ACADEMIC PRESS

Available at www.ComputerScienceWeb.com
POWERED BY SCIENCE DIRECT®

Graphical Models 65 (2003) 131–148

Graphical Models

www.elsevier.com/locate/gmod

Topological quadrangulations of closed triangulated surfaces using the Reeb graph

Franck Hétroy* and Dominique Attali

LIS Laboratory, Domaine Universitaire, BP 46, 38402 Saint Martin d'Hères cedex, France

Received 15 May 2002; received in revised form 4 September 2002; accepted 17 October 2002

Abstract

Although surfaces are more and more often represented by dense triangulations, it can be useful to convert them to B-spline surface patches, lying on quadrangles. This paper presents a method for the construction of coarse topological quadrangulations of closed triangulated surfaces, based on Morse theory. In order to construct on the surface a quadrangulation of its canonical polygonal schema, we compute first a Reeb graph then a canonical set of generators embedded on the surface. Some results are shown on different surfaces.

Keywords: Triangulated surfaces; Quadrangulation; Surface generators; Reeb graph; Dijkstra's algorithm

1. Introduction

Surfaces of arbitrary shape and topology are often represented in computer graphics by large fine triangle meshes, obtained with complex data acquisition hardware. Such discrete representations of surfaces are not easy to store and handle because of the huge amount of data (up to several millions of triangles). Many works trying to reduce the number of triangles have been carried out [11]. Other approaches convert these dense triangle mesh representations to other suitable models, such as parametric surfaces [6,14].

Our goal is to create a new discrete description of a surface involving large non planar quadrangles, in order to later fit B-spline surface patches on them (one patch

[*] Corresponding author.
E-mail addresses: franck.hetroy@lis.inpg.fr (F. Hétroy), dominique.attali@lis.inpg.fr (D. Attali).

1524-0703/03/$ - see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S1524-0703(03)00005-5

per quadrangle). We decided to work with quadrangles instead of triangles or a mix of both, because schemes based on 4-sided patches are more common. The problem of fitting spline patches on quadrangles has been well studied, see works by Peters [19] or Lai and Schumaker [15]. We focus on the determination of the quadrangles.

To define such a description of a surface, strategies can be very different whether we try to optimize the shape of the quadrangles, the number of quadrangles or the number of vertices whose valence is different from four. The last point is important in our case since problems generally occur at the vicinity of those vertices, when analyzing continuity between B-spline surface patches. Several techniques exist to convert a triangulated surface into a quadrilateral mesh (see [18] for a survey on mesh generation). But using these methods which focus on the geometry of the surface, the number of quadrangles often remains high. Theoretical results also exist on quadrangulations of surfaces with a minimum number of quadrangles [10]. As far as we are concerned, we want both a small number of quadrangles and a method to construct them on the surface.

In this paper we define topological quadrangulations on connected compact surfaces and construct them on triangulated surfaces. The combinatorial relationships between quadrangles will only depend on the topology of the surface, but the geometrical properties (such as the length of the edges or the location of the vertices) will depend on its geometry. In other words, the topological quadrangulation seen as a combinatorial mesh will only depend on the genus of the surface (two homeomorphic surfaces, e.g., a sphere and a cube, will have the same topological quadrangulation) whereas its embedding in $\mathbb{R}^3$ will depend on the geometry of the surface.

Our definition minimizes the number of vertices whose valence is different from four in a topological quadrangulation. We could have chosen instead to minimize the maximum valence of a vertex; the first strategy comes to concentrate problems on a few places whereas the last method comes to distribute problems all over the surface. Moreover, the number of quadrangles in our quadrangulation will only depend on the genus of the surface and will be small.

This paper is organized as follows. In Section 2 we recall some results about the topology of surfaces and we define a topological quadrangulation of a connected, compact surface as a 2D-quadrangulation of its canonical polygonal schema. Then in Section 3 we define a Reeb graph, based on a distance function to a source point, and embedded on the triangulated surface. We prove that this Reeb graph is linked to the location of the holes in the surface. It helps us to detect an embedding of the canonical generators on the triangulated surface, from which we deduce the topological quadrangulation of the triangulated surface. Finally in Section 4 we show some results on several triangulated surfaces. We conclude and discuss future work in Section 5.

2. Topological quadrangulation of a surface

We call *surface* a connected, compact 2-manifold embedded in $\mathbb{R}^3$. This definition implies that a surface is orientable, closed (i.e., without boundary).

2.1. Topological classification of surfaces

Two surfaces are topologically equivalent if and only if they are homeomorphic. We recall here the classification theorem for surfaces and some definitions. A good introduction to combinatorial topology and a detailed proof of the classification theorem for surfaces can be found in [8].

Let M be an orientable, connected, compact 2-manifold. There exists an integer $g \geqslant 0$ such that M is homeomorphic to a sphere with g handles. g is called the *genus* of M.

