Mesh Connectivity Compression Using Convection Reconstruction

Raphaëlle Chaine, Pierre-Marie Gandoin, Céline Roudet

Mesh Connectivity Compression Using Convection Reconstruction

Raphaëlle Chaine[1], Pierre-Marie Gandoin[2] and Céline Roudet[1]
[1]LIRIS - UCB Lyon 1 - Bâtiment NAUTIBUS - 8 bd Niels Bohr - 69622 Villeurbanne Cedex - France
[2]LIRIS - Université Lumière Lyon 2 - 5 av Pierre Mendès-France - 69676 Bron Cedex - France
FirstName.LastName@liris.cnrs.fr

Abstract

During a highly productive period running from 1995 to about 2002, the research in lossless compression of 3D meshes mainly consisted in a hard battle for the best bitrates. But for a few years, compression rates seem stabilized around 1.5 bit per vertex for the connectivity coding of usual meshes, and more and more work is dedicated to remeshing, lossy compression, or gigantic mesh compression, where memory and CPU optimizations are the new priority. However, the size of 3D models keeps growing, and many application fields keep requiring lossless compression. In this paper, we present a new contribution for single-rate lossless connectivity compression, which first brings improvement over current state of the art bitrates, and secondly, does not constraint the coding of the vertex positions, offering therefore a good complementarity with the best performing geometric compression methods. The initial observation having motivated this work is that very often, most of the connectivity part of a mesh can be automatically deduced from its geometric part using reconstruction algorithms. This has already been used within the limited framework of projectable objects (essentially terrain models and GIS), but finds here its first generalization to arbitrary triangular meshes, without any limitation regarding the topological genus, the number of connected components, the manifoldness or the regularity. This can be obtained by constraining and guiding a Delaunay-based reconstruction algorithm so that it outputs the initial mesh to be coded. The resulting rates seem extremely competitive when the meshes are fully included in Delaunay, and are still good compared to the state of the art in the general case.

CR Categories: I.3.5 [Computing Methodologies]: Computer Graphics—Computational Geometry and Object Modeling

Keywords: Mesh, Compression, Reconstruction, Lossless, Connectivity

1 Introduction

For several years, meshes have played a leading role in computer graphics, and their ability to model the world throws them in the heart of advanced applications in all the fields of science, arts or leisure. If some of these applications can tolerate a limited loss of information (provided that this loss is well controlled and does not damage a certain visual realism), the others, for practical or even legal reasons, impose to work continuously with exact copies of original objects. In this case, the only way of optimizing the storage and the transmission of 3D information is to have recourse to lossless compression methods.

A mesh is defined by a set of points (we speak of the geometry of the mesh), and by a combinatorial structure describing the relations between these points, using geometric objects of higher dimensions : edges, facets, polyhedra (we speak of the connectivity or the topology of the mesh). To this fundamental information are sometimes added attributes allowing to improve the rendering : normals, colors or textures. If we put aside these attributes (they relate only to a subset of the meshes met in practice), all the difficulty for compressing 3D objects is to process with an equal efficiency the geometry and the connectivity of a wide class of meshes. However, most of the methods proposed so far stress on one of these two aspects (usually the connectivity), generally to the detriment (or at least, not to the advantage) of the other one, which is constrained by a description order rarely optimal in terms of compression. In this article, we propose a new single-rate connectivity coder which is not only general and efficient, but also does not impose any constraint on the coding of the geometry. In particular, it is totally compatible with the position coders that currently propose the best compression rates.

We took it as given that very often, the main part of a mesh connectivity can be deduced from its vertex set using reconstruction methods. This remark has already been made in the past, but the problem was to find an effective way of describing the difference between the initial mesh and its reconstructed version. Indeed, within the framework of lossless compression, it is indispensable to be able to correct these errors in order to obtain a decoded mesh perfectly identical to the original mesh. But rather than describe the difference to the initial object at the end of the reconstruction phase, we suggest adapting an existing algorithm so that it can accept occasional codes modifying its default behavior at the moments when it would commit an "error" of reconstruction with regard to the initial mesh.

Having placed our work in the historical context of 3D compression and 3D reconstruction (Section 2), we will expose the general principle of the method (Section 3), first within a restricted framework, then generalized to arbitrary triangular meshes (regarding the genus, the number of connected components, the regularity and the manifoldness). Then we will detail the coding techniques and optimization steps leading to better compression rates (Section 4), before presenting some comparative results and comments on the method performances (Section 5).

2 Context and Previous Works

2.1 Mesh Compression

As mentioned above, a mesh is composed of both a geometric part (the vertex positions), and a topological part (the vertex connectivity). Now, by looking at the whole scientific production in mesh compression since 1995 until now, we notice clearly that the connectivity coding has motivated most of the proposed methods. The usual scheme consists in describing the combinatorial structure by enumerating its vertices in a precise order, designed to minimize the size of the code : each vertex is coupled with a variable size symbol defining the way it is connected to the rest of the mesh. Consequently, the geometric coder has to work with this predefined

order of the vertices, but it can exploit the connectivity to predict their location. The position of the currently transmitted vertex is estimated from its neighbors, and the prediction error is the sole information to be coded. But the order imposed by the connectivity coding is not necessarily favorable to this prediction, as shown by the results obtained on usual meshes by classical algorithms : the best of them reduce the connectivity information to less than 1 or 2 bits per vertex, while for the geometric part, the rates rarely go down under 90% of the initial size (except for very low quantizations). Among the works that follow this principle, we focus here in single-rate methods, which code and decode the mesh in one pass and do not allow progressive visualization [Deering 1995; Evans et al. 1996; Taubin and Rossignac 1998; Gumhold and Strasser 1998; Touma and Gotsman 1998; Li and Kuo 1998; Bajaj et al. 1999a; Bajaj et al. 1999b; Gumhold et al. 1999; Rossignac 1999; Rossignac and Szymczak 1999; King and Rossignac 1999; Isenburg and Snoeyink 1999; Isenburg 2000; Alliez and Desbrun 2001; Lee et al. 2002; Coors and Rossignac 2004; Kaelberer et al. 2005; Jong et al. 2005]. To understand how these methods compare, the reader can also refer to the following surveys [Alliez and Gotsman 2004; Gotsman et al. 2002; Peng et al. 2005]. It is worth to mention here that the relative stability in compression rates observed these days can be explained by the new interest of the community for gigantic meshes : in this framework, the challenge is to improve the efficiency of the compression in terms of CPU and memory requirements rather than in terms of bitrates [Isenburg and Gumhold 2003; Isenburg and Lindstrom 2005; Isenburg et al. 2005].

