

MULTIRESOLUTION MODELING AND INTERACTIVE DEFORMATION OF LARGE 3D MESHES

Jens Vorsatz and Hans-Peter Seidel

Max-Planck-Institut für Informatik

Keywords: Multiresolution, Modeling, Mesh, Deformation

Abstract Due to their simplicity triangle meshes are often used to represent geometric surfaces. Their main drawback is the large number of triangles that are required to represent a smooth surface. This problem has been addressed by a large number of mesh simplification algorithms which reduce the number of triangles and approximate the initial mesh. Hierarchical triangle mesh representations provide access to a triangle mesh at a desired resolution, without omitting any information.

In this paper we demonstrate how a hierarchical structure of a mesh can be derived for arbitrary meshes to enable intuitive and efficient modifications without restrictions on the underlying connectivity. We combine mesh reduction algorithms and constrained energy minimization to decompose the given mesh into several frequency bands and focus on a stabilizing technique to encode the geometric difference between the levels.

Introduction

Modification of complex 3D geometric shapes is a challenging task required for a wide variety of applications, for instance animation and design. Usually, this is done by editing a freeform surface, which represents the outer skin of a solid object. Like their real world equivalent, the surfaces often carry detail information on various scales such as skin wrinkles in the context of animation or e.g. the company sign on an engine hood. Certainly it is desirable to preserve these features while editing the global shape of the surface. During the last years, hierarchical multiresolution representations of geometric shape have become the de facto standard for those purposes.

The basic idea is to separate the high frequency detail from the low frequency shape and encode each detail level relative to a local coordinate

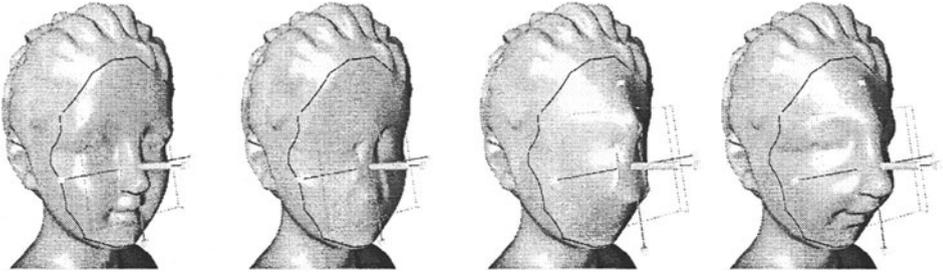


Figure 1 Multiresolution edit: In the area of interest (blue line), the original mesh (outer left) is decomposed into two frequency bands. The geometric difference between the high frequency detail (outer left) and the low frequency shape (center left) is stored with respect to local frames. A designer changes the low frequency shape by shifting the green polygon (center right). Adding the previously stored detail information yields the final result (outer right).

frame induced by the lower frequency shape. That way modifications on a coarser level can be propagated to the finer ones (cf. Fig. 1). More formally, given an arbitrary surface \mathcal{S}_m , a multiresolution *decomposition* consists of a sequence of topologically equivalent surfaces $\mathcal{S}_{m-1}, \dots, \mathcal{S}_0$ with decreasing level of geometric detail. The difference $\mathcal{D}_i = \mathcal{S}_{i+1} - \mathcal{S}_i$ between two successive surfaces is the *detail* on level i which is added or removed when switching between the two approximations. The reconstruction $\mathcal{S}_m = \mathcal{S}_i + \mathcal{D}_i + \dots + \mathcal{D}_{m-1}$ of the original surface \mathcal{S}_m can start on any level of detail \mathcal{S}_i . Multiresolution modeling means that on some level of detail, the surface \mathcal{S}_i is replaced by \mathcal{S}'_i . This operation does not have any effect on $\mathcal{S}_0, \dots, \mathcal{S}_{i-1}$ but \mathcal{D}_{i-1} and hence $\mathcal{S}_{i+1}, \dots, \mathcal{S}_m$ change since the (unchanged) detail information $\mathcal{D}_i, \dots, \mathcal{D}_{m-1}$ is now added to the modified base surface \mathcal{S}'_i for the reconstruction of \mathcal{S}'_m . In order to guarantee the intuitive preservation of the shape characteristics after a modification on some lower level of detail, this basic setting has to be extended such that the detail information \mathcal{D}_i is encoded with respect to *local frames*. These frames are aligned to the surface geometry of \mathcal{S}_i .

