulation to yield the initial dissection, and parametrizing each patch using *harmonic mapping*. In contrast, the MAPS scheme proposed by Lee et al. [Lee et al. 1998] employed a mesh simplification algorithm to yield an initial dissection. Their algorithm is also automatic and more practical than Eck's. The *hierarchical conformal mapping* is performed during mesh simplification and a patch parametrization is obtained. Guskov et al. [Guskov et al. 2000] proposed a new remeshing process to construct a compact representation called a *Normal Mesh*. The above works focus on the cases of single object. Praun et al. [Praun et al. 2001] illustrated that shortest paths probably lead to problems for finding an initial dissection, and proposed a modified shortest path algorithm to trace curves between given feature points and yeild a base domain. They also focus on multiple objects, but users are required to provide a patch layout for the common base domain and feature points identificaton on each object.

The correspondence problem in the mesh morphing is naturally related to parametrizations. The survey of 3D morphing can be found in [Lazarus and Verroust 1998], and an extensive reviews of mesh morphing can be found in [Alexa 2001]. Particularly, Alexa pointed out in [Alexa 2001] that the remeshing approach is appealing for morphing applications, because it allows the size of the representation mesh to be scaled. On the contrary, the conventional merging approach generates a more complicated intermediate representation. Kanai et al. [Kanai et al. 1998] used a single patch and the patch will be parametrized by harmonic mapping. In their recent work [Kanai et al. 2000], the user first defines a set of corresponding feature vertices and they then applied their approximate shortest path algorithm [Kanai and Suzuki 2001] to find an initial dissection. The works of Gregory et al. [Gregory et al. 1998] and Zöckler et al. [Zöckler et al. 2000] also require users to manually provide initial dissections. Among the methods proposed, Lee et al. utilized their MAPS to the mesh morphing application [Lee et al. 1999]. But the base domains of the two input meshes are different—they are not consistently parametrized. They need a *"meta-mesh"* as an intermediate representation. Unfortunately this algorithm does not scale well, because the meta-mesh is more complicated than the original two models. Michikawa et al. [Michikawa et al. 2001] proposed *multiresolution interpolation meshes(MIMesh)* representation for mesh morphing. An interface is designed for users to define a common patch layout on both source and target meshes. Each patch is then parametrized and a surface fitting is performed to produce a semi-regular mesh for the intermediate mesh. The method can be extended to multi-target morphing.

3 Consistent Mesh Dissection and Parameterization

3.1 Floater's patch parameterization

The patch parameterization builds a bijective mapping between a region $\mathcal{R}$ of the original input 3D mesh and a region $\mathcal{B}$ of the 2D parameter plane, i.e., find a bijective map $u : \mathcal{R} \rightarrow \mathcal{B}$. The map u is fixed by a boundary condition $\partial\mathcal{R} \rightarrow \partial\mathcal{B}$. For the parameterization proposed by Floater [Floater 1997], we first fix the parameters of quadrilateral patch boundary vertices on a unit square, and then the function u needs to satisfy the following equation for each of the interior vertices v_i:

$$u(v_i) = \sum_{k \in \mathcal{V}(i)} \lambda_{i,k} u(v_k), \tag{1}$$

$$\sum_{k \in \mathcal{V}(i)} \lambda_{i,k} = 1, \qquad i = 1, 2, ..., n \tag{2}$$

where $\mathcal{V}(i)$ is the 1-ring neighborhood of the vertex v_i, and n is the number of the vertices of the patch. To derive the mapping u, we first compute for each i the $\lambda_{i,k}$, $k \in \mathcal{V}(i)$, that satisfies shape-preserving criteria and $\sum_{k \in \mathcal{V}(i)} \lambda_{i,k} = 1$, then solve the linear system of Eq. 1 for $u(v_k)$. The weight finding procedure of $\lambda_{i,k}$, for each i, consists of two stages. The first one is to compute a local parameterization for v_i, followed by computing the $\lambda_{i,k}$ from the local parameterization.

The main advantage of Floater's parameterization is that the $\lambda_{i,k}$ are always positive, because they are all derived from barycentric coordinates. This guarantees the mapping is bijective and a sparse linear system derived from Eq. 1 can be solved robustly by the iterative biconjugate gradient method.