2.2 Prioritizing the Geometric Coding

After this first wave of works, in the knowledge that the geometry of a mesh weighs generally much more than its connectivity, the researchers began to propose methods giving the priority to geometric compression. Some of them deviate from the lossless framework by proposing algorithms that often impose a complete remeshing of the object before applying spectral decomposition tools like wavelets, subdivision schemes, or classic geometric methods of compression [Gross et al. 1996; Certain et al. 1996; Lounsbery et al. 1997; Staadt et al. 1998; Guskov et al. 2000; Khodakovsky et al. 2000; Khodakovsky and Guskov 2000; Karni and Gotsman 2000; Szymczak et al. 2002; Attene et al. 2003].

On the other hand, the works introduced by Gandoin and Devillers and improved by Peng and Kuo [Devillers and Gandoin 2000; Gandoin and Devillers 2002; Peng and Kuo 2005] describe a progressive and lossless compression method, centered on the geometry, which interests us particularly within the framework of this article. Indeed, this method was originally designed to code an unstructured point cloud, with the underlying idea that in many cases, the connectivity information could be deduced from the geometry thanks to reconstruction algorithms. So, the first version of this method [Devillers and Gandoin 2000] proposed a multiresolution compression algorithm for an unstructured point cloud, guaranteeing a theoretical minimal gain of $n(log_2 n - 2.4)$ (where n denotes the number of points), and very competitive practical performances, even compared to the best single-rate algorithms. The method was then extended to deal with the connectivity compression, while remaining centered on the geometry [Gandoin and Devillers 2002]. Indeed, the kd-tree decomposition scheme involved in the heart of the geometric coder, is enriched with a connectivity coder that uses some of the classical mesh simplification operations introduced by Hoppe *et al.* [Hoppe 1996; Popović and Hoppe 1997]). Eventually, the method has recently been resumed by Peng and Kuo [Peng and Kuo 2005] who improve its performances by imposing a priority on the cell subdivision order in the octree, and by proposing a more effective prediction model for the point distribution in the sub-cells generated by a subdivision.

2.3 Compression and Reconstruction

The idea to entrust a reconstruction algorithm with the task of computing the connectivity of a mesh from its vertices has already been used in the context of mesh compression. Within the framework of terrain models, Kim *et al.* [Kim et al. 1999] suggest transmitting only a fraction of the edges composing the object, and using constrained Delaunay triangulation — in 2D, since the terrain models are projectable — to find the whole connectivity. Devillers and Gandoin [Devillers and Gandoin 2000] also mention the compression of GIS as an application of their geometric coder, and develop an edge coder well suited for this context, which results in compression rates as low as 0.2 bits per vertex for the complete connectivity. Besides, Devillers *et al.* show that the minimal set to be transmitted to guarantee the exact reconstruction of the initial model by constrained Delaunay triangulation is constituted by the edges that are non locally Delaunay [Devillers et al. 2003].

Unfortunately, the generalization from 2.5D to 3D meshes is not straightforward : since the mesh is not projectable any more, the use of the constrained Delaunay triangulation is impossible. Nevertheless, the idea to use a reconstruction method driven by a partial description of the connectivity remains extremely promising in terms of compression results. Provided that it could be possible to find a method both powerful and capable of taking advantage of partial connectivity data to guarantee an exact reconstruction whatever the initial model may be.

A first attempt to use a reconstruction algorithm to encode connectivity information has been proposed by Lewiner *et al.* [Lewiner et al. 2005]. The geometry is coded independently, through a kd-tree based algorithm derived from [Gandoin and Devillers 2002] and [Botsch et al. 2002] and the connectivity is coded through an advancing front triangulation inspired of the ball-pivoting strategy [Bernardini et al. 1999]. This algorithm requires a code for each edge of an active border that is initialized to a triangle. If the geometry of the mesh meets good sampling conditions, the entropy of each code will be extremely low. In Section 5, we compare this method to ours regarding the compression rates.

2.4 Reconstruction by Convection

The problem of reconstructing a surface from a set of points has received considerable attention during the last decade [Mencl and Müller 1998]. Interesting and outstanding algorithms have been issued both in computer graphics and computational geometry, but we have decided to focus on the algorithms of the second category, since their combinatorial concerns are more suitable for lossless compression purposes. Most of computational geometry algorithms exploit the geometric properties of structures such as the Delaunay triangulation, the Voronoi diagram or the power diagram of the input point set, assuming auspicious properties on the way they were sampled on the surface (ε-sample [Amenta and Bern 1999]). A consistent set of facets can then be extracted from the geometric structure, sometimes using global heuristics or local analysis to solve ambiguities. The existing algorithms are numerous and an attempt to classify and distinguish them has recently been proposed in the state of the art by Cazals and Giesen [Cazals and Giesen 2002]. Most of these algorithms produce triangular meshes, which makes them good candidates to use for compression purposes.

The convection algorithm we use in our method is based on a similar notion of flow as it was developped in the Wrap algorithm by Edelsbrunner [Edelsbrunner 2002], and the flow complex algorithm by Giesen and John [Giesen and John 2002]. Indeed, this reconstruction algorithm has been inspired from a convection process described by Zhao *et al.* [H.K.Zhao et al. 2001]. They use it to ini-

tialize their surface before running an energy minimization process in the level set framework. Given an evolving surface S enclosing an input point set P, the convection process makes each point of S move inwards, along the direction of its normal, towards its closest point in P. However, the numerical scheme they propose can be translated in Delaunay, to make it depend on the geometry of the input data set only, and not on the precision of some grid around the surface. A demonstration of this result is presented by Chaine [Chaine 2003] together with a subsequent convection algorithm. A Delaunay triangulation of P is a partition of space into tetrahedra so that the ball circumscribed to each tetrahedron does not contain any point of P. The convection process is run directly in the 3D Delaunay triangulation of P, with an evolving surface S composed of oriented Delaunay facets. S is initialized to the convex hull of P. An oriented facet of S that does not meet the oriented Gabriel property — *i.e.* its inwards diametral half-sphere is not empty — is attracted towards the 3 inner facets in the incident Delaunay inner tetrahedra (see Fig. 1, for an illustration in 2D where the evolving surface is replaced by an evolving curve, and the Delaunay tetrahedra are replaced by Delaunay triangles). During the convection process, thin parts can appear (see Fig. 1, c and d), on which a 2D surface based version of the convection is run. A deeper explanation of this algorithm will be presented in Section 3 while revisiting it for compression purposes.

