

# Robust Multi-Band Detail Encoding for Triangular Meshes of Arbitrary Connectivity

Jens Vorsatz, Leif P. Kobbelt

Max-Planck-Institut für Informatik, Computer Graphics Group  
Im Stadtwald, 66123 Saarbrücken, Germany  
Email: vorsatz@mpi-sb.mpg.de

## Abstract

The flexibility coming along with the simplicity of their base primitive and the support by today's graphics hardware, have made triangular meshes more and more popular for representing complex 3D objects. Due to the complexity of realistic datasets, a considerable amount of work has been spent during the last years to provide means for the modification of a given mesh by intuitive metaphors, i.e. large scale edits under preservation of the detail features. In this paper we demonstrate how a hierarchical structure of a mesh can be derived for arbitrary meshes to enable intuitive modifications without restrictions on the underlying connectivity, known from existing subdivision approaches. We combine mesh reduction algorithms and constrained energy minimization to decompose the given mesh into several frequency bands. Therefore, a new stabilizing technique to encode the geometric difference between the levels will be presented.

## 1 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 and it is desirable to preserve these features under a modification of the global shape.

During the last years, hierarchical representations of geometric shape have become the de facto standard for those purposes. The basic idea is to encode a high-frequency detail level relative to a coordinate frame induced by a coarser approximation of the original shape such that modifications on a coarser level can be propagated to the finer ones. Pioneering work in this area was done by Forsey and Barthels in [6, 7], where they used hierarchical polynomial patches (H-Splines) to represent and edit a surface. 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 one of the reasons, why the interest in surface representations based on triangular meshes increased over the last years. Generalizing the patch-based concepts, the wide family of subdivision techniques [1, 2, 3, 16, 12] start with a coarse base mesh approximating a geometric shape of arbitrary topology and refine 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. When refining a mesh, the position of the inserted control-vertices is predicted by the smoothing-rule of a subdivision-scheme. A detail vector (relative to a local coordinate system induced by the 'parent triangle') will be added to reduce the approximation error. That way, the

new vertex is linked to the coarser level (and follows modifications, if the global shape changes). Storing the base mesh and the sequence of detail vectors for a fixed subdivision scheme leads to a hierarchical representation of the original shape [17, 21]. In practice, one is often given a fine mesh acquired for example from a laser-range scan and an expensive remeshing process [4, 15] is needed to obtain subdivision-connectivity.

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 [8, 10, 14]. 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 it minimizes some global energy functional. Only recently, a couple of new techniques have been proposed [13, 11, 9]. The second step removes the original vertex and encodes the position with respect to its minimized counterpart.

This would require a minimization process for every single vertex. One could also apply the functional 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 changes 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.

## 2 Detail encoding

As mentioned before, we cannot simply store the detail vectors with respect to a global coordinate system but have to define them with respect to local frames which are aligned to the low-frequency geometry [6, 7, 18, 19, 20]. This guarantees the 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 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 then 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




Figure 1: The position of a vertex 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).

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  $p$  of  $\mathcal{M}_{m+1}$  with a corresponding base point  $q$  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  $q$  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 1).

A technique used in [11] is based on the construction of a local quadratic interpolant  $F$  to the low-frequency geometry. For a vertex  $p \in \mathcal{M}_{m+1}$  it is based on the closest triangle  $T \in \mathcal{M}'_{m+1}$  and its adjacent vertices, which can be found in linear time by a simple local search procedure, starting from  $p$ 's corresponding vertex  $p' \in \mathcal{M}_{m+1}$ . Since now a local parameterization is given, parameter values  $(u, v)$  defining the base point  $q$  can be found by Newton-iteration. We start from the center of  $T$  at  $q_0 = F(\frac{1}{3}, \frac{1}{3})$ ,  $q_{n+1}$  is defined by the projection of  $p$  into the tangent plane of  $F$  at  $q_n$ . 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

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

with detail vector  $d = p - q_n$ , which is perpendicular (within a prescribed tolerance) to  $F(u_n, v_n)$  after a few steps. The absolute value of




Figure 2: 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.

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 1 middle).

We therefore propose a different approach which adapts the basic idea of Phong-shading [5] 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  $\triangle(\mathbf{a}, \mathbf{b}, \mathbf{c})$  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  $(\alpha, \beta)$  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

$(\triangle\alpha, \triangle\beta)$  between two consecutive steps can be denoted as follows.

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

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

In case one of the barycentric coordinates of the resulting point  $q$  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. 1 depicts the situation schematically and Fig. 2 shows an example edit where the piecewise constant normal field causes mesh artifacts which do not occur if the Phong normal field is used.