3.2 Quinary Patch Subdivision

Highly stretched parameterization is not desirable and will affect the subsequent remeshing process significantly. We employ TMPM's L^2 stretch metric [Sander et al. 2001] to evaluate the patch's stretch ratio. If the stretch ratio of a patch P, say $L^2(P)$, exceeds a pre-defined threshold τ, it will be further subdivided into five patches by our quinary subdivision. Fig. 1 illustrates the quinary subdivision. A

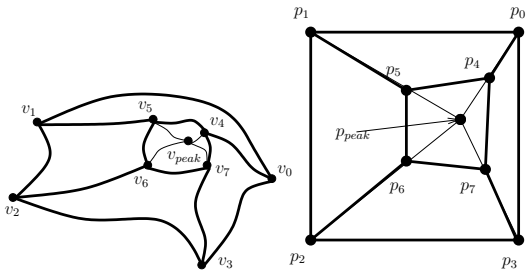

(a) 3D domain (b) 2D parameter domain

Figure 1: Quinary subdivision illustration

patch with four corners, v_0, v_1, v_2, and v_3, in R^3 is denoted as $P_{\{0,1,2,3\}}$. Likewise, a patch is parameterized on a unit square with four corners, p_0, p_1, p_2, and p_3, where $p_0 = u(v_0)$, $p_1 = u(v_1)$, $p_2 = u(v_2)$, and $p_3 = u(v_3)$, in R^2 is denoted as $S_{\{0,1,2,3\}}$. To apply our quinary subdivision to a patch $S_{\{0,1,2,3\}}$, we first find the greatest stretched face t_{peak} in R^2, which is T_{peak} in R^3, then simply take the centroid p_{peak} of t_{peak} as the peak point. For each corner p_i, $i =$

$0, 1, 2, 3$, of the patch, we find an intermediate point p_{i+4} on the line segment connecting p_{peak} and p_i as follows:

$$p_{i+4} = (1 - \omega_i)p_i + \omega_i p_{peak}, \quad i = \{0, 1, 2, 3\}, \quad 0 \leq \omega_i \leq 1, \tag{3}$$

for some w_0, w_1, w_2, and w_3. Because of the exponential uptrend of the face stretch ratio, we simply choose a constant ω for ω_i, $i = 0, 1, 2, 3$, as

$$\omega = \psi + \log_{10}(L^2(t_{peak})), \tag{4}$$

where ψ is a bias and is taken as 0.85 in our current implementation. Now we have the four new corners, and the patch is subdivided into five new patches $S_{\{0,1,5,4\}}$, $S_{\{1,2,6,5\}}$, $S_{\{2,3,7,6\}}$, $S_{\{3,0,4,7\}}$, and $S_{\{4,5,6,7\}}$. The peak point finding and linear combination process ensure the subdivision be uniquely derived. As in [Guskov et al. 2000], a straight line, which may go through faces, in the parameter domain will be a fair curve in 3D domain. By using the inverse mapping from 2D to 3D, the boundary curves of the newly created patches can be determined. We then employ a simple algorithm to marshal faces to their respected new patches. Pick a seed face by inverse mapping the 2D point, which is the average position of four corners of the new patches in the parameter domain, and recursively collect all its neighboring faces until the boundary curves are met. Finally the original patch can be removed. Fig. 2 shows the process of quinary subdivision on a cat head model.

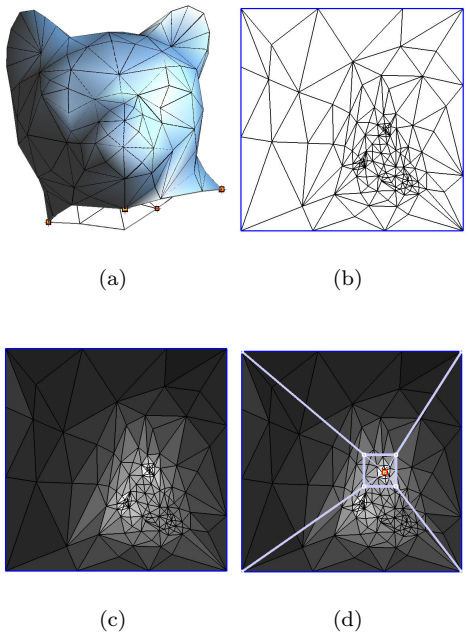

Figure 2: (a) A cat head model. (b) The parameterization. (c) The face stretch ratio is visualized as face color. (d) The quinary subdivision.

3.3 Dissection on a Single Object

It's obvious that the quinary subdivision has recursive structure. If the newly created patches are still not desirable, they can be subdivided recursively. In order to perform quinary subdivision, we have to first divide a genus-zero model into two patches. This step requires users to specify four "seed points" on the surface. When four seed points are correctly given, we perform Dijkstra's shortest path algorithm [Cormen et al. 1990] based on geodesic distance to trace from one seed point to another seed point to form a closed loop. Now with this closed curve, the original model is dissected into two "seed patches" and the respected faces are well marshaled; as shown in Fig. 3(a-b).