Fundamental work in this area, based on splines and wavelets, was done by Forsey and Barthels [Forsey and Bartels, 1988, Forsey and Bartels, 1995], Lee et al. [Lee et al., 1997, Lee, 1999], and Gortler and Cohen [Gortler and Cohen, 1995]. Though splines have a straight forward shape control mechanism based on control vertices, it is well-known to be rather complicated to preserve boundary conditions when handling complex geometry.

This is just one of the reasons, why the interest in surface representations in the context of multiresolution editing based on triangulation

lar meshes has increased over the last years. Generalizing the patch-based concepts, the wide family of subdivision techniques starts with a coarse base mesh approximating a geometric shape of arbitrary topology and refines it iteratively. An exponential number of vertices is introduced to capture finer detail information, until a prescribed tolerance is reached. This bottom-up approach generates the so-called subdivision-connectivity, which means, that sub-regions of the refined mesh which correspond to a single triangle in the base mesh have the connectivity of regular grids. To separate the high-frequency from the low-frequency, again, one defines decomposition and reconstruction operations. The reconstruction operator is given by the underlying subdivision scheme. To transform a mesh \mathcal{M}_m to the next refinement level $\mathcal{M}'_{m+1} = \mathbf{S}\mathcal{M}_m$ one applies the subdivision operator \mathbf{S} and moves the obtained control vertices by adding the associated detail vectors: $\mathcal{M}_{m+1} = \mathcal{M}'_{m+1} + \mathcal{D}_m$. One can think of the decomposition operator as the inverse of the subdivision operator, i.e. given a fine mesh \mathcal{M}_{m+1} one has to find a mesh \mathcal{M}_m such that $\mathcal{M}_{m+1} \approx \mathbf{S}\mathcal{M}_m$. This can be achieved by solving a minimization problem or more efficiently by discrete fairing [Zorin et al., 1997]. In this case the detail vectors $\mathcal{D}_m := \mathcal{M}_{m+1} - \mathbf{S}\mathcal{M}_m$ become as small as possible.

To achieve the desired multiresolution edit, one shifts a control vertex \mathbf{p}_i^m in \mathcal{M}_m . This has influence on several control vertices in the finer levels and causes a smooth bump in the resulting surface while maintaining the high-frequency detail information. The underlying low frequency geometry on each level can be computed by applying the reconstruction operator \mathbf{S} without detail reconstruction ($\mathcal{D}_m := 0$).

One problem which is inherent to multiresolution representations of freeform geometry based on subdivision surfaces is the fixed support of a modification. If control vertices are used as handles to modify the surface on a certain level of detail, the region that actually changes its shape is determined by the support of the associated basis functions. One could simulate more flexibility by moving several vertices at a time, but this annihilates the mathematical elegance of the representation.

Moreover, in practice it is rather unlikely that a mesh with subdivision-connectivity is given as input. For instance 3D acquisition of geometric data with a range scanning device followed by a triangulation or the conversion of CAD data, often result in a triangle mesh with arbitrary connectivity. For this reason sophisticated schemes have been presented to approximate an arbitrary input mesh with one having subdivision-connectivity [Eck et al., 1995, Lee et al., 1998, Kobbelt et al., 1999]. But besides being computationally expensive, the conversion is always a resampling process which gives rise to sampling artifacts. A popular

way to avoid the described problem is to build the hierarchical structure the other way around i.e from fine to coarse. For this, techniques which adapt the mesh-complexity to the available hardware resources emerging from another branch in computer graphics can be used. Multiple levels of resolution are produced by incrementally decimating the fine mesh [Garland and Heckbert, 1997, Hoppe, 1996, Kobbelt et al., 1998a, Lindstrom, 2000]. This is often done by applying a decomposition operator, that successively collapses edges and removes the redundant vertices and faces. To capture the detail information, which would be lost otherwise, again, detail vectors have to be stored. For a hierarchical representation, a proper reconstruction has to be ensured. Hence, we need a base point, where the detail vector could be attached to. In contrast to the subdivision scheme, where the base point is predicted by the subdivision operator, no such point exists for the coarse to fine approach, since the mesh-connectivity does not provide the necessary regular structure. For this reason, a vertex removal is split into two steps. First, the original position is altered such that local fairness is achieved. Only recently, a couple of new techniques have been proposed [Taubin, 1995, Kobbelt, 1997, Kobbelt et al., 1998b, Guskov et al., 1999, Schneider and Kobbelt, 2001]. The second step removes the original vertex and encodes the position with respect to its minimized counterpart.