M can be represented canonically using a $4g$-gon G, whose successive edges are marked with letters $a_1, b_1, a_1^{-1}, b_1^{-1}, \ldots, a_g, b_g, a_g^{-1}, b_g^{-1}$, and such that for all $1 \leqslant i \leqslant g$, edges a_i and b_i are oriented counterclockwise and edges a_i^{-1} and b_i^{-1} are oriented clockwise. G is named the *canonical polygonal schema* or *canonical fundamental polygon* of M. If we identify edges a_i and a_i^{-1}, and edges b_i and b_i^{-1} of G according to their orientation, $\forall 1 \leqslant i \leqslant g$, we obtain a space homeomorphic to M. Also, all vertices of G become a single point under this identification. After identification, the edges of G form $2g$ curves on M, which are called the *canonical generators* of M.

In 1998, Dey and Guha [5] analyzed the generators of the first homology group of compact connected 3-manifold in $\mathbb{R}^3$ and proved that a canonical set of generators of a surface (considered as the boundary of a 3-manifold) embedded in $\mathbb{R}^3$ is composed of g *longitudinal* generators and g *latitudinal* generators. The first surround the g holes and the last turn around each handle associated with a hole. See Fig. 1 for an example and [5] for more details.

2.2. Pre-quadrangulation of a surface

Definition 1 (*Quadrangle*). Let M be a surface and $\square$ be the closed square of $\mathbb{R}^2$, i.e., the set of points $(x_1, x_2) \in \mathbb{R}^2$ such that $0 \leqslant x_1 \leqslant 1$ and $0 \leqslant x_2 \leqslant 1$. We define a *quadrangle* on M as the image $f(\square)$ of a homeomorphism $f : \square \rightarrow M$. Quadrangles have four edges and four vertices which are the images of the edges and vertices of $\square$.

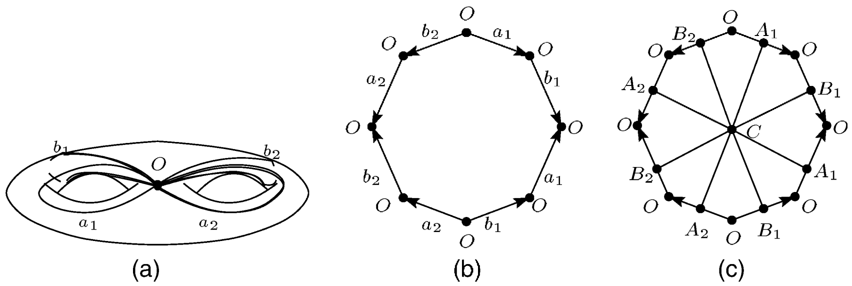

Fig. 1. Canonical set of generators (a), canonical polygonal schema (b), and pre-quadrangulation of a 2-torus (c). a_1 and a_2 are the longitudinal generators, and b_1 and b_2 are the latitudinal ones.

In this paper, we consider special decompositions of surfaces into vertices, edges, and quadrangles that we call *pre-quadrangulations*.

Definition 2 (*Pre-quadrangulation*). A *pre-quadrangulation* Q of a surface M is a decomposition of M into a finite number of vertices, edges, and quadrangles such that:

1. Q is a cell complex, that is to say a collection of cells (a cell is either a vertex or an edge or a quadrangle) such that their interiors are disjoint and the boundary of each cell is made of cells of lower dimensions;
2. any pair of vertices is joined by at most two edges;
3. any two quadrangles of Q are either disjoint or meet in a common vertex or intersect along one common edge or intersect along two non-consecutive common edges.

Our definition includes quadrangulations of surfaces as defined in [10], but is more general: unlike quadrangulations, two vertices in a pre-quadrangulation can be joined by two edges and two quadrangles can intersect along two edges. However, we can still patch a pre-quadrangulation with B-spline surfaces. For example, Fig. 2 shows a pre-quadrangulation of the torus: we have four distinct quadrangles, all having the same four vertices A, B, C and D; vertices A and B are joined by two different edges.

Proposition 3. *Let Q be a pre-quadrangulation of a surface M. Let V be the number of vertices in Q, E the number of edges, F the number of quadrangles and χ_M the Euler characteristic of M: $\chi_M = V - E + F$.*
A lower bound for the number of quadrangles needed to construct Q is

$$F_{\min} = \left\lceil \frac{1}{2}(\sqrt{9-8\chi_M}+3) \right\rceil - \chi_M. \tag{1}$$

Example 4. For a torus ($\chi_M = 0$), we have $F \geqslant 3$. For a 2-torus ($\chi_M = -2$), we have $F \geqslant 6$.

Proof. According to Rule (2) of Definition 2, E is less than or equal to twice the number of pairs of vertices in Q, which is $\frac{V(V-1)}{2}$: $E \leqslant V(V-1)$. Since each edge belongs to two faces, and each face has four distinct edges, $2E = 4F$. We thus have $F = V - \chi_M$ and $F \leqslant \frac{V(V-1)}{2}$. Consequently, $2V - 2\chi_M \leqslant V(V-1)$. This quadratic inequality has the solution: $V \geqslant \lceil \frac{1}{2}(\sqrt{9-8\chi_M}+3) \rceil$. Finally, $F = V - \chi_M \geqslant \lceil \frac{1}{2}(\sqrt{9-8\chi_M}+3) \rceil - \chi_M$. □

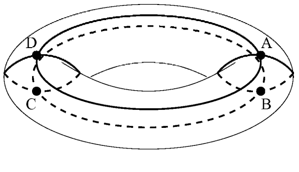

Fig. 2. Pre-quadrangulation of a torus.