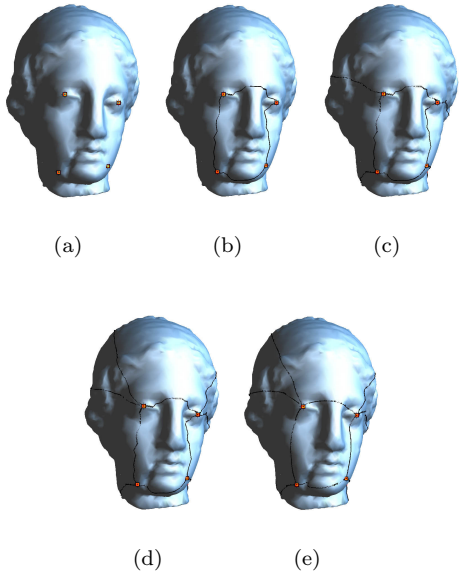

Figure 3: (a) Four seed points specified. (b) Two seed patches created. (c) Subdivision on the larger patch. (d) Result of recursive quinary subdivision. (e) The relaxed version.

Now we have two patches in hand, and we put them into a list. For each patch in the list, we first check their stretch ratios. If there is a patch exceeding a user-specified threshold, the patch is removed from the list for further subdivision. The resulting newly created patches are then put into the list. The process is repeated until all patches in the list satisfy the stretch ratio test, and the dissection of a single object is found; as shown in Fig. 3(c-d). We can also perform patch boundary relaxation and global vertex relaxation (similar to [Guskov et al. 2000]) to yield a better dissection. Fig. 3(e) shows the relaxed dissection of Fig. 3(a). In practice, only initially closed curves traced out by the shortest path algorithm necessarily require relaxation. All other curves resulting from the quinary subdivision are inverse mapped from the line segments on the parameter domain, so they are often fair enough and require no further relaxation.

3.4 Common Dissection for Multiple Objects

The recursive quinary subdivision for the seed patches can be recorded as a "quinary tree". In order to yield a common dissection for multiple objects, the recursive subdivision processes must be the same. That is, their quinary trees must be the same. To this end, we can simply take the union of these quinary trees, denoted as Q_{union}. In this extent, the initial seed points act as feature corresponding points. Because of

the initial seed points, there will be one unique union of the quinary trees. For an object and its respected quinary tree Q_i, we examine the difference between Q_i and Q_{union}, and further deliberately subdivide the patch if there is a newly added node or subtree. This process can be performed at the end or iteratively in the course of quinary subdivision. For each iteration, subdivide corresponding patches on other objects whenever a patch of an object needs to be subdivided. Repeat the process until all patches of all objects satisfy the stretch ratio test.

3.4.1 Extra Feature Correspondences

The derivation of base domain based on the quinary subdivision can be fully automatic with only four initial seed points specified by the user. If the users require extra feature corresponding points to be specified on the models for better performance on applications such as mesh metamorphosis, the quinary subdivision rule and parameterization can be enhanced to take into account the alignment of those additional feature points. Insufficient feature points may result in an unexpected metamorphosis sequence in such applications. Without proper alignment, the specified feature points in correspondence, however, could belong to different dissected patches, or belong to the same patch but possess different parametric values. Such cases often occur when feature points are specified too close together or the input models are too dissimilar, Fig. 4 illustrates such situations. To prevent from this situation, a *foldover-free warping* in the parameter domain should be performed before testing for each quinary subdivision.

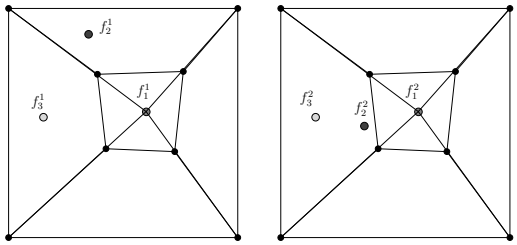

Figure 4: Feature points f_2^1 and f_2^2 are marshaled into different patches during quinary subdivision.