This would require a fairing step for every single vertex. One could also apply the fairing operation to *all* vertices before storing the detail information to lower the computational costs. This would lead to a two-band representation, i.e. a smoothed version, and the original mesh linked by the detail vectors. In practice, a multi-band hierarchy, similar to a level of detail representation would be desirable. This could reflect the multiple scales of features on the surface to stabilize the modeling-process on the one hand and keep down the costs on the other hand.

Hence, to build an appropriate hierarchical structure of a triangular mesh for our modeling purposes, we have to solve two problems. First, we have to choose the right intermediate frequency-bands, such that a modification of a coarser level will lead to reasonable alternation of the finer ones. On the other hand, the detail has to be encoded with respect to a proper base point, to ensure a stable reconstruction. The following sections discuss several approaches for both problems.

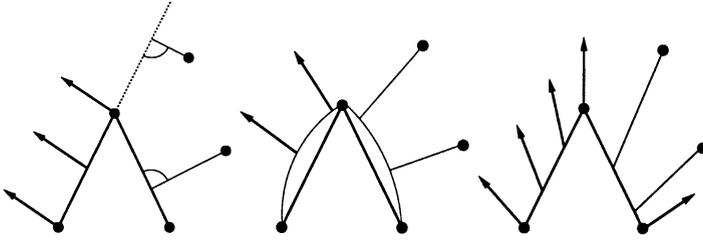


Figure 2 The position of a vertex (loose points) in the original mesh (high-frequency geometry) is given by a base point on the low-frequency geometry plus a displacement in normal direction. There are many ways to define a normal field on a triangle mesh. With piecewise constant normals (left) we do not cover the whole space and hence we sometimes have to use virtual base points with negative barycentric coordinates. The sketch shows, that this can lead to non intuitive reconstructions, if the 'base mesh' is for example flattened out. The use of local quadratic patches and their normal fields (center) somewhat improves the situation but problems still occur since the overall normal field is not globally continuous. Such difficulties are avoided if we generate a Phong-type normal field by blending estimated vertex normals (right).

1. DETAIL ENCODING

As mentioned before, we cannot simply store the detail vectors with respect to a global coordinate system but we have to define them with respect to local frames which are aligned to the low-frequency geometry. This guarantees an intuitive detail preservation under modification of the global shape. Usually, the associated local frame for each vertex has its origin at the location predicted by the smoothing operator (reconstruction operator with suppressed detail). However, in many cases this can lead to rather long detail vectors with a significant component within the local tangent plane. Since we prefer short detail vectors for stability reasons, it makes sense to use a different origin for the local frame. In fact, the optimal choice is to find that point on the low-frequency surface whose normal vector points directly to the original vertex. In this case, the detail is not given by a three dimensional vector $(\Delta x, \Delta y, \Delta z)^T$ but rather by a base point $\mathbf{p} = \mathbf{p}(u, v)$ on the low-frequency geometry plus a scalar value h for the displacement in normal direction. If a local parameterization of the surface is available, the base point \mathbf{p} can be specified by a two-dimensional parameter value (u, v) .

The general setting for detail computation is that we have given two meshes \mathcal{M}_{m+1} and \mathcal{M}'_{m+1} where \mathcal{M}_{m+1} is the original data while \mathcal{M}'_{m+1} is reconstructed from the low-frequency approximation \mathcal{M}_m with suppressed detail, i.e. for coarse-to-fine hierarchies, the mesh \mathcal{M}'_{m+1} is generated by applying a stationary subdivision scheme and for fine-to-coarse

hierarchies \mathcal{M}'_{m+1} is optimal with respect to some global bending energy functional. Encoding the geometric difference between both meshes requires to associate each vertex \mathbf{p} of \mathcal{M}_{m+1} with a corresponding base point \mathbf{q} on the continuous (piecewise linear) surface \mathcal{M}'_{m+1} such that the difference vector between the original point and the base point is parallel to the normal vector at the base point. An arbitrary point \mathbf{q} on \mathcal{M}'_{m+1} can be specified by a triangle index i and barycentric coordinates within the referred triangle.