We do not know if a general formula gives the minimum number of quadrangles in a pre-quadrangulation of a surface. In this paper we will construct a pre-quadrangulation of a surface with $4g$ quadrangles and $2g+2$ vertices.

Since a pre-quadrangulation is a cell complex and a surface is orientable, we also have the Euler relation: $\chi_M = 2(1-g)$.

A vertex in a pre-quadrangulation is said to be *ordinary* if its valence (or degree) is 4. Vertices with a valence different from 4 are said to be *extraordinary*. We have the following remarkable relation:

Proposition 5. *Let Q be a pre-quadrangulation of M, g the genus of M, V_{e} the number of extraordinary vertices in Q, and $v_1, \ldots, v_{V_{\mathrm{e}}}$ the valences of these extraordinary vertices. We have*:

$$\sum_{i=1}^{V_{\mathrm{e}}}(v_i - 4) = 8(g-1). \tag{2}$$

Proof. We rank the vertices of the pre-quadrangulation from 1 to V starting with the extraordinary vertices. $\forall 1 \leqslant i \leqslant V$, let v_i be the valence for the ith vertex of the pre-quadrangulation: $\forall V_{\mathrm{e}} + 1 \leqslant i \leqslant V, v_i = 4$. Since M is closed, v_i is also the number of faces incident in the ith vertex. Since each face has four distinct vertices, we have $\sum_{i=1}^{V} v_i = 4F$.

We can split the sum in two, separating the extraordinary vertices from the ordinary ones: $\sum_{i=1}^{V_{\mathrm{e}}} v_i + 4(V - V_{\mathrm{e}}) = 4F$. This leads to: $4F = \sum_{i=1}^{V_{\mathrm{e}}}(v_i - 4) + 4V$. We have the Euler relation $\chi_M = V - E + F = 2(1-g)$. Since each edge belongs to two faces and each face has four distinct edges, $2E = 4F$. Thus $8(1-g) = 4(V - E + F) = 4(V - F)$.

Finally, $8(g-1) = \sum_{i=1}^{V_{\mathrm{e}}}(v_i - 4)$. □

Proposition 5 shows that the sum of the orders (that is to say, the valences minus four) of all extraordinary vertices is a surface invariant. This also shows that, except for surfaces with genus equal to 1, any pre-quadrangulation of a surface contains extraordinary vertices.

2.3. Quadrangulation of the canonical polygonal schema

We now define a pre-quadrangulation of a surface M with genus $g > 0$ using the notions introduced above. Our pre-quadrangulation has $4 - 2\chi_M = 4g$ quadrangles and $4 - \chi_M = 2g + 2$ vertices. It is the embedding of a quadrangulation of the canonical polygonal schema on the surface.

Let G be the canonical polygonal schema of M and $a_1, b_1, a_1^{-1}, b_1^{-1}, \ldots, a_g, b_g, a_g^{-1}, b_g^{-1}$ the successive labels of its edges. Let $\mathcal{O}_1, \ldots, \mathcal{O}_{4g}$ be the vertices of G. We choose a vertex $\mathcal{A}_i$ (resp. $\mathcal{B}_i$) on each edge a_i (resp. b_i) of G, and a vertex $\mathcal{A}_{i+g}$ (resp. $\mathcal{B}_{i+g}$) on each edge a_i^{-1} (resp. b_i^{-1}) of G. We also choose a vertex C

inside the polygon. We join C and each $\mathcal{A}_i$ (resp. $\mathcal{B}_i$), $1 \leqslant i \leqslant 2g$, together to create an edge.The quadrangulation Q of G is defined as follows:

- the vertices of Q will be $\mathcal{O}_1, \ldots, \mathcal{O}_{4g}, \mathcal{A}_1, \ldots, \mathcal{A}_{2g}, \mathcal{B}_1, \ldots, \mathcal{B}_{2g}, C$;
- the edges of Q will be the edges of G and the edges created between C and each $\mathcal{A}_i$ or $\mathcal{B}_i$;
- the quadrangles of Q will be the quadrangles defined by these vertices and edges inside the polygon.

This quadrangulation of G has $8g + 1$ vertices, $12g$ edges (one edge of G correspond to two edges of Q, between some $\mathcal{A}_i$ or $\mathcal{B}_i$ and two $\mathcal{O}_j$), and $4g$ quadrangles.

It corresponds to a pre-quadrangulation of M with $2g + 2$ vertices, $8g$ edges and $4g$ quadrangles, since under identification of the edges of G, all vertices $\mathcal{O}_i$ become a single point O, and for all $1 \leqslant i \leqslant g$, $\mathcal{A}_i$ and $\mathcal{A}_{i+g}$ (resp. $\mathcal{B}_i$ and $\mathcal{B}_{i+g}$) become a single point A_i (resp. B_i).

An example is shown on Fig. 1c, or Fig. 3 (left).