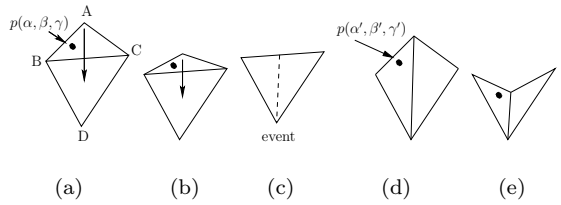

Figure 5: Triangulation over time. (a) Vertex A tends to move down. (b) Check if there is an event occuring during the movement. (c) An event, identifying the degeneration of triangle ABD, is detected. (d) The local triangulation is altered to prevent from the degeneration, and the parameters are recomputed by new barycentric coordinates. (e) Warped result.

We take the approach proposed in [Fujimura and Makarov 1998]. We briefly describe the main idea of the foldover-free warping algorithm. Given patches S^i, $i = 1, ..., n$, in R^2 and the associated set of feature points $\{f_1^i, ..., f_e^i\}$, where n is the number of input meshes and e is the number of feature points in S^i, we first compute the averaged feature point position as

$$\overline{f_k} = \frac{1}{n}\sum_{i=1}^{n} f_k^i, \qquad k = 1, ..., e.$$

For each S^i, we first construct a *warp mesh* for it by a 2D Delaunay triangulation which takes four corners, and f_1^i,..., and f_e^i as input. All other points will be marshaled into their respective enclosing triangles and their barycentric coordinates are also computed. The objective of a warping in parameter domain of S^i is to move feature point f_k^i to $\overline{f_k}$ for all i and k and recompute all other parameters which will still keep it as a bijective mapping. Therefore, the algorithm is just performed to each individual patch, not to all patches simultaneously. During the deformation of the warp mesh by the movement of the feature points, some triangles of the warp mesh may degenerate and start folding over. We call such situation an *event*. In order to prevent triangles from folding over, an algorithm called *triangulation over time* will be performed. Before starting the deformation, we detect if there will be an event by employing a binary search between the source position and the destination position of an arbitrary feature point f_t. If an intermediate position of the event is found, we alter the local triangulation and recompute all barycentric coordinates affected by this alteration. Then we move f_t to the position of the event, which was detected but will not occur this time. Then all parameters are recomputed by using the new barycentric coordinates. The influence range of the warping is not global and only parameters marshaled in this range will be affected. Furthermore, if there is more than one feature point, the triangulation over time algorithm will process them one by one in an arbitrary order. Although the processing order will affect the final distribution of the non-feature points, what we wish is to keep the mapping bijective. After processing one feature point, we set the position as the source position and repeat the event detection for other feature points. The process terminates when all feature points reach their destination positions.

If no event is found, simply deform the warp mesh to destination and recompute all parameters. Because the barycentric coordinates make the parameters inside a triangle bijective and the triangulation of the warp mesh is also bijective, we can be sure that the warped parameters are still bijective. Fig. 5 illustrates the triangulation over time. Fig. 6 shows an example of foldover-free warping in the parameter domain. An example is shown in Fig. 7, in which, besides four limbs, more feature points are specified for mouths, tail of the triceratops and hip of the horse, horn of the triceratops and ear of the horse. Fig. 8 shows the result of QS applied to the models of Venus, Venus head, and Isis.

4 Remeshing

4.1 Uniform Remeshing

With the parametrization ready for each patch of the base domain, we can start the remeshing process. For each patch, take the regular grid point parameters and compute their inverse-mapped 3D position. For an arbitrary parameter (vertices in parameter domain), we have to determine which

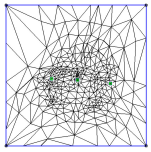

(a) Patch A.

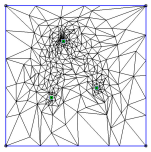

(b) Patch B.

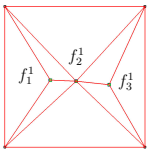

(c) Warp mesh of patch A.

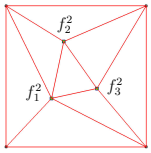

(d) Warp mesh of patch B.

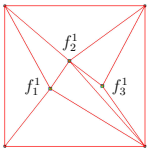

(e) Triangulation is altered during deformation.

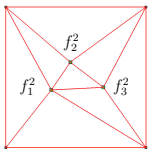

(f) After deformation.

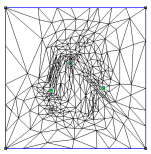

(g) Warped parameters.

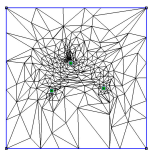

(h) Warped parameters.

Figure 6: Foldover-free warping in the parameter domain.