To actually compute the detail coefficients, we have to define a normal field on the mesh \mathcal{M}'_{m+1} . The most simple way to do this is to use the normal vectors of the triangular faces for the definition of a piecewise constant normal field. This projection can be computed efficiently and works fine, if the resulting coefficient is short compared to the edges of the assigned triangle and if \mathcal{M}'_{m+1} is sufficiently smooth. But since the orthogonal prisms spanned by a triangle mesh do not completely cover the vicinity of the mesh, we have to accept negative barycentric coordinates for the base points if it does not lie within such a prism. This leads to non-intuitive detail reconstruction if the low-frequency geometry is modified (cf. Fig 2).

A technique used in [Kobbelt et al., 1998b] is based on the construction of a local quadratic interpolant \mathbf{F} to the low-frequency geometry. For a vertex $\mathbf{p} \in \mathcal{M}_{m+1}$ it is based on the closest triangle $\mathbf{T} \in \mathcal{M}'_{m+1}$ and its adjacent vertices, which can be found in linear time by a simple local search procedure, starting from \mathbf{p} 's corresponding vertex $\mathbf{p}' \in \mathcal{M}_{m+1}$. Since now a local parameterization is given, parameter values (u, v) defining the base point \mathbf{q} can be found by Newton-iteration. We start from the center of \mathbf{T} at $\mathbf{q}_0 = \mathbf{F}(\frac{1}{3}, \frac{1}{3})$, \mathbf{q}_{n+1} is defined by the projection of \mathbf{p} into the tangent plane of \mathbf{F} at \mathbf{q}_n . In terms of parameter values (u, v) , this leads to the simple update rule $(u_{n+1}, v_{n+1}) \leftarrow (u_n, v_n) + (\Delta u, \Delta v)$, where $(\Delta u, \Delta v)$ is the solution of the linear system

$$\begin{pmatrix} F_u^T F_u & F_u^T F_v \\ F_u^T F_v & F_v^T F_v \end{pmatrix} \begin{pmatrix} \Delta u \\ \Delta v \end{pmatrix} = \begin{pmatrix} F_u^T d \\ F_v^T d \end{pmatrix} \quad (1)$$

with detail vector $\mathbf{d} = \mathbf{p} - \mathbf{q}_n$, which is perpendicular (within a prescribed tolerance) to $\mathbf{F}(u_n, v_n)$ after a few steps. The absolute value of the displacement-coefficient h is set to $\|\mathbf{d}\|$ and has to be multiplied by -1 if $d^T(f_u(u_n, v_n) \times f_v(u_n, v_n)) < 0$. Although this reduces the number of pathological configurations with negative barycentric coordinates for the base point, we still observe artifact in the reconstructed high-frequency surface which are caused by the fact that the resulting global normal field of the combined local patches is not continuous (cf. Fig 2 middle).

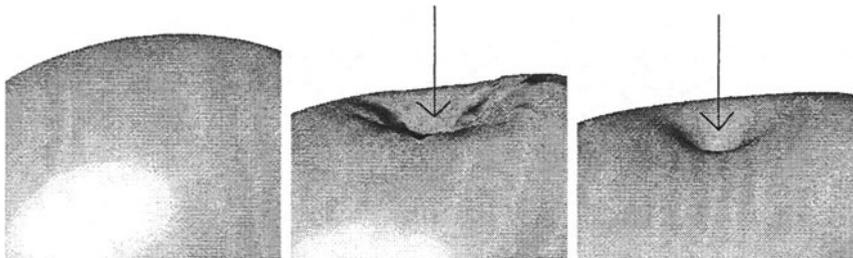


Figure 3 The original shape (left) is modified by pushing a single vertex while minimizing a membrane energy functional. A piecewise linear normal field leads to undesirable mesh artifacts (middle), while storing detail information with respect to a Phong normal field (left) performs a satisfying modification.