The number of quadrangles, $4g$, is close to the lower bound of Proposition 3. The only extraordinary vertices of the pre-quadrangulation are O and C, which are connected to each A_i and each B_i twice: their valence is equal to $4g$ (if $g = 1$, our pre-quadrangulation does not contain any extraordinary vertex). We have chosen to minimize the number of extraordinary vertices (with respect to the number of vertices), but it implies in return the valence of these extraordinary vertices is high, according to Proposition 5.

Let us observe that the two vertices O and C play the same role in the pre-quadrangulation. Indeed, the quadrangulation of the canonical polygonal schema with O or with C as the base-point leads to the same pre-quadrangulation of M. We can go from one quadrangulation to the other simply by mentally cutting each quadrangle along its edges and pasting the quadrangles together, according to the orientation of the edges. See Fig. 3 for an example on a 2-torus.

We will call this particular pre-quadrangulation of the surface based on a quadrangulation of its canonical polygonal schema a *topological quadrangulation* of the surface.

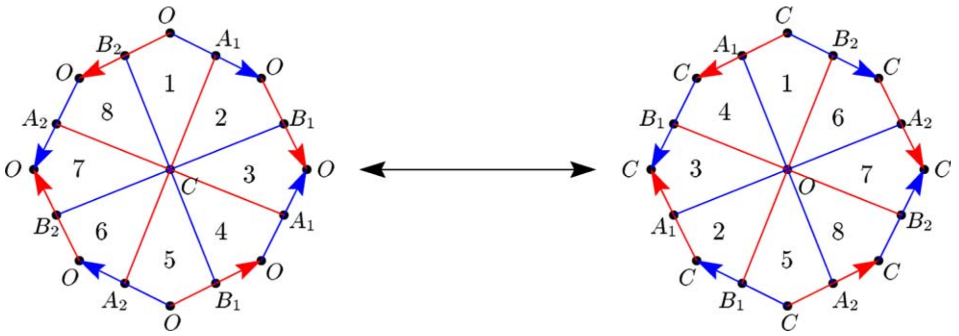

Fig. 3. Two different representations of the same pre-quadrangulation of a double torus. Quadrangles are numbered from 1 to 8.

3. Construction of a topological quadrangulation on a triangulated surface

This section is devoted to the computation of the topological quadrangulation on a triangulated surface, deduced from the construction of the canonical polygonal schema on the surface.

3.1. Reeb graph

To construct a canonical set of surface generators on a triangulated surface, we use a topological structure for compact manifolds called the Reeb graph. The Reeb graph, named after Reeb [20], is a powerful tool to describe topological features on surfaces. The Reeb graph combined with the Morse theory is increasingly used in computer graphics, mainly for terrain analysis and shape modeling [2,13,21,22]. See [8] for an introduction about Morse theory and the Reeb graph.
A Reeb graph is defined as follows:

Definition 6. Let $f : M \to \mathbb{R}$ be a real-valued continuous function on a compact manifold M. The *Reeb graph of* f is the quotient space of M by the equivalence relation $\sim$ defined by:

$x_1 \sim x_2 \iff f(x_1) = f(x_2)$ and x_1 and x_2 are in the same connected component of $f^{-1}(f(x_1))$.

Fig. 4 shows a Reeb graph of a height function on a torus (figure inspired by [8]).

Nodes correspond to splits or merges of connected components of $x \mapsto f^{-1}(f(x))$. There exists two kinds of nodes in a Reeb graph:

Definition 7. Nodes incident to only one edge will be called *extremal nodes*. They represent local extrema of the function f.

Nodes incident to at least three edges will be called *internal nodes*. They represent saddle points of the function f.

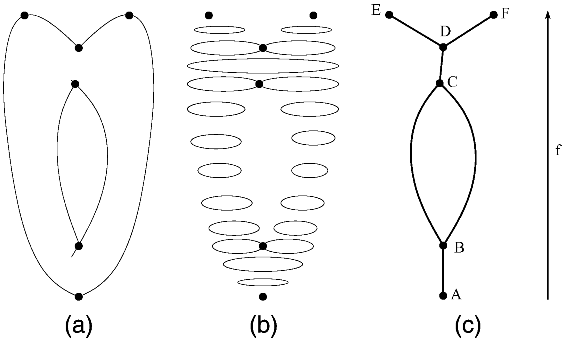

Fig. 4. A torus and the critical points of its height function f (a), connected components of some level sets for f (b), and the Reeb graph of f on this torus (c).

For example, on Fig. 4, nodes A, E, and F are extremal nodes and nodes B, C, and D are internal nodes.

Points at which the topology of the level sets $x \mapsto f^{-1}(f(x))$ change are called *critical points* of f.

3.2. Cycles in the Reeb graph and canonical generators

From now on, we consider only triangulated surfaces. The Reeb graph we construct will be useful for our purpose because its cycles will help us to locate the generators on the surface. The following theorem is related to works by Wood et al. [24], but is more general.

Theorem 8. *Let M be a triangulated surface in $\mathbb{R}^3$, with genus g. We assume f is linear inside each triangle of M. We also assume f is general for M, that is to say $f(v) \neq f(w)$ for any edge (v, w) of the triangulated surface M, and all singularities of f are simple (i.e., nodes of the Reeb graph are incident to exactly three edges).*

Then the number of independent cycles in the Reeb graph of f on M is equal to the genus g of M, and each cycle corresponds to exactly one hole in M.