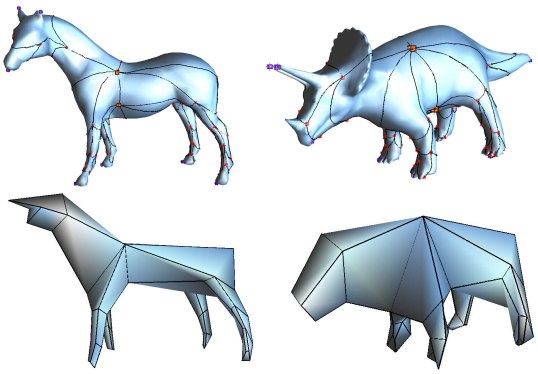

Figure 7: Common dissection and base domains of a horse model and a triceratops model.

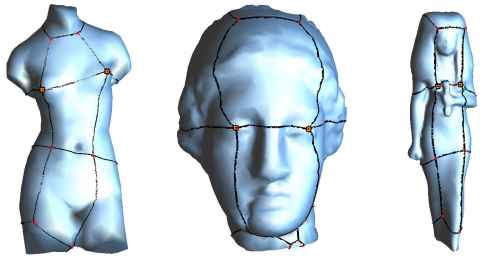

Figure 8: Common dissection of Venus, Venus head, and Isis models.

face contains it, which is a point location problem. Standard computational geometry algorithms can be employed. In our current implementation, we first perform a spatial partitioning technique on 2D parameter domain and the 2D quadtree for which all faces in the parameter domain will be marshaled into quadtree leaf node will be constructed. For a point location query, traverse the quadtree to find its corresponding leaf node, and each faces belonging to this leaf node will be tested. We restrict the face capacity of a leaf node, and the time complexity of a point-location query will be bounded to the depth of the quadtree.

After the face is found, simply use the barycentric coordinates to compute its 3D position. If you want to resample to level n, the total amount of the samples is $(2^n + 1)^2 m$, where m is the number of patches. All required geometry information can be resampled, and stored into a 2D array or a quadtree data structure depending on the applications. No matter what 2D array or tree structure is used, only geometry information needs to be stored by arranging the data in a special order. The resulting file size will be small. As for multiple objects, the final remeshes will have the same connectivity structures. Fig. 9 shows the result of uniform remeshing derived from the dissection in Fig. 7.

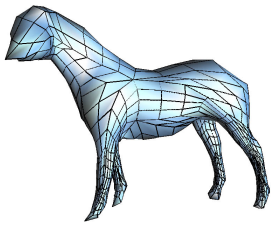

(a) Horse: M^2

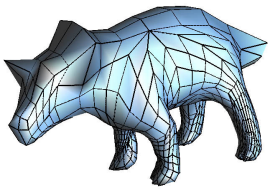

(b) Triceratops: M^2

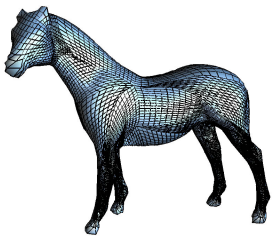

(c) Horse: M^4

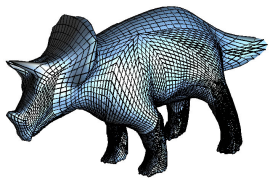

(d) Triceratops: M^4

Figure 9: The uniform remeshing.

4.2 Adaptive Remeshing

Uniform remeshing has the drawback that in order to resolve a small local feature on the original mesh, one may need to subdivide to a very fine level. This total number of faces will be quadrupled. We now describe a simple method to build the adaptive remesh within a conservative error bound.

For a given input mesh $\mathcal{M}$, a base domain consisting of some quad-faces corresponding to patches are constructed. For a given quad-face q, we find a best-fitting plane g, and measure the minimum Euclidian distance between the plane g and each vertex in the patch P_q associated with quad-face q. Let $d(v)$ be the minimum Eucilid distance for each $v \in P_q$ and each $p \in g$.

$$d(v) = \min_{p \in g} |v - p|$$

We define the error function $E(q)$ for each quad-face q as the maximum of $d(v)$ for all $v \in P_q$; i.e.,

$$E(q) = \max_{v \in P_q} d(v) \tag{5}$$

Fig. 10 illustrates the error function.

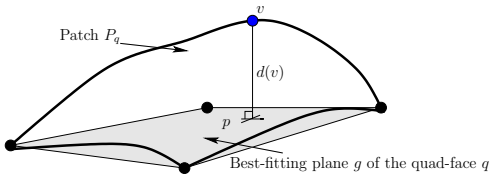

Figure 10: Error function definition.