We therefore propose a different approach which adapts the basic idea of Phong-shading [Foley et al., 1990] where normal vectors are prescribed at the vertices of a triangle mesh and a continuous normal field for the interior of the triangular faces is computed by linearly blending the normal vectors at the corners. We use the same search procedure as described above and obtain a triangle $\Delta(\mathbf{a}, \mathbf{b}, \mathbf{c})$ with the associated normal vectors $N_{\mathbf{a}}$, $N_{\mathbf{b}}$, and $N_{\mathbf{c}}$. For each interior point

$$\mathbf{q} = \alpha \mathbf{a} + \beta \mathbf{b} + \gamma \mathbf{c}$$

with $\alpha + \beta + \gamma = 1$ we find the associated normal vector $N_{\mathbf{q}}$ by

$$N_{\mathbf{q}} = \alpha N_{\mathbf{a}} + \beta N_{\mathbf{b}} + \gamma N_{\mathbf{c}}.$$

When computing the detail coefficients for a given point \mathbf{p} we have to find the base point \mathbf{q} such that

$$(\mathbf{p} - \mathbf{q}) \times N_{\mathbf{q}}$$

has all three coordinates vanishing. By plugging in the definition of \mathbf{q} and $N_{\mathbf{q}}$ and eliminating $\gamma = 1 - \alpha - \beta$ we obtain a bivariate quadratic function

$$F : (u, v) \rightarrow \mathbf{R}^3$$

and we have to find the parameter value (α, β) such that $F(\alpha, \beta) = (0, 0, 0)^T$. Again, this can be accomplished by performing several steps of Newton-iteration. Notice that F can be interpreted as a quadratic surface patch in \mathbf{R}^3 which passes through the origin. The Taylor-coefficients

of F can explicitly be given by

$$\begin{aligned}
 F(0,0) &= W + WW \\
 F_u(0,0) &= U + UW - W - 2WW \\
 F_v(0,0) &= V + VW - W - 2WW \\
 F_{uu}(0,0) &= UU - UW + WW \\
 F_{uv}(0,0) &= UV - UW - VW + 2WW \\
 F_{vv}(0,0) &= VV - VW + WW
 \end{aligned}$$

where

$$\begin{aligned}
 U &= \mathbf{p} \times N_{\mathbf{a}} \\
 V &= \mathbf{p} \times N_{\mathbf{b}} \\
 W &= \mathbf{p} \times N_{\mathbf{c}} \\
 UU &= N_{\mathbf{a}} \times \mathbf{a} \\
 VV &= N_{\mathbf{b}} \times \mathbf{b} \\
 WW &= N_{\mathbf{c}} \times \mathbf{c} \\
 UV &= (N_{\mathbf{b}} \times \mathbf{a}) + (N_{\mathbf{a}} \times \mathbf{b}) \\
 UW &= (N_{\mathbf{c}} \times \mathbf{a}) + (N_{\mathbf{a}} \times \mathbf{c}) \\
 VW &= (N_{\mathbf{c}} \times \mathbf{b}) + (N_{\mathbf{b}} \times \mathbf{c})
 \end{aligned}$$

This leads to a similar update rule as described in 1. Starting with $(\alpha_0, \beta_0) = (\frac{1}{3}, \frac{1}{3})$, the difference $(\Delta\alpha, \Delta\beta)$ between two consecutive steps can be denoted as follows.

$$\begin{aligned}
 \Delta\alpha &= (F_u^T F_v \cdot F_v^T F - F_v^T F_v \cdot F_u^T F) / s \\
 \Delta\beta &= (F_u^T F_v \cdot F_u^T F - F_u^T F_u \cdot F_v^T F) / s
 \end{aligned}$$

with $s = F_u^T F_u \cdot F_v^T F_v - (F_u F_v)^2$.

In case one of the barycentric coordinates of the resulting point \mathbf{q} is negative, we continue the search for a base point in the corresponding neighboring triangle. Since the Phong normal field is globally continuous we always find a base point with positive barycentric coordinates. Fig. 2 depicts the situation schematically and Fig. 3 shows an example edit where the piecewise constant normal field causes mesh artifacts which do not occur if the Phong normal field is used.

2. HIERARCHY LEVELS

For coarse-to-fine hierarchies the levels of detail are determined by the uniform refinement operator. Starting with the base mesh \mathcal{M}_0 , the m th refinement level is reached after applying the refinement operator m times. For fine-to-coarse hierarchies there is no such canonical choice for the levels of resolution. Hence we have to figure out some heuristics to define such levels.