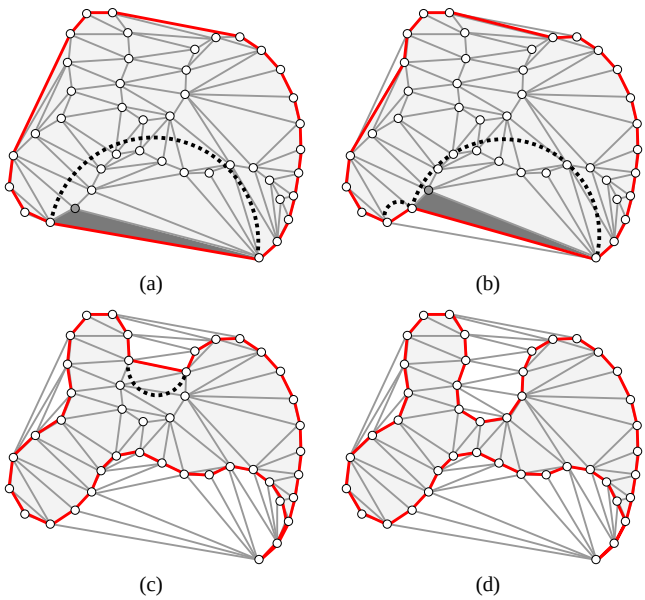

Figure 1: Geometric convection on a 2D point set : (a) the evolving curve C is initialized to the convex hull, (b) C locally evolves at the level of an edge iff the half-circle associated to this edge is not empty, (c) result of the initial convection process, (d) the convection process is locally enhanced to hollow a pocket out.

An interesting property of the convection algorithm is that it is driven locally, in the Delaunay triangulation of the points, without involving a global heuristic. The topology of the evolving surface may change during the convection process, so that it can handle surfaces with boundaries and surfaces of high genus. A drawback of the convection process is that it can locally be stuck in presence of pockets [Edelsbrunner et al. 1998] hiding facets, but a simple and local point density analysis permits to hollow them out [Chaine 2003].

3 Principle of the Compression Algorithm

3.1 Compliant Meshes

The benefit of using a 3D reconstruction method for compression purposes is double : first, this allows to obtain very low costs for the coding of the connectivity — ideally, a null cost if the reconstruction algorithm is able to find the exact structure of the original mesh by itself —, and secondly, unlike the previous methods of topologic compression, no constraint is imposed on the order of the vertices to the geometric coder, which constitutes an important efficiency token.

The main difficulty consists in being able to help the reconstruction algorithm at a reasonable cost : indeed, it is highly improbable to design an algorithm capable of finding the complete connectivity of a mesh from its vertex set. It is thus necessary to be able to alter the course of the reconstruction process, by occasionally changing the default behavior of the algorithm to drive it in a sure way to a result known in advance : the structure of the initial mesh.

Among the plethora of available methods, the reconstruction by convection is distinguishable from others by two important assets : its qualities in terms of reconstruction — practical accuracy and faithfulness to the original model, handling of complex topologies, computation times —, and above all, its ability to be effectively driven by means of simple instructions. The first asset guarantees that the algorithm will not need to be guided too often (small number of codes), the second guarantees that this can be done at lower cost (compactness of the codes).

In return, as many algorithms in computational geometry, the convection algorithm is based on a Delaunay triangulation, and it may be difficult to force it out of this structure. This imposes a condition over the initial mesh : all its facets have to belong to the 3D Delaunay triangulation of its vertex set. It is quite a strong property that is not verified by all the 3D objects met in practice. We will see in second stage (Section 3.3) how to break it.

Intuitively, the reconstruction by convection consists in embedding the vertex set in a block of marble that will be sculpted gradually, facet after facet, tetrahedron after tetrahedron, until it takes on the exact shape of the initial object. The algorithm begins by preparing the sculpting material : an enclosing block which is nothing else than the convex hull of the point cloud, as well as the network of galleries through which it is going to dig until the initial shape is reached. This network is composed of the tetrahedra of the 3D Delaunay triangulation, and every stage towards the object shape consists in examining a new facet of the current surface and deciding if it is necessary or not to open it and excavate the inside of its associated tetrahedron. When a facet is opened, it is removed from the current surface of the 3D object under construction, and replaced by the three other facets of the excavated tetrahedron.

As mentioned above, the criterion that decides on this opening is purely local : it consists in observing whether the Gabriel half-sphere associated to the current oriented facet contains or not the fourth vertex of the tetrahedron. If it is the case, the oriented facet is opened and replaced in the current surface by the 3 other facets of the associated tetrahedron. (If this one is already excavated, the surface locally vanishes through the current facet.) Otherwise, it is maintained on the surface.

Note that if all the facets of the initial object are in its 3D Delaunay triangulation, they are reachable by this sculpture process from the convex hull. The problem is to make sure that the algorithm will not dig a facet belonging to the initial mesh, or that, conversely, it will not be retained before having reached such a facet. Hence the need for additional codes allowing to modify the behavior of

the convection algorithm, and to drive it towards the object to be coded. Even if the convection algorithm was designed to compute with a good accuracy the structure of a mesh sufficiently dense, it will most likely need occasional assistance to guarantee the perfect reconstruction of any mesh.

To present our method in a progressive way, we focus in this section on a mesh M verifying the following two assumptions : all its facets are in the 3D Delaunay triangulation of its vertex set, and it is manifold, that is to say without borders nor thin parts. (This is what we call a "compliant" mesh.)

Under these assumptions, here is the general principle of our compression algorithm :

1 Build the 3D Delaunay triangulation D of the point set P,

2 Mark the facets of D belonging to the initial mesh M,

3 Initialize the current surface S with the convex hull of P (included in D),

4 Launch the convection process on S : every oriented facet f in S is examined in turn, and depending on whether its Gabriel half-sphere is empty or not, f is, by default, maintained on S or replaced by its 3 neighbors in D. To modify this default behavior of the convection reconstruction, it is necessary to locate the oriented facets of S for which the algorithm has to act differently. It is thus indispensable to define a completely deterministic traversal of the facets in S, so that the nth oriented facet met during the compression would also be the nth oriented facet met during the decompression. More generally, it is necessary to ensure the synchronization of the algorithms of compression and decompression so that the nth action of the coder matches the nth action of the decoder. Thanks to these reference marks, the behavior of the convection process can be safely altered in the following two circumstances :

   a) when the convection asks for the opening of a facet f that belongs to M, the coder forbids this opening, and codes the index of f (more exactly, the moment when f is met in the algorithm) to warn the decoder that at this precise moment of the reconstruction, it has to break the rules of the convection algorithm. We will call this a RDG event (Retain Despite the Geometry),

   b) at the end of this first step, it is possible that some oriented facets of S — those whose Gabriel half-spheres are empty and thus the convection did not decide to dig —, do not belong to M. Therefore, it is necessary to specify to the decompression algorithm that these oriented facets must be forced, against the convection rules : for each remaining oriented facet of S not belonging to M, the coder transmits its index (*i.e.* the moment when it has been met) and relaunch the convection process by forcing its opening. We will call this a ODG event (Open Despite the Geometry). Note that by relaunching the convection process locally, some facets that were due to be forced can disappear by autointersection of S.