The error function can be normalized by the diagonal length of the bounding box of the input mesh, denoted as $B(\mathcal{M})$; i.e.,

$$E_N(q) = \frac{E(q)}{B(\mathcal{M})} \tag{6}$$

We begin the remeshing process with base domain mesh and construct a quadtree root for each quad-face of base domain mesh. Then we evaluate the error of quad-faces in each quadtree based on the error function $E_N(q)$. If the error of a quad-face q, $E_N(q)$, is exceeding a pre-defined error bound ϵ, the quad-face is further refined and its geometry information is resampled. In other words, we take the quad-face to the next finer level and four new children will be attached into the quadtree. The process is performed recursively and the adaptive remesh is constructed. However, there will be lots of *T-vertices*, which appear along the boundaries between quad-faces of different levels. We first force the level difference between neighboring quad-faces to be at most one by deliberately refining the quad-face of the coarser level. We then perform adaptive subdivision to quad-faces of the coarser level.

To perform adaptive remeshing consistently on multiple objects, we can simply take the maximum of the error of each corresponding quad-faces in the multiple objects. The result will still be consistent.

$$E_C = \max_{1 \le i \le n} (E_N(q_i)) \tag{7}$$

where n is the number of objects. This will also result in adaptive remeshing with the same connectivity structure. If there are some local geometric features in one object, which correspond to a flat region in another object, the flat region will be forced to be refined and still result in lots of faces. Fig. 11 shows the result of the consistently adaptive remeshing.

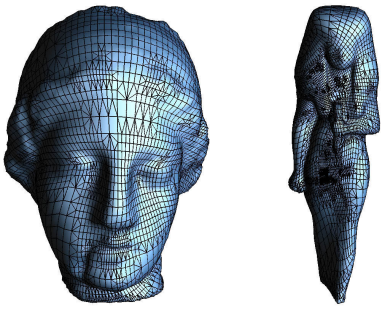

Figure 11: Adaptive remeshing of Venus and Isis models ($\epsilon = 0.0025$).

4.3 Experimental Results

We have implemented quinary subdivision and remeshing as described above. The applications was written in C++ using standard computational geometry data structures (for example, half-edge data structures). The operating system is Microsoft Windows 2000. The graphics rendering is based on OpenGL with nVidia extensions. The remeshing results are evaluated by IRI-CNR Metro tool [Cignoni et al. 1998]. We show the percentage of mean square error (L^2) normalized by the diagonal length of the bounding box of the models. Roughly speaking, the L^2 error below 0.05% is considered to be a nice approximation.

Table 1 shows the result of uniform and adaptive remeshing. Table 2 shows the approximate errors. Table 3 shows calculation time, which is performed on a 750MHz Pentium III machine.

5 Mesh Metamorphosis

The semi-regular meshes constructed from the remeshing process for multiple objects have the same connectivity. We can simply linearly interpolate their positions. A sequence of intermediate morphed objects will be generated, and they still have multiresolution structures. Fig. 12 shows a multiresolution metamorphosis using uniform remeshing. A multiresolution metamorphosis using adaptive remeshing is depicted in Fig. 13. Fig. 14 shows the morphing sequence from a pig model to a triceratops model. Based on the constructed semi-regular meshes obtained based on the common dissection shown in Fig. 8, we can produce any morphing sequence among these objects. Fig. 15 shows the result.

6 Concluding Remarks

The correspondence establishment among multiple objects is a versatile algorithm in computer graphics and geometry computing, especially in morphing applications. Other applications such as geometry processing can also benefit from the correspondences establishment. However, considerable effort is required by users suppress the establishment. A simpler and intuitive framework is necessary for alleviating the effors.

We have proposed a systematic method called recursive quinary subdivision to find a common dissection for multiple objects with only four feature points specified by the user. Extra feature points in correspondences can also be specified for semantics and aligned during the subdivision. Based on this dissection, uniform and adaptive remeshing can be

Model	Polygon number								
	original	uniform remeshing							adaptive remeshing ($\epsilon = 0.0025$)
		M^0	M^1	M^2	M^3	M^4	M^5	M^6	
Venus Isis	10000	14	56	224	896	3854	14336	57344	8656

Table 1: Statistics of polygon numbers.

Model	L^2 error(%)							
	uniform remeshing							adaptive remeshing ($\epsilon = 0.0025$)
	M^0	M^1	M^2	M^3	M^4	M^5	M^6	
Venus	7.97	2.46	0.83	0.26	0.086	0.033	0.012	0.047
Isis	2.64	1.32	0.66	0.19	0.073	0.026	0.009	0.037

Table 2: Statistics of errors.