Proof. Since f is general for M and linear between its vertices, all points of M which are not vertices of the triangulation must be regular points (i.e., not critical points).

We recall here that the Euler characteristic of the union of two polyhedra is the sum of the Euler characteristics of these polyhedra minus the Euler characteristic of their intersection. Thus if these polyhedra intersect along a disjoint union of closed polygonal lines, the Euler characteristic of their union is the sum of the Euler characteristics of the two polyhedra, since the Euler characteristic of a closed polygonal line is equal to 0.

Let $\alpha_0 < \alpha_1 < \cdots < \alpha_{n-1}$ be the critical values of f (that is to say, the values corresponding to the nodes of the Reeb graph). We will sweep the surface starting from s. $\forall \alpha_0 < \alpha \leqslant \alpha_{n-1}$, we define $M_\alpha = \{x \in M, f(x) \leqslant \alpha\}$, χ_α to be the Euler characteristic of M_α and g_α its genus.

Since M_α is an orientable connected 2-manifold with boundary, we have

$$\chi_\alpha = 2 - 2g_\alpha - b_\alpha, \tag{3}$$

where b_α is the number of boundary components of M_α.

Let C_α be the boundary of M_α: $C_\alpha = f^{-1}(\alpha)$. If α is not a critical value, we can observe that each connected component of C_α is a closed polygonal line, since f is general and linear on each edge and inside each triangle of M. If α is a critical value, there exists one connected component of C_α which is the union of two single closed polygonal lines meeting at a common (saddle) point (only one connected component since f is general, and only two polygonal lines since α is a simple singularity), see Fig. 5.

For $\alpha_0 \leqslant \alpha < \alpha_1$, C_α is a single closed polygonal line. Filling C_α with a disk adds 1 to the Euler characteristic of M_α, since the Euler characteristic of a (closed) disk is

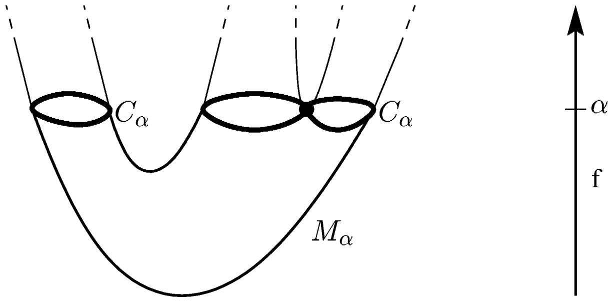

Fig. 5. The surface M_α and its boundary C_α, where α is a critical value of f.

equal to 1. It also leads to a surface homeomorphic to a sphere. Since the Euler characteristic of a sphere is equal to 2, $\chi_\alpha = 2 - 1 = 1$.

Let us now see what happens when α increases past a critical value α_i of f. Let $\varepsilon < \min_{1 \leqslant j \leqslant n-1}(\alpha_j - \alpha_{j-1})$ and $N_\varepsilon = \{x \in M, |f(x) - \alpha_i| \leqslant \varepsilon\}$. Let P_i be the associated critical point. P_i is a vertex of M. We have the following possible cases:

- P_i corresponds to a local maximum of f. We meet the end of an edge in the Reeb graph. A connected component of C_α is closing (see Fig. 6a), and $\chi_{\alpha_i} = \chi_{\alpha_{i-1}} + 1$.
- P_i is a saddle point of f in which one connected component of C_α split in two connected components (see Fig. 6b). In the Reeb graph, one edge splits in two edges.

We have $M_{\alpha_i+\varepsilon} = M_{\alpha_i-\varepsilon} \cup N_\varepsilon$, and $M_{\alpha_i-\varepsilon} \cap N_\varepsilon$ is a disjoint union of closed polygonal lines. Thus $\chi_{\alpha_i+\varepsilon} = \chi_{\alpha_i-\varepsilon} + \chi_{N_\varepsilon}$. Since α_i is a simple singularity, N_ε has exactly three boundary components, which are closed polygonal lines, and if we fill each of these boundary components with a disk we obtain a surface homeomorphic to a sphere (see Fig. 7). Thus $\chi_{N_\varepsilon} = 2 - 3 = -1$, and $\chi_{\alpha_i+\varepsilon} = \chi_{\alpha_i-\varepsilon} - 1$.

Since one connected component of C_α split in two connected components, we have $b_{\alpha_i+\varepsilon} = b_{\alpha_i-\varepsilon} + 1$. Eq. (3) leads to $g_{\alpha_i+\varepsilon} = g_{\alpha_i-\varepsilon}$.