A detailed description of step 4 can be given through a recursive function *Convection* applied to the current surface S. For each oriented facet f, let f_{asso} denote the oriented facet associated to f, that is to say the oriented facet constructed on the same vertices as f, but with the opposite orientation. When f and f_{asso} are both in the surface S, that means they belong to a thin part of S. Besides, let f_1, f_2 and f_3 denote the 3 neighbors of f in the 3D Delaunay triangulation D, *i.e.* the 3 other facets of the tetrahedron that is removed when f is opened. Finally, we would like to draw the reader's attention to the fact that the reference marks transmitted to the decoder are not exactly absolute moments of events in the compression process, but rather intervals between two such moments (which explains the presence of the instructions "$time \leftarrow 0$" in the detailed algorithms). This well-known technique of differential coding allows to reduce the size of the transmitted codes.

```
Function Convection(S : Surface) :
  while S ≠ ∅ do
    f ← pop first oriented facet in S
    if Gabriel half-sphere of f is empty then
      push f at the end of S_temp
    else
      if f ∈ M then {RDG event}
        output time
        time ← 0
        push f at the end of S_final
      else
        if f such that f_asso ∈ S (resp. S_temp) then
          remove f_asso from S (resp. S_temp)
        else
          push {f_1, f_2, f_3} at the end of S
        end if
        time ← time + 1
      end if
    end if
  end while
```

With this recursive definition of the *Convection* function, the step 4 of our algorithm amounts to the following :

```
Main function :
  S ← convex hull of P
  S_border ← ∅ {set of facets creating a thin part}
  S_final ← ∅ {set of facets that are in M}
  S_temp ← ∅ {set of facets candidates to be in M}
  time ← 0
  Convection(S)
  while S_temp ≠ ∅ do
    f ← pop first oriented facet in S_temp
    if f such that f_asso ∈ S_temp then
      remove f_asso from S_temp
      push {f, f_asso} at the end of S_border
    else
      if f ∈ M then
        push f at the end of S_final
        time ← time + 1
      else {ODG event}
        output time
        time ← 0
        S ← {f_1, f_2, f_3}
        Convection(S)
      end if
    end if
  end while
```

At the end of the compression algorithm, the convection has been driven towards the initial mesh M, and the correcting data have been stored for the decompression algorithm. However, a last stage remains that consists in cleaning thin parts possibly generated by the convection. Indeed, since we first assumed that the initial mesh was manifold, it suffices to delete all these thin parts, that is to say each facet of S whose associated facet is also on S. In the algorithm described above, the thin parts exactly match the set S_{border}.

3.2 Non Manifold Meshes

In this section, we are going to adapt the previous algorithm to non manifold, thin parts meshes. It suffices to modify the final stage : it is no longer possible to delete the thin parts created by the algorithm, because facets belonging to the initial mesh could be lost. This stage is thus replaced by a new convection process, but this time in its 2D version, and on the thin parts only. In this framework, the convection updates a curve C constituted by the edges composing the boundaries of the thin parts. The oriented edges of C are examined in turn : an oriented edge whose Gabriel half-circle is empty will be kept on C, whereas an oriented edge in the opposite case will be removed and replaced in C by the two other oriented edges of its incident triangle (e_1 and e_2 in the algorithm detailed below).

The general principle remains the same as in the 3D version of the algorithm, and each edge that is opened against the convection rules (ODG event) launches a new 2D convection process. Note that there is no set C_{border} similar to the previous S_{border}. Here is a detailed version of the portion of the compression algorithm dedicated to the processing of the thin parts. It adds to the steps 1 to 4 described in the previous section :

5 Processing of thin parts

```
Function Convection_2D(C : Curve) :
  while C ≠ ∅ do
    e ← pop first oriented edge in C
    if Gabriel half-circle of e is empty then
      push e at the end of C_temp
    else
      if e ∈ M then {RDG event}
        output time
        time ← 0
        push e at the end of C_final
      else
        if e such that e_asso ∈ C (resp. C_temp) then
          remove e_asso from C (resp. C_temp)
        else
          push {e_1, e_2} at the end of C
        end if
        time ← time + 1
      end if
    end if
  end while

Main 2D function :
  C ← boundaries of S_border
  C_final ← ∅; C_temp ← ∅; time ← 0
  Convection_2D(C)
  while C_temp ≠ ∅ do
    e ← pop first oriented edge in C_temp
    if e such that e_asso ∈ C_temp then
      remove e_asso from C_temp
    else
      if e ∈ M then
        push e at the end of C_final
        time ← time + 1
      else {ODG event}
        output time
        time ← 0
        C ← {e_1, e_2}
        Convection_2D(C)
      end if
    end if
  end while
```

3.3 Non Delaunay Meshes

As previously described, the method can only be applied to meshes whose facets all belong to the 3D Delaunay triangulation of their vertex set. Indeed, we saw that this structure constituted the support of the convection algorithm, and that it was thus impossible for the current surface to reach a non Delaunay facet. Nevertheless, most of the meshes met in practice contain a fraction of non Delaunay facets (see the unfavorable case of Fig. 2). So, to make the method widely usable, it is necessary to manage the coding of such facets. A simple and efficient way to do this is to code explicitely all the non Delaunay facets of the initial mesh in the same time as its vertices, using the method of Gandoin and Devillers [Gandoin and Devillers 2002], which we will note GD in the following. More precisely, non Delaunay facets constitute patches on the surface of the mesh. Before launching the convection process, the connectivity of these patches is transmitted to the GD coder. Then, the algorithm previously described is applied to the mesh minus the non Delaunay facets. Consequently, even if the initial mesh is manifold, the convection process is going to perform on a mesh with boundaries. As shown in the previous section, it is no theoretical problem, but risks though to result in a significant increase of the number of driving codes. Indeed, each facet of the mesh will be reached twice, first through a facet oriented outwards, then through the associated oriented facet lying on the internal side of the object surface. We propose several heuristics in order to limit this phenomenon. The intuitive idea is to block temporarily the convection surface when it is about to cross a patch — a facet opening is delayed if the tetrahedron behind it intersects the mesh —, so as to favor the discovery of the initial mesh facets from outside (see Fig. 3). Thus, transmitting the total number of mesh facets to the decoder will allow to stop the convection process as soon as they are all discovered, which will drastically reduce the number of corrective codes.