Model	size	size of M^0	time(sec)	
			dissection	remeshing
Venus+Isis	10000+10000	14	11.967	1.522
horse+human	5000+5000	106	6.169	4.156
horse+triceratops	2000+5660	70	3.836	2.904
bunny	69664	34	143.967	10.145

Table 3: Statistics of calculation time. The remeshing is uniform and up to level 5.

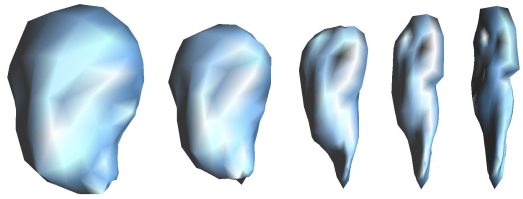

(a) M^2: 224 faces

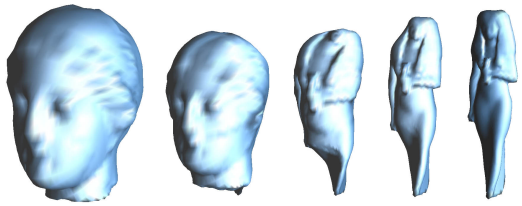

(b) M^4: 3584 faces

Figure 12: Multiresolution metamorphosis using uniform remeshing.

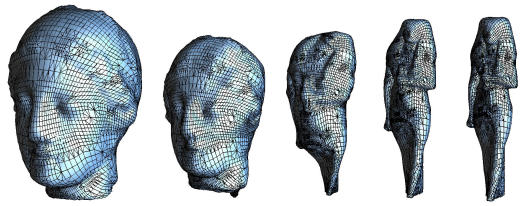

Figure 13: Metamorphosis by adaptive remeshing ($\epsilon = 0.0025$: 8656 faces).

performed to yeild a set of semi-regular meshes. Moveover, geometric details can easily be resampled and stored as normal maps to improve the visual effects using the modern graphics hardware. We have demonstrated the 3D mesh morphing application between two or more objects using the correspondence established by the common dissection and remeshing.

References

Alexa, M. 2001. Mesh morphing. In *EUROGRAPHICS'01 State of The Art Report*.

Cignoni, P., Rocchini, C., and Scopigno, R. 1998. Metro: Measuring error on simplified surfaces. *Computer Graphics Forum 17*, 2, 167–174.

Cormen, T. H., Leiserson, C. E., and Rivest, R. L. 1990. *Introduction to Algorithms*. The MIT Press.

Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., and Stuetzle, W. 1995. Multiresolution analysis of arbitrary meshes. In *Proceedings of SIGGRAPH 95*, 173–182.

Floater, M. S. 1997. Parametrization and smooth approximation of surface triangulation. *Computer Aided Geometric Design 14*, 3, 231–250.

Fujimura, K., and Makarov, M. 1998. Foldover-free image warping. *Graphical models and image processing 60*, 2 (March), 100–111.

Gregory, A., State, A., Lin, M., Manocha, D., and Livingston, M. 1998. Feature-based surface decomposition for correspondence and morphing between polyhedra. In *Computer Animation'98*, IEEE Computer Society, 64–71.

Guskov, I., Vidimče, K., Sweldens, W., and Schröder, P. 2000. Normal meshes. In *Proceedings of SIGGRAPH 2000*, 95–102.

Kanai, T., and Suzuki, H. 2001. Approximate shortest path on polyhedral surface and its applications. *Computer-Aided Design 33*, 11 (September), 801–811.