In [Kobbelt et al., 1998b] a simple two-band decomposition has been proposed for the modeling, i.e. the high frequency geometry is given

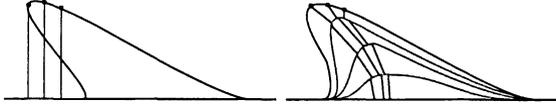


Figure 4 If the high-frequency detail cannot be projected onto the successive level (left), intermediate levels have to be inserted to guarantee a feasible detail reconstruction (right).

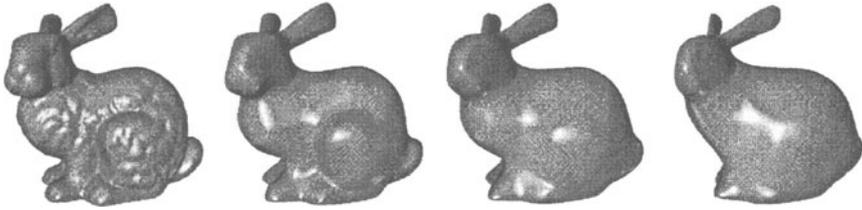


Figure 5 Four versions of the Stanford bunny. The smoother versions are generated by applying mesh decimation down to a certain target complexity and re-inserting the vertices under minimization of some discrete fairness functional. The degree by which geometric detail is removed depends on the coarseness of the base mesh. Notice that all shown meshes share the same connectivity.

by the original mesh and the low-frequency geometry is the solution of some constrained optimization problem. This simple decomposition performs well if the original geometry can be projected onto the low-frequency geometry without self-intersections. Fig 4 schematically shows a configuration where this is not satisfied and consequently the detail feature does not deform intuitively with the change of the global shape.

This effect can be avoided by introducing several intermediate levels of detail, i.e., by using a true multi-band decomposition. The definition of the Phong-type normal field introduced in the last section provides the means to guarantee a stable reconstruction. The number of hierarchy levels has to be chosen such that the $(i + 1)$ st level can be projected onto level i without self-intersection. Detail information has to be computed for every intermediate level.

Intermediate levels can be generated by the following algorithm (see Fig. 5). We start with the original mesh and apply an incremental mesh decimation algorithm which performs a sequence of edge collapse operations. When a certain mesh complexity is reached, we perform the reverse sequence of vertex split operations which reconstructs the original mesh connectivity. The position of the re-inserted vertices is found

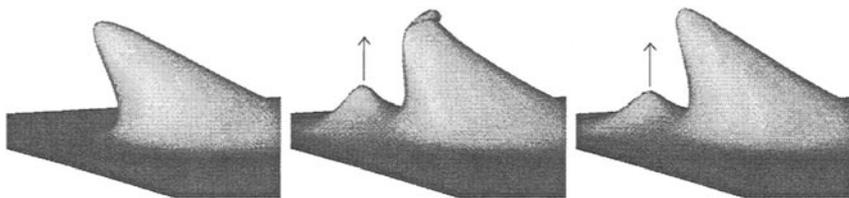


Figure 6 Starting from the original shape (left), a two-band decomposition (middle) can lead to long detail-vectors and hence to exaggerated modifications or even self-intersections for relatively small edits. Multiple levels of detail avoid these artifacts and the modifications behave in a natural fashion (right).

by solving a global bending energy minimization problem [Kobbelt, 1997, Kobbelt et al., 1998b, Guskov et al., 1999, Schneider and Kobbelt, 2001]. The mesh that results from this procedure is a smoothed version of the original mesh where the degree by which detail information has been removed depends on the target complexity of the decimation algorithm.

Suppose the original mesh has n_m vertices, where m is the number of intermediate levels that we want to generate. We can compute the meshes $\mathcal{M}_m, \dots, \mathcal{M}_0$ with fewer detail by applying the above procedure where the decimation algorithm stops at a target resolution of n_m, \dots, n_0 remaining vertices respectively. The resulting meshes yield a multi-band decomposition of the original data. When a modeling operation changes the shape of \mathcal{M}_0 we first reconstruct the next level \mathcal{M}'_1 by adding the stored detail vectors and then proceed by successively reconstructing \mathcal{M}'_{i+1} from \mathcal{M}'_i .