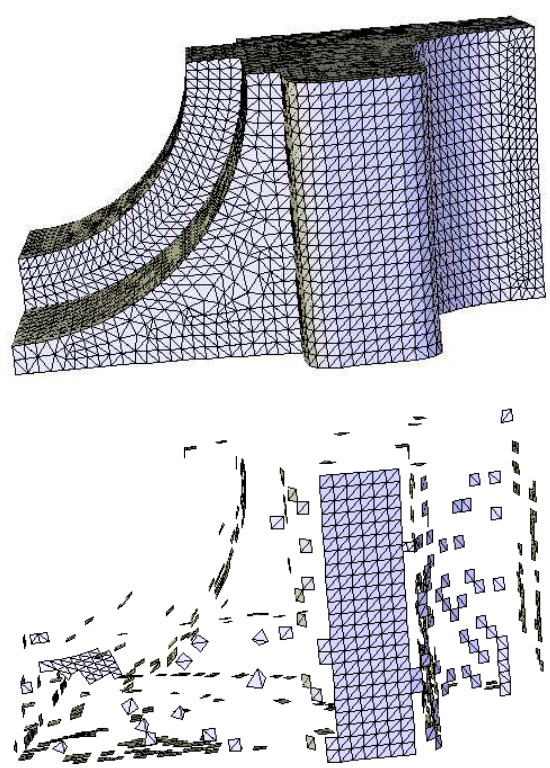

Figure 2: Fandisk : complete (12946 facets), then showing the non Delaunay facets only (1104 facets representing 8.5% of the set)

In the decoding stage, the vertex set is first obtained from the GD decoder, as well as the patches connectivity. Then the mesh without the patches is reconstructed by the algorithm of driven convection. At last, the patches connectivity is merged into the reconstructed leaky mesh. We thus see that our method improves the results of the method of Gandoin and Devillers only for meshes whose facets are mainly Delaunay. In the limit case where all the mesh facets would be non Delaunay, we would come across the GD coder compression rates.

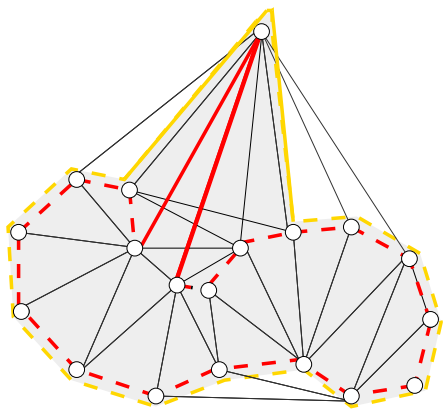

Figure 3: 2D Example : the original curve to be coded is in red thick line, including non Delaunay patches (parts in solid line). The Delaunay triangulation is in black thin line. The current evolving convection curve, in orange thick line, is temporarily blocked (parts in solid line) to avoid intersecting the patches. The hidden Delaunay edges of the original curve will thus be discovered later, when the convection process will be relaunched to intersect the patches.

4 Coding and Optimization

We saw in section 3 that the codes resulting from the compression algorithm are a sequence of positive or null integers, representing the number of facets met between two special interventions of the algorithm, that is to say between two facets for which the algorithm does not follow the rules of the reconstruction by convection (ODG and RDG events). To minimize the output size, we have chosen to transmit these numbers to an arithmetic coder, which is able to code a sequence of symbols in a nearly optimal way, given a probability model. Thus, if $p(s)$ denotes the probability of appearance of the symbol s, the arithmetic coder will encode s on $log_2(1/p(s)) + \varepsilon$. In particular, this entropic coder is capable of coding a symbol on a fractional number of bits.

The main difficulty consists in defining a good probability model for the integers to be coded. The first solution consists in computing statistical data for each mesh and transmitting the probability table in the header of the compressed file. But ideally, to save the transmission of such a table, we would like to model for a wide class of meshes the behavior of the sequence to be coded, or alternatively, to design an adaptive model recomputing the probability of a symbol each time it occurs. Several kinds of model can be used, according to the size of the context from which the probability is estimated. For the order-0 model, each integer has an absolute probability for the whole coding sequence, independent from the context. For the order-1 model, the probability of an integer depends on the value of the previously transmitted integer, and so on. We have finally opted for an adaptive order-1 model specifically designed to take advantage of the particular structure of the integer sequence. For instance, this model handles the long runs of null integer frequently occuring in the sequence.

As a matter of fact, these runs often correspond to large patches of facets that have to be forced (ODG events) because the Gabriel criterion is too restrictive compared to the local density of the mesh. Statistics show that the convection process is rarely wrong when it decides to open a facet; on the contrary, at the end of the convection, a lot of facets have been retained but do not belong to the initial mesh. To encourage the convection algorithm to open more facets, we have relaxed the Gabriel criterion under some conditions. Intuitively, a facet will be opened when its size is "big" with regard to the vertex density in its neighborhood [Chaine 2003]. By setting this ratio to a near optimal value, the improvements can be drastic, particularly in the case of poorly sampled meshes where the number of driving codes can be lowered by about 50%.

5 Experimental Results

The table 1 shows the results of the method applied to some usual meshes (the rates are in bits per vertex). The objects *fandisk*, *blob*, and *horse* are there essentially for comparison purposes, since they have been used for example by Touma and Gotsman [Touma and Gotsman 1998]. The Stanford *bunny* is also a very classical mesh, widely used in 3D compression for about ten years. However, these models are not very representative of today meshes, whose number of vertices is much higher. The last three meshes are more typical of what can be found nowadays. The *hand* (see Fig. 6) and *max_planck* (see Fig. 5) can be found on *http://shapes.aim-at-shape.net*, while the *triple_hecate* (see Fig. 4) comes from *Le Louvre C2RMF lab*.

model (number of vertices)	number of Delaunay (non Del.) facets	connectivity rate in bpv (Del. + non Del. facets)	computing time in seconds (compression / decompression)
fandisk (6475)	11842 (1104)	2.08 (1.22 + 0.86)	14.45s / 14.48s
blob (8036)	15226 (842)	2.68 (1.98 + 0.70)	16.69s / 14.64s
horse (19851)	38173 (1525)	2.00 (1.48 + 0.52)	26.38s / 22.01s
bunny (35947)	71890 (0)	0.08 (0.08 + 0.00)	17.60s / 13.66
triple_hecate (75729)	151462 (0)	0.04 (0.04 + 0.00)	46.30s / 37.15s
max_planck (199169)	391181 (6862)	0.90 (0.69 + 0.21)	200.87s / 156.75s
hand (327290)	649922 (4674)	0.19 (0.11 + 0.08)	341.52s / 287.57s
ajax (273383)	547117 (0)	0.14 (0.14 + 0.00)	153.63s / 117.13s