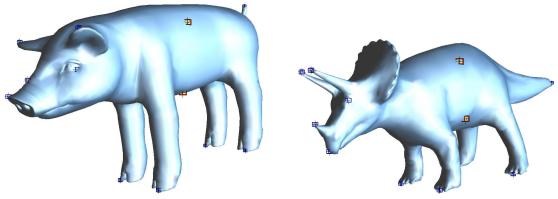

(a) User-specified feature points

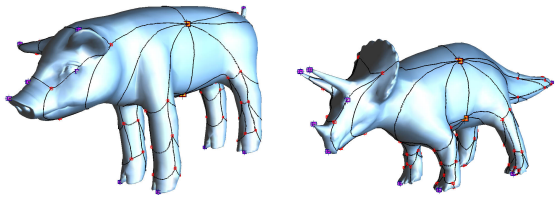

(b) Common dissection

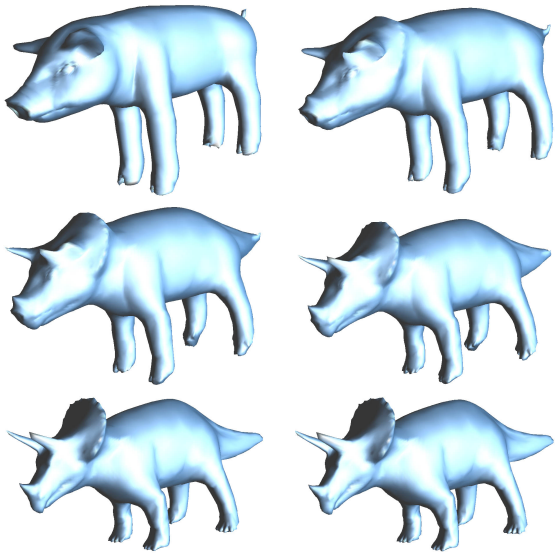

(c) Morphing sequence

Figure 14: Morphing from a pig model to a triceratops model.

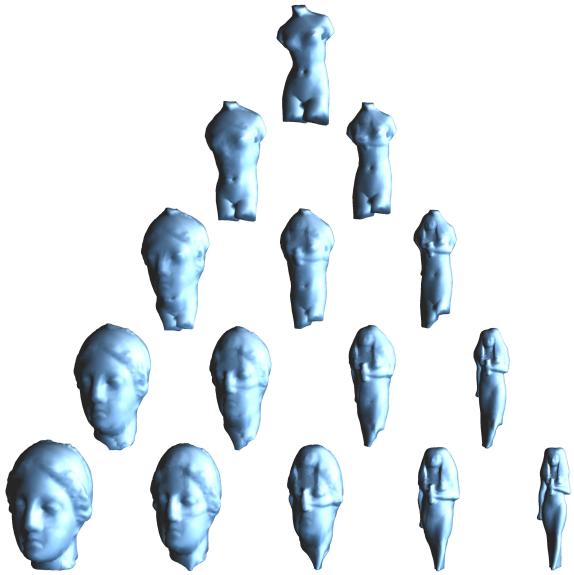

Figure 15: Multi-target metamorphosis (M^5: 14336 faces).

Kanai, T., Suzuki, H., and Kimura, F. 1998. Three-dimensional geometric metamorphosis based on harmonic maps. *The Visual Computer 14*, 4, 166–176.

Kanai, T., Suzuki, H., and Kimura, F. 2000. Metamorphosis of arbitrary triangular meshes. *IEEE Computer Graphics & Applications 20*, 2 (March/April), 62–75.

Khodakovsky, A., Schröder, P., and Sweldens, W. 2000. Progressive geometry compression. In *Proceedings of SIGGRAPH 2000*, 271–278.

Lazarus, F., and Verroust, A. 1998. Three-dimensional metamorphosis: a survey. *The Visual Computer 14*, 4, 373–389.

Lee, A., Sweldens, W., Schröder, P., Cowsar, L., and Dobkin, D. 1998. MAPS: Multiresolution adaptive parameterization of surfaces. In *Proceedings of SIGGRAPH 98*, 95–104.

Lee, A., Dobkin, D., Sweldens, W., and Schröder, P. 1999. Multiresolution mesh morphing. In *Proceedings of SIGGRAPH 99*, 343–350.

Lounsbery, M., DeRose, T., and Warren, J. 1997. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics 16*, 1, 34–73.

Lounsbery, M. 1994. *Multiresolution Analysis for Surfaces of Arbitrary Topological Type*. PhD thesis, Department of Computer Science and Engineering, University of Washington.

Michikawa, T., Kanai, T., Fujita, M., and Chiyokura, H. 2001. Multiresolution interpolation meshes. In *Proceedings of Pacific Graphics'01*, IEEE Computer Society Press, 60–69.

Praun, E., Sweldens, W., and Schröder, P. 2001. Consistent mesh parameterizations. In *Proceedings of SIGGRAPH 2001*, 179–184.

Sander, P. V., Snyder, J., Gortler, S. J., and Hoppe, H. 2001. Texture mapping progressive meshes. In *Proceedings of SIGGRAPH 2001*, 409–416.

Zöckler, M., Stalling, D., and Hege, H. C. 2000. Fast and intuitive generation of geometric shape transitions. *The Visual Computer 16*, 5, 241–253.