The remaining question is how to determine the numbers n_i . A simple way to do this is to build a geometric sequence with $n_{i+1}/n_i = \text{const}$. This mimics the exponential complexity growth of the coarse-to-fine hierarchies. Another approach is to stop the decimation every time a certain average edge length \bar{l}_i in the remaining mesh is reached.

A more complicated heuristic tries to equalize the sizes of the differences between levels, i.e., the sizes of the detail vectors. We first compute a multi-band decomposition with, say, 100 levels of detail where we choose $\sqrt[3]{\bar{n}_i} = \text{const.}$. For every pair of successive levels we can compute the average length of the detail vectors (displacement values). From this information we can easily choose appropriate values $n_j = \bar{n}_j$, such that the geometric difference is distributed evenly among the detail levels.

In practice it turned out that about five intermediate levels is usually enough to guarantee correct detail reconstruction. Fig. 6 compares the

results of a modeling operation based on a two-band and a multi-band decomposition.

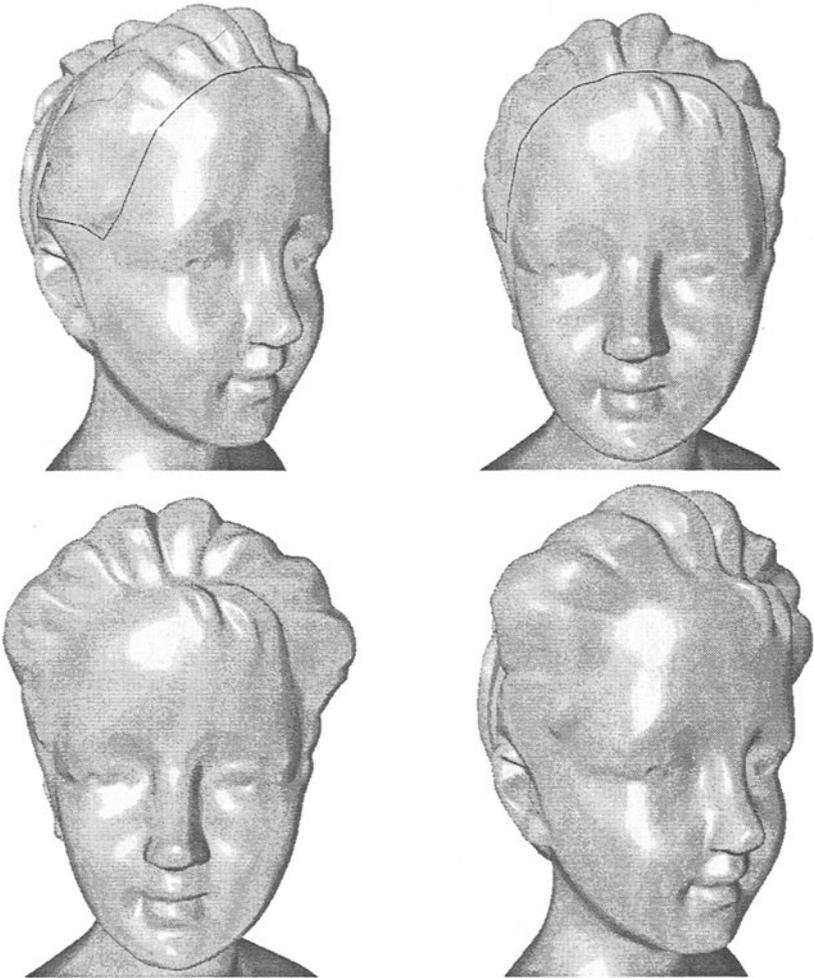


Figure 7 Multiresolution editing of a bust model. The area within the blue line is decomposed into two frequency-bands. The user changes the handle polygon (green) and thus changes the low-frequency surface on top of which the recorded detail based on a Phong-type normal field is reconstructed (lower row).

3. CONCLUSION

We have shown how one can derive a robust hierarchical structure of a triangle mesh with arbitrary connectivity. This enables efficient edits of a complex mesh in an intuitive manner. A designer can apply global deformations without losing detail information inherent to the surface. In particular, we have focused on a method to encode high-frequency detail with respect to a low-frequency base mesh. However, the user can still apply particular edits, where undesirable effects like self intersection of detail vectors during the reconstruction process happen. Moreover, due to the fixed mesh-connectivity, extreme stretches of triangles can occur. We are currently developing a system, which handles changes of the mesh during the modeling process, i.e. insertion of vertices, where the mesh is locally stretched and vertex removal, where the triangle size undergoes a given threshold. We are also keeping track of a promising approach to avoid self intersection without changing the mesh-connectivity.