Table 1: Experimental results on usual meshes

The first remark is that the rates obtained for the first three models are not especially competitive with regard to the best current methods (for example, [Kaelberer et al. 2005], whose algorithm yields 0.74 bpv for the *fandisk*, and 0.96 bpv for the *horse*). The main reason is that the vertex set of these meshes are clearly not ε-sample and therefore, not favorable to the convection algorithm (but note that on these meshes, we gain on the geometry coding by using the GD coder, as shown in [Gandoin and Devillers 2002]). On the contrary, the *hand* and *max_planck* are correctly sampled and give a quite good idea of what the method can achieve when the point densities are reasonable. The *triple_hecate* and *ajax* models are highly compliant meshes, not only because their facets are fully in Delaunay, but also because they were constructed from a set of points,

using some reconstruction algorithm similar to the convection process. However, this kind of mesh is more and more widespread since a large class of objects are obtained from scanning and reconstruction.

Besides, we can notice that the rates associated to the coding of non Delaunay facets are quite bad facing their small number, and heavy penalize the global rates. Indeed, the GD coder is not optimal for sparse connectivity, and it should be possible to find a better way to code this small set of facets.

The last column of the table shows the computing user times in seconds for the connectivity compression / decompression of the meshes on a Pentium IV 3.0 Ghz 2 Go RAM computer. Globally, these times are rather high compared to classical methods that explicitly encode the whole connectivity, using some traversal through the mesh vertices. One can incriminate the precomputation and the traversal of the Delaunay tetrahedrization required by the convection process. This constraint particularly intends our method to applications where storage space or network bandwidth are more limited resources than processing power. However, a second version of the algorithm could be developped where Delaunay computation is not explicit any more. The current version also encounters timings limitations when the mesh is not entirely included in Delaunay. This is due to intersection determinations between the Delaunay tetrahedrization and non-Delaunay mesh facets.

We can compare our algorithm to the only other compression method using reconstruction through the 3 meshes we have in common : the method [Lewiner et al. 2005] obtains 1.19 bpv for the horse, 2.63 bpv for the fandisk, and 1.18 bpv for the bunny, which is globally slightly higher than the rates of the table 1. Regarding the methods [Alliez and Desbrun 2001; Lee et al. 2002; Kaelberer et al. 2005], derived from the Touma and Gotsman's coding principle [Touma and Gotsman 1998], they obtain rates around 1.5 bpv for usual meshes, and can achieve very low rates for highly regular meshes where vertex degrees are almost constant. We don't have any result of these algorithms for meshes fully included in Delaunay, but there is no particular reason why their rates would be better in such cases.

Another limitation of the algorithm concerns the memory footprint. This footprint is that of a Delaunay triangulation enriched with mesh information. The compression algorithm is implemented in CGAL [CGAL] and the Delaunay tetrahedrization size is proportional to the number of Delaunay cells and the number of Delaunay vertices. Note that the Delaunay tetrahedrization of a surface based point set meeting good sampling conditions is nearly linear in the number of points.

6 Conclusion and Future Work

We have presented a new method of lossless single-rate connectivity compression based on a semiautomatic reconstruction process able to deduce most of the mesh connectivity from its sole vertex set, and occasionally guided through compact codes that alter its default behavior when necessary. This method is originally suitable for Delaunay embedded meshes. We have also described a generalization removing this constraint inherent to the convection algorithm, using a separated coding of non Delaunay patches, as well as heuristics minimizing the number of driving codes during the convection process. So the final algorithm is able to compress any kind of 3D mesh, including non manifold ones, without any constraint on the topological genus or the number of connected components. After a probabilistic modelization and an entropic coding of the output, the numerical results show a substancial improvement above the current state of the art, with an average rate of 1 bit per vertex, reaching below 0.1 bpv for well sampled meshes. In addition, the algorithm can be used in parallel with currently most performing geometric compression methods, which results in very competitive overall rates. The main limitations of the method are its computing times, relatively high compared to some of the state-of-the-art algorithms, and its memory footprint, which exclude for now the compression of gigantic meshes on standard 32-bit machines. As a result, the optimization of the compression, and above all, of the decompression algorithm remains a primary perspective for this work.

The compression rates could probably be improved through more accurate probability models for the entropic coder, and even more, through a better coding of the non Delaunay facets. Similarily, it could be possible to keep improving the behavior of the convection algorithm, by opening the facets more accurately, according to some smarter criterion than the local vertex distribution. However, we are currently working on the generalization of this work to multiresolution, which is surely the most important perspective of this paper. This will be made possible by progressive insertion of vertices in the convection surface [Allègre et al. 2005], under the constraints imposed by a kd-tree type progressive geometric coder. About the GD coder, a last significant perspective consists in integrating it to the convection process, in order to improve the geometric prediction by using connectivity information.

Acknowledgement

This research takes place into the ACI projects Art3D and Eros3D, supported by the French Centre National de la Recherche Scientifique (CNRS). Some of the 3D models appearing in this paper are provided courtesy of AIM@Shape (IMATI and INRIA). The triple_hecate model is provided courtesy of Christian Lahanier from C2RMF (Musée du Louvre, département des Antiquités Etrusques et Romaines, marbre, MA 2594, C2RMF : FZ31452).