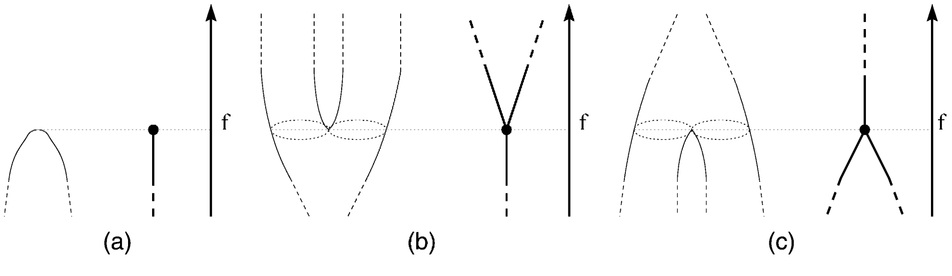

Fig. 6. Sweep of M with respect to f. The closing of a connected component of M corresponds to an external node in the Reeb graph (a); one connected component becoming two connected components corresponds to an "opening" internal node in the Reeb graph (b); two connected components becoming one connected component corresponds to a "closing" internal node in the Reeb graph (c).

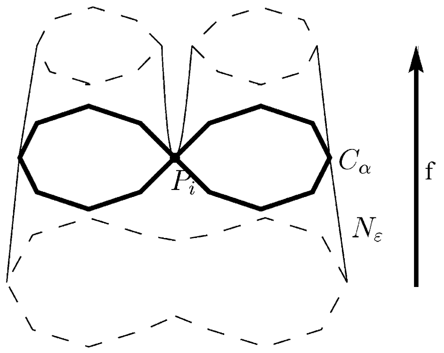

Fig. 7. Connected component of C_α, where α is a critical value of f (bold), and the set N_ε (between the dotted polylines).

- P_i is a saddle point of f in which two connected components of C_α merge in one connected component (see Fig. 6c). In the Reeb graph, two edges merge in one edge, and the number of independent cycles increases by 1.

  For the same reason as in the previous case, $\chi_{\alpha_i+\varepsilon} = \chi_{\alpha_i-\varepsilon} - 1$.

  Since two connected components of C_α merge in one connected component, we have $b_{\alpha_i+\varepsilon} = b_{\alpha_i-\varepsilon} - 1$. Eq. (3) leads to $g_{\alpha_i+\varepsilon} = g_{\alpha_i-\varepsilon} + 1$.

The number of independent cycles in the Reeb graph and the genus of the surface M_α increase at the same time: *a cycle in the Reeb graph corresponds to a hole in M.* Now let a be the number of local maxima of f, b the number of saddle points where one connected component of the boundary of M_α becomes two connected components, and c the number of saddle points where two connected components of the boundary of M_α become one connected component. When we sweep the Reeb graph, we meet $1 + 2b + c$ beginnings of edges, and $a + b + 2c$ ends of edges. Consequently,

$$1 + 2b + c = a + b + 2c, \tag{4}$$

that is to say

$$a - b = 1 - c. \tag{5}$$

Moreover, the Euler characteristic χ_M of M is equal to $\chi_{\alpha_{n-1}}$. We thus have

$$\chi_M = \chi_{\alpha_0} + a - b - c. \tag{6}$$

Since $\chi_{\alpha_0} = 1$ and following from Eqs. (5) and (6),

$$\chi_M = 2 - 2c. \tag{7}$$

Since M is orientable, $\chi_M = 2(1 - g)$, and we can deduce from Eq. (7) that

$$g = c. \tag{8}$$

But the number of cycles in the Reeb graph is also c. Consequently, the number of cycles in the Reeb graph is equal to the genus of the surface. Since for each cycle we have found a corresponding hole in the surface, we can claim that each cycle corresponds to exactly one hole in M. □

3.3. Construction of a Reeb graph embedded on a triangulated surface

3.3.1. Choice of the function f

To define our Reeb graph, we must choose a function f. As said before, the Reeb graph is mainly used in terrain analysis. In this case f is often a height function, because we have a given orientation where the height is meaningful. In the case of surfaces, we do not have such a particular orientation, and people often choose f as a distance function to a given source point on the surface [1,16,24]. Another possibility is to integrate this function over the surface, in order to avoid the choice of the source point [12]. Since we are only interested in the topology of the surface and not the exact location of the critical points of f, we do not need to choose a very subtle (but computationally more expensive) function f. That is why we choose f as *the shortest path distance to a given source point on the surface*:

let s be a vertex on M (the *source point*); for each vertex v of M, we define $f(v)$ as the shortest edge path distance from s to v and we extend f linearly to a continuous function over all points of M.

The distance between two neighbouring vertices is the length of the edge between these two vertices, whereas in [1] it is equal to 1. If f is not general for M, we can turn f into a general function simply by a small perturbation of the vertices, as described in [7]. This method has been used in a similar problem by [3].

In practice multi-saddle points can exist. Multi-saddle points are points where several splits or merges of connected components of $x \mapsto f^{-1}(f(x))$ can occur. In order to break them into multiple simple singularities, we refer to Carr et al. algorithm [3].

The choice of the source point is not critical for our topological purpose, but the number of edges in the Reeb graph, therefore the computational time of our algorithm, depends on its location. We will choose the source point s as the furthest vertex from an arbitrary vertex [16]: in practice s will be at the end of some long branch of the surface; consequently the direction induced by f is geometrically meaningful. For our sample surfaces, we observed that the number of local extrema for f is usually smaller than with other source points, inducing in return the amount of edges in the Reeb graph to remain small.

3.3.2. Algorithm

To avoid confusion between the nodes (resp. edges) of the Reeb graph and the vertices (resp. edges) of the mesh, the first will be called Reeb nodes (resp. Reeb edges) and the last vertices (resp. edges).

To construct the Reeb graph of f on the surface, we must first compute the value of f at all the vertices of the mesh, then detect the critical points of f and finally link them to create the Reeb edges. Actually, we will do the three steps at the same time, using Dijkstra's algorithm. Dijkstra's algorithm gives at each step and for each vertex X a temporary value $v(X)$ and a predecessor vertex.

With each vertex X of the triangulation we associate its *value* $v(X)$ (at the end of the algorithm $v(X)$ will be equal to the value of f at X), its *predecessor* according to Dijkstra, and a list of *labels*. Labels will be used to identify Reeb edges: two vertices

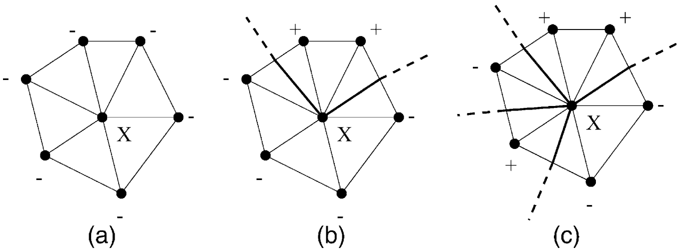

Fig. 8. Sign changes $n(X)$ for $f - f(X)$ around a vertex X of the triangulated surface. $n(X) = 0$ (a); $n(X) = 2$ (b); $n(X) = 4$ (c).

X and Y will have the same label if and only if the segment $[f(X), f(Y)]$ contains no critical values of f and X and Y belong to the same connected component of $f^{-1}([f(X), f(Y)])$. At the beginning, the value of each vertex is $+\infty$ and vertices have no label, except the source point whose value is 0 and which has one label.

At each step, the currently processed vertex is the vertex X with minimal value: we know that $f(X) = v(X)$. Values of all neighbouring vertices of X are updated using Dijkstra's algorithm.

We then compute the number of sign change $n(X)$ for $v - v(X)$ when we visit all neighbouring vertices. Let Y be a neighbouring vertex of X. Since we use Dijkstra's algorithm, if $v(Y) < v(X)$, we also have $f(Y) < f(X)$. Thus $n(X)$ is also the number of sign change for $f - f(X)$.

If $n(X) = 0$, X is a local extremum for f, that is to say an extremal Reeb node. Its only label will be equal to the (first) label of its predecessor. If $n(X) = 2$, X is not a Reeb node. Its only label will also be equal to the (first) label of its predecessor. Finally if $n(X) \geqslant 4$, X is an internal Reeb node. We compute its local level set and its connected components (see below). Each vertex of each upper connected component will be given a new label (the same for all vertices in the same connected component). X will be given all labels of its upper and lower connected components. See Fig. 8 for some examples: (a), X is a local maximum for f; (b) X is not a Reeb node; (c) X is an internal Reeb node.

A Reeb edge between vertices A and B with $f(A) < f(B)$ will consist of two parts, as depicted in Fig. 9. The first part is a path made from B to its successive predecessors through Dijkstra's algorithm until A or a vertex with the corresponding label is found. The second part is the path made from this vertex to A by following a connected component of A (if A is not the source point, A is an internal node of the graph, so its local level set has been computed).

3.3.3. Local level sets

The *local level set* $LS(X)$ *of* X is a set of pairs of vertices (Y, Z), where $f(Y) > f(X)$ and $f(Z) < f(X)$ (see Fig. 10). An *upper* (resp. *lower*) *connected component* of X will be a connected component of the set of all Y (resp. Z) in $LS(X)$ such that $f(Y) > f(X)$ (resp. $f(Z) < f(X)$). Note that a connected component of X will be a set of vertices, and not a set of pairs of vertices.

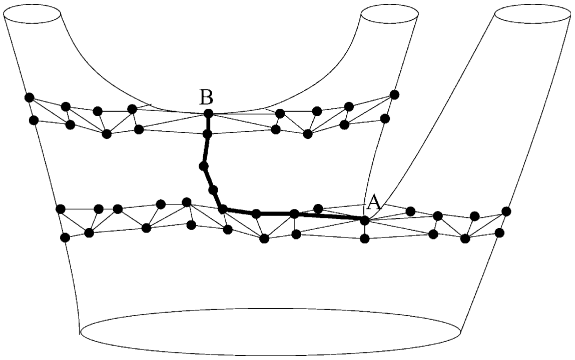

Fig. 9. An edge between two Reeb nodes A and B (bold) and the corresponding local level sets.