References

- [Eck et al., 1995] Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., and Stuetzle, W. (1995). Multiresolution Analysis of Arbitrary Meshes. In *Computer Graphics (SIGGRAPH 95 Proceedings)*, pages 173–182.
- [Foley et al., 1990] Foley, van Dam, Feiner, and Hughes (1990). *Computer Graphics*. Addison Wesley.
- [Forsey and Bartels, 1988] Forsey, D. R. and Bartels, R. H. (1988). Hierarchical B-spline refinement. In *Computer Graphics (SIGGRAPH 88 Proceedings)*, pages 205–212.
- [Forsey and Bartels, 1995] Forsey, D. R. and Bartels, R. H. (1995). Surface Fitting with Hierarchical Splines. *ACM Transactions on Graphics*, 14(2):134–161.
- [Garland and Heckbert, 1997] Garland, M. and Heckbert, P. S. (1997). Surface Simplification Using Quadric Error Metrics. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 209–218.
- [Gortler and Cohen, 1995] Gortler, S. J. and Cohen, M. F. (1995). Hierarchical and variational geometric modeling with wavelets. In *1995 Symposium on Interactive 3D Graphics*, pages 35–42.
- [Guskov et al., 1999] Guskov, I., Sweldens, W., and Schröder, P. (1999). Multiresolution signal processing for meshes. In *Computer Graphics (SIGGRAPH 99 Proceedings)*, pages 325–334.
- [Hoppe, 1996] Hoppe, H. (1996). Progressive Meshes. In *Computer Graphics (SIGGRAPH 96 Proceedings)*, pages 99–108.

- [Kobbelt, 1997] Kobbelt, L. (1997). Discrete fairing. In *Proceedings of the Seventh IMA Conference on the Mathematics of Surfaces*, pages 101–131.
- [Kobbelt et al., 1998a] Kobbelt, L., Campagna, S., and Seidel, H.-P. (1998a). A general framework for mesh decimation. In *Proceedings of the Graphics Interface Conference '98*, pages 43–50.
- [Kobbelt et al., 1998b] Kobbelt, L., Campagna, S., Vorsatz, J., and Seidel, H.-P. (1998b). Interactive multi-resolution modeling on arbitrary meshes. In *Computer Graphics (SIGGRAPH 98 Proceedings)*, pages 105–114.
- [Kobbelt et al., 1999] Kobbelt, L., Vorsatz, J., Labsik, U., and Seidel, H.-P. (1999). A shrink wrapping approach to remeshing polygonal surfaces. *Computer Graphics Forum, Proceedings of Eurographics '99*, 18(3):119–130.
- [Lee et al., 1998] Lee, A., Sweldens, W., Schröder, P., Cowsar, L., and Dobkin, D. (1998). Multiresolution adaptive parameterization of surfaces. In *Computer Graphics (SIGGRAPH 98 Proceedings)*, pages 95–104.
- [Lee, 1999] Lee, S. (1999). Interactive multiresolution editing of arbitrary meshes. *Computer Graphics Forum, Proceedings of Eurographics '99*, 18(3):73–82.
- [Lee et al., 1997] Lee, S., Wolberg, G., and Shin, S. Y. (1997). Scattered data interpolation with multilevel B-splines. *IEEE Transactions on Visualization and Computer Graphics*, 3(3):228–244.
- [Lindstrom, 2000] Lindstrom, P. (2000). Out-of-core simplification of large polygonal models. In *Computer Graphics (SIGGRAPH 00 Proceedings)*, pages 259–262.
- [Schneider and Kobbelt, 2001] Schneider, R. and Kobbelt, L. (2001). Geometric fairing of irregular meshes for free-form surface design. *CAGD (to appear)*.
- [Taubin, 1995] Taubin, G. (1995). A signal processing approach to fair surface design. In *Computer Graphics (SIGGRAPH 95 Proceedings)*, pages 351–358.
- [Zorin et al., 1997] Zorin, D., Schröder, P., and Sweldens, W. (1997). Interactive multiresolution mesh editing. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 259–268.