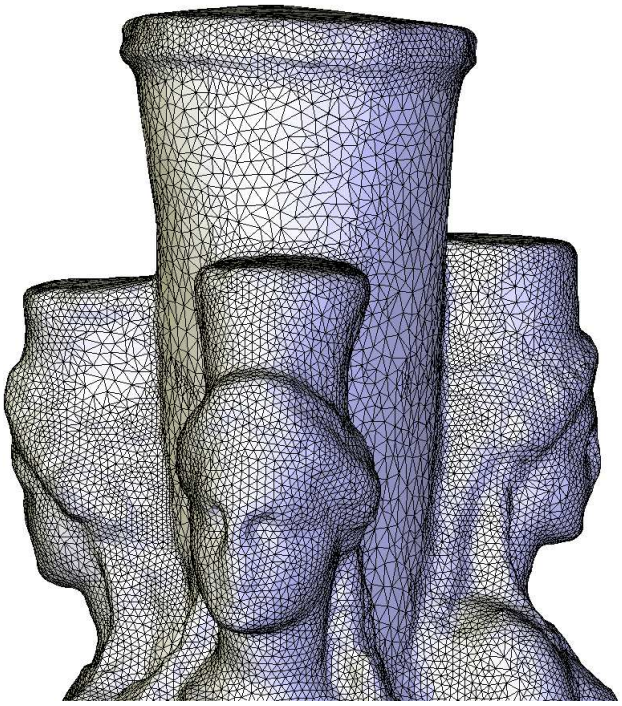

Figure 4: triple_hecate (detail) : 75729 vertices and 151462 facets coded with 0.04 bits per vertex

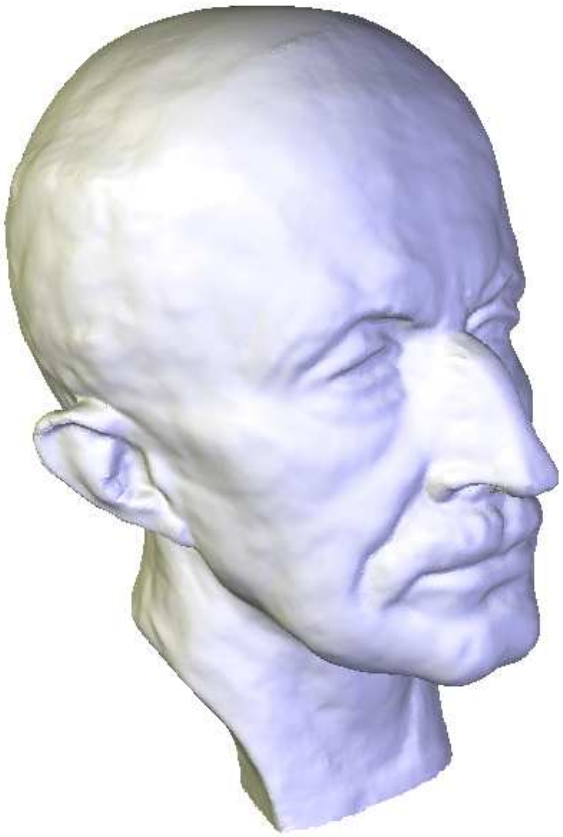

Figure 5: max_planck : 199169 vertices and 398043 facets coded with 0.90 bits per vertex

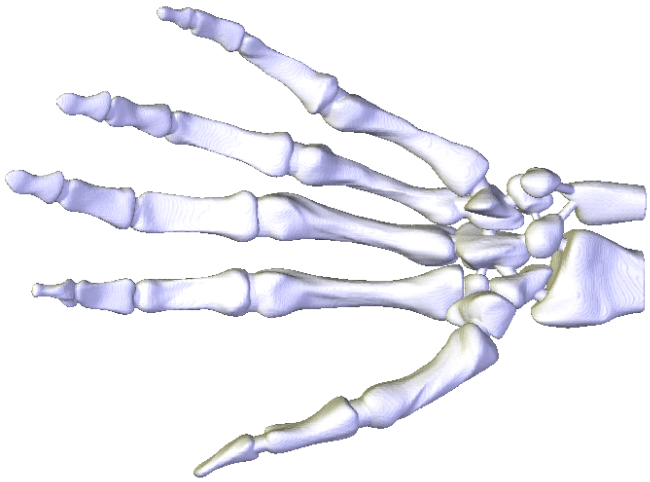

Figure 6: hand : 327290 vertices and 654596 facets coded with 0.19 bits per vertex

References

ALLÈGRE, R., CHAINE, R., AND AKKOUCHE, S. 2005. Convection-driven dynamic surface reconstruction. In *Proc. Shape Modeling International, IEEE Computer Society Press*, 39–48.

ALLIEZ, P., AND DESBRUN, M. 2001. Valence-driven connectivity encoding for 3d meshes. In *Eurographics 2001 Conference Proc.*

ALLIEZ, P., AND GOTSMAN, C. 2004. *Recent Advances in Compression of 3D Meshes*. Springer.

AMENTA, N., AND BERN, M. 1999. Surface reconstruction by voronoi filtering. *Discrete Comput Geom. 22*, 4, 481–504.

ATTENE, M., FALCIDIENO, B., SPAGNUOLO, M., AND ROSSIGNAC, J. 2003. Swingwrapper : Retiling triangle meshes for better edgebreaker colmpression. *ACM Transactions on Graphics 22*, 4, 982–996.

BAJAJ, C., CUTCHIN, S., PASCUCCI, V., AND ZHUANG, G. 1999. Error resilient streaming of compressed vrml. Tech. rep., University of Texas, Austin.

BAJAJ, C., PASCUCCI, V., AND ZHUANG, G. 1999. Single resolution compression of arbitrary triangular meshes with properties. *Computational Geometry : Theory and Applications*.

BERNARDINI, F., MITTLEMAN, J., RUSHMEIER, H., SILVA, C., AND TAUBIN, G. 1999. The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Visualization and Computer Graphics 5*, 4 (/), 349–359.

BOTSCH, M., WIRATANAYA, A., AND KOBBELT, L., 2002. Efficient high quality rendering of point sampled geometry.

CAZALS, F., AND GIESEN, J. 2002. Delaunay triangulation based surface reconstruction : Ideas ans algorithms. Tech. rep., Technical Report 5393, INRIA.

CERTAIN, A., POPOVIĆ, J., DEROSE, T., DUCHAMP, T., SALESIN, D., AND STUETZLE, W. 1996. Interactive multiresolution surface viewing. In *SIGGRAPH 96 Conference Proc.*

CGAL. *http://www.cgal.org*.

CHAINE, R. 2003. A geometric convection approach of 3-d reconstruction. In *Eurographics Symposium on Geometry Processing, Aachen, Germany*, 218–229.

COORS, V., AND ROSSIGNAC, J. 2004. Delphi : Geometry-based connectivity prediction in triangle mesh compression. *The Visual Computer 20*, 8-9, 507–520.

DEERING, M. 1995. Geometry compression. In *SIGGRAPH 95 Conference Proc.*, 13–20.

DEVILLERS, O., AND GANDOIN, P.-M. 2000. Geometric compression for interactive transmission. In *IEEE Visualization 2000 Conference Proc.*

DEVILLERS, O., ESTKOWSKI, R., GANDOIN, P.-M., HURTADO, F., RAMOS, P., AND SACRISTÁN, V. 2003. Minimal set of constraints for 2D constrained Delaunay reconstruction. *International Journal of Computational Geometry and Applications 13*, 391–398.