### 3 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 [11] a simple two-band decomposition has been proposed for the modeling, i.e. the high frequency geometry is given 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 3 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 pro-




Figure 3: If the high-frequency detail cannot be projected onto the successive level (top), intermediate levels have to be inserted to guarantee a feasible detail reconstruction (bottom).

jected 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. 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 by solving a global bending energy minimization problem [13, 11, 9]. 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  $n_{i+1}/n_i =$




Figure 4: 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).

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  $\sqrt[i]{\bar{n}_i} = \text{const.}$ . 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  $n_j = \bar{n}_{i_j}$  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. 4 compares the results of a modeling operation based on a two-band and a multi-band decomposition.

## 4 Conclusion and Future research

We have presented a new method to encode high-frequency detail with respect to a low-frequency base mesh. Now, we are able to perform a robust true multi-band decomposition for a given fine triangular mesh of arbitrary connectivity. This leads to intuitive modifications of global shapes under preservation of detail features. However,

the user can still apply particular edits, where undesirable effects like self intersection of detail vectors during the reconstruction process happen, or, 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

- [1] E. Catmull and J. Clark. Recursively Generated B-Spline Surfaces on Arbitrary Topological Meshes. *Computer Aided Design*, 10(6):239–248, November 1978.
- [2] D. Doo and M. Sabin. Behaviour of recursive division surfaces near extraordinary points. *CAD*, 1978.
- [3] N. Dyn, D. Levin, and J. A. Gregory. A butterfly subdivision scheme for surface interpolation with tension control. *ACM Transactions on Graphics*, 9(2):160–169, April 1990.
- [4] Matthias Eck, Tony DeRose, Tom Duchamp, Hugues Hoppe, Michael Lounsbery, and Werner Stuetzle. Multiresolution Analysis of Arbitrary Meshes.

- In *Computer Graphics (SIGGRAPH 95 Proceedings)*, pages 173–182, 1995.
- [5] Foley, van Dam, Feiner, and Hughes. *Computer Graphics*. Addison Wesley, 1990.
  - [6] David R. Forshey and Richard H. Bartels. Hierarchical B-spline refinement. In *Computer Graphics (SIGGRAPH 88 Proceedings)*, pages 205–212, 1988.
  - [7] David R. Forshey and Richard H. Bartels. Surface Fitting with Hierarchical Splines. *ACM Transactions on Graphics*, 14(2):134–161, April 1995.
  - [8] Michael Garland and Paul S. Heckbert. Surface Simplification Using Quadric Error Metrics. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 209–218, 1997.
  - [9] I. Guskov, W. Sweldens, and P. Schröder. Multiresolution signal processing for meshes. Preprint.
  - [10] Hugues Hoppe. Progressive Meshes. In *Computer Graphics (SIGGRAPH 96 Proceedings)*, pages 99–108, 1996.
  - [11] L. Kobbelt, S. Campagna, J. Vorsatz, and H-P. Seidel. Interactive multi-resolution modeling on arbitrary meshes. In *Computer Graphics (SIGGRAPH 98 Proceedings)*, pages 105–114, 1998.
  - [12] Leif Kobbelt. Interpolatory Subdivision on Open Quadrilateral Nets with Arbitrary Topology. In *Computer Graphics Forum, Proceedings of Eurographics '96*, pages C407–C420, 1996.
  - [13] Leif Kobbelt. Discrete fairing. In *Proceedings of the Seventh IMA Conference on the Mathematics of Surfaces*, pages 101–131, 1997.
  - [14] Leif Kobbelt, Swen Campagna, and H-P. Seidel. A general framework for mesh decimation. In *Proceedings of the Graphics Interface conference '98*, 1998.
  - [15] A. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. Multiresolution adaptive parameterization of surfaces. In *Computer Graphics (SIGGRAPH 98 Proceedings)*, pages 95–104, 1998.
  - [16] Charles Loop. Smooth spline surfaces over irregular meshes. In *Computer Graphics Proceedings, Annual Conference Series*, pages 303–310. ACM Siggraph, 1994.
  - [17] Michael Lounsbery, Tony DeRose, and Joe Warren. Multiresolution Analysis for Surfaces of Arbitrary Topological Type. *ACM Transactions on Graphics*, 16(1):34–73, January 1997.
  - [18] David Luebke and Carl Erikson. View-Dependent Simplification of Arbitrary Polygonal Environments. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 199–208, 1997.
  - [19] Jarek Rossignac. Simplification and Compression of 3D Scenes, 1997. Tutorial Eurographics '97.
  - [20] Greg Turk. Re-Tiling Polygonal Surfaces. In *Computer Graphics (SIGGRAPH 92 Proceedings)*, pages 55–64, 1992.
  - [21] Denis Zorin, Peter Schröder, and Wim Sweldens. Interactive multiresolution mesh editing. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 259–268, 1997.




Figure 5: Some snapshots of a modeling session based on the new multi-band hierarchy. The nose of the original bust model was transformed in various ways e.g. scaled (top,right) and translated (bottom row). Notice, how naturally the features of the face are changed.