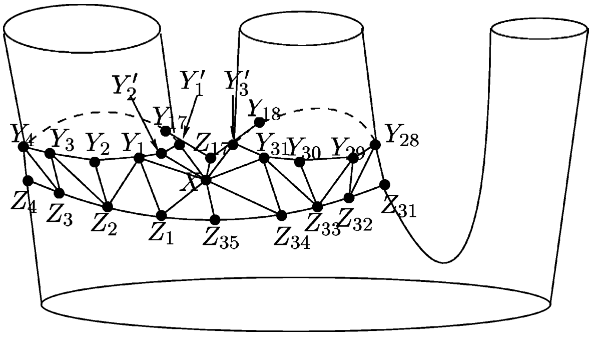

Fig. 10. Local level set of an internal node X and its connected components. Here we have two upper connected components $(Y_1, \ldots, Y_{17} Y_1' Y_2'$ and $Y_{18}, \ldots, Y_{31} Y_3')$ and one lower connected component $(Z_1, \ldots, Z_{17}, \ldots, Z_{35})$.

To compute $LS(X)$ and the connected components of X, we use a very simple algorithm. Basically, we start from X and we follow the level set $f^{-1}(f(X))$. We enumerate the edges intersected by this level set and put them in a list.

```
$LS := 0$
Repeat
  Choose two adjacent vertices $Y$ and $Z$ in the neighbourhood of $X$ such that
  $f(Y) > f(X)$ and $f(Z) < f(X)$
  Add $(Y, Z)$ to $LS$
  $avoid\_vertex := X$
  $new\_vertex := Y$
  While $new\_vertex$ is different from $X$ do:
    $new\_vertex$ :=the vertex adjacent to $Y$ and $Z$ which is not $avoid\_vertex$
    If $f(new\_vertex) < f(X)$ then
      Add $(Y, new\_vertex)$ to $LS$
      $avoid\_vertex := Z$
      $Z := new\_vertex$
    Elsif $f(new\_vertex) > f(X)$ then
      Add $(new\_vertex, Z)$ to $LS$
      $avoid\_vertex := Y$
      $Y := new\_vertex$
    End if.
  End while.
Until it is no more possible to choose such two adjacent vertices.
```

At the end of each *while* loop, we have two new connected components of X, which are not complete. To complete them and possibly merge them, we compute the vertex NV adjacent to X and Y (where Y is the last vertex of a connected component) which is not the previous vertex of the connected component. We add it to the connected component, replace Y by this new vertex and do the same until NV is in $LS(X)$. If the last vertex NV (which is in $LS(X)$) is also in the connected component, this connected component is complete. If not, we must merge it with the connected component which contains NV.

3.4. Construction of the canonical generators embedded on a triangulated surface

Only a few algorithms have been proposed to construct generators. Dey and Schipper [4] compute a polygonal schema of a triangulated surface in $O(n)$ time; unfortunately their polygonal schema is not necessarily canonical. Vegter and Yap's algorithm [23] runs in $O(g \cdot n)$ time and space (which is optimal) and construct a canonical set of generators. Lazarus et al. [17] simplified this algorithm and compared it to another algorithm based on Brahana's method.

Providing that the genus of the surface is strictly positive, finding a cycle basis $\mathcal{B}$ on our constructed Reeb graph will give us information about the location of surface holes, and thus about surface generators.

To find $\mathcal{B}$, we use a traditional algorithm from graph theory: we find a maximal forest (a forest is a subgraph with no cycle) on the Reeb graph, then for each edge which is not in this forest we find a minimal cycle containing this edge. The set of all minimal cycles found is $\mathcal{B}$. Such algorithms are detailed in graph theory books, e.g., [9].

Once we have found $\mathcal{B}$ we can construct a canonical set of surface generators: for each cycle $\mathcal{C}$ in $\mathcal{B}$ there exists one longitudinal generator and one latitudinal generator. $\mathcal{C}$ is the longitudinal generator. To construct the latitudinal generator, we find the cycle vertex V with the minimal value, and then follow an upper connected component of V (we have computed the local level set of V when we have constructed the Reeb graph).

3.5. Construction of the topological quadrangulation

Now that we have computed a canonical set of generators on the surface, the construction of our pre-quadrangulation will be easy if we choose judiciously its vertices O, A_k, B_k, and C (cf. Section 2.3 for the notations).

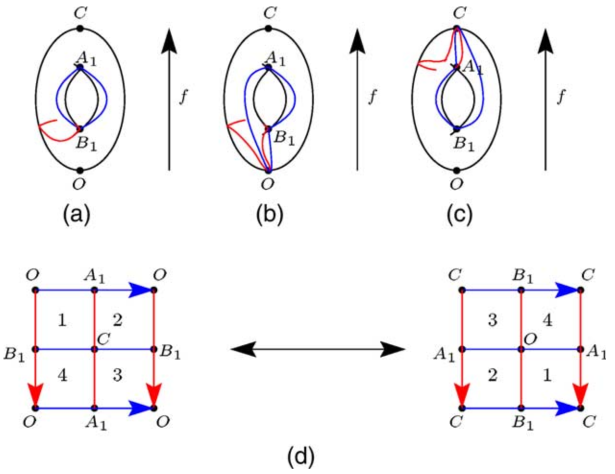

Fig. 11. Generators of a torus (a), edges OA_i and OB_i of the topological quadrangulation (b), edges CA_i and CB_i of the topological quadrangulation (c). The bottom figure (d) shows the quadrangulation "centered" either at C or at O (quadrangles are numbered from 1 to 4).

We take O equals to the source point and C equals to the vertex for which f is maximum. Then