EDELSBRUNNER, H., FACELLO, M., AND J.LIANG. 1998. On the definition and the construction of pockets in macromolecules. *Discrete Appl. Math 88*, 83–102.

EDELSBRUNNER, H. 2002. Surface reconstruction by wrapping finite point sets in space. *B. Aronov, S. Basu, J. Pach and Springer-Verlag M. Sharir, editors, Ricky Pollack and Eli Goodman Festschrift*, 379–404.

EVANS, F., SKIENA, S., AND VARSHNEY, A. 1996. Optimizing triangle strips for fast rendering. In *IEEE Visualization 96 Conference Proc.*

GANDOIN, P.-M., AND DEVILLERS, O. 2002. Progressive lossless compression of arbitrary simplicial complexes. In *ACM SIGGRAPH Conference Proc.*

GIESEN, J., AND JOHN, M. 2002. Surface reconstruction based on a dynamical system. In *Proc. Eurographics*.

GOTSMAN, C., GUMHOLD, S., AND KOBBELT, L. 2002. *Simplification and Compression of 3D Meshes*. Springer.

GROSS, M. H., STAADT, O. G., AND GATTI, R. 1996. Efficient triangular surface approximations using wavelets and quadtree data structures. *IEEE Transactions on Visualization and Computer graphics*.

GUMHOLD, S., AND STRASSER, W. 1998. Real time compression of triangle mesh connectivity. In *SIGGRAPH 98 Conference Proc.*, 133–140.

GUMHOLD, S., GUTHE, S., AND STRASSER, W. 1999. Tetrahedral mesh compression with the cut-border machine. In *IEEE Visualization 99 Conference Proc.*

GUSKOV, I., VIDIMCE, K., SWELDENS, W., AND SCHRÖDER, P. 2000. Normal meshes. In *SIGGRAPH 2000 Conference Proc.*

H.K.ZHAO, OSHER, S., AND FEDKIW, R. 2001. Fast surface reconstruction using the level set method. In *Proceedings of IEEE Workshop on Variational and Level Set Methods in Computer Vision (VLSM)*.

HOPPE, H. 1996. Progressive meshes. In *SIGGRAPH 96 Conference Proc.*

ISENBURG, M., AND GUMHOLD, S. 2003. Out-of-core compression for gigantic polygon meshes. In *SIGGRAPH 2003 Conference Proc.*

ISENBURG, M., AND LINDSTROM, P. 2005. Streaming meshes. In *IEEE Visualization 2005 Conference Proc.*

ISENBURG, M., AND SNOEYINK, J. 1999. Mesh collapse compression. In *Symposium on Computational Geometry*.

ISENBURG, M., LINDSTROM, P., AND SNOEYINK, J. 2005. Streaming compression of triangle meshes. In *Eurographics Symposium on Geometry Processing*.

ISENBURG, M. 2000. Triangle fixer : Edge-based connectivity encoding. In *16th European Workshop on Computational Geometry Proc.*

JONG, B. S., YANG, W. H., TSENG, J.-L., AND LIN, T. W. 2005. An efficient connectivity compression for triangular meshes. In *ACIS-ICIS*, 583–588.

KAELBERER, F., POLTHIER, K., REITEBUCH, U., AND WARDETZKY, M. 2005. Freelence - coding with free valences. In *Eurographics 2005 Conference Proc.*

KARNI, Z., AND GOTSMAN, C. 2000. Spectral compression of mesh geometry. In *SIGGRAPH 2000 Conference Proc.*

KHODAKOVSKY, A., AND GUSKOV, I. 2000. Normal mesh compression. In *to appear*.

KHODAKOVSKY, A., SCHRÖDER, P., AND SWELDENS, W. 2000. Progressive geometry compression. In *SIGGRAPH 2000 Conference Proc.*

KIM, Y.-S., PARK, D.-G., JUNG, H.-Y., AND CHO, H.-G. 1999. An improved TIN compression using Delaunay triangulation. In *Pacific Graphics 99 Conference Proc.*

KING, D., AND ROSSIGNAC, J. 1999. Guaranteed 3.67v bit encoding of planar triangle graphs. In *Canadian Conference on Computational Geometry Proc.*

LEE, H., ALLIEZ, P., AND DESBRUN, M. 2002. Angle-analyzer: A triangle-quad mesh codec. In *Eurographics 2002 Conference Proc.*

LEWINER, T., CRAIZER, M., LOPES, H., PESCO, S., VELHO, L., AND MEDEIROS, E. 2005. Gencode : Geometry-driven compression in arbitrary dimension and co-dimension. In *Sibgrapi*.

LI, J., AND KUO, C.-C. J. 1998. A dual graph approach to 3d triangular mesh compression. In *IEEE International Conference on Image Processing Proc.*

LOUNSBERY, M., DEROSE, T., AND WARREN, J. 1997. Multiresolution analysis for surfaces of arbitrary topological type. *ACM Transactions on Graphics*.

MENCL, R., AND MÜLLER, H. 1998. Interpolation and approximation of surfaces from three-dimensional scattered data points. In *State of the Art Reports, Eurographics*, 51–67.

PENG, J., AND KUO, C.-C. J. 2005. Geometry-guided progressive lossless 3d mesh coding with octree decomposition. In *ACM SIGGRAPH Conference Proc.*

PENG, J., KIM, C.-S., AND KUO, C.-C. J. 2005. Technologies for 3d mesh compression : A survey. *ELSEVIER Journal of Visual Communication and Image Representation 16*, 6, 688–733.

POPOVIĆ, J., AND HOPPE, H. 1997. Progressive simplicial complexes. In *SIGGRAPH 97 Conference Proc.*

ROSSIGNAC, J., AND SZYMCZAK, A. 1999. Wrap&zip : Linear decoding of planar triangle graphs. *Computational Geometry : Theory and Applications*.

ROSSIGNAC, J. 1999. Edgebreaker : Connectivity compression for triangle meshes. *IEEE Transactions on Visualization and Computer Graphics*, 47–61.

STAADT, O. G., GROOS, M. H., AND WEBER, R. 1998. Multiresolution compression and reconstruction. In *Eurographics 98 Conference Proc.*

SZYMCZAK, A., ROSSIGNAC, J., AND KING, D. 2002. Piecewise regular meshes : Construction and compression. *Graphical Models 64*, 3-4, 183–198.

TAUBIN, G., AND ROSSIGNAC, J. 1998. Geometric compression through topological surgery. *ACM Transactions on Graphics 17*, 2.

TOUMA, C., AND GOTSMAN, C. 1998. Triangle mesh compression. In *Graphics Interface 98 Conference Proc.*, 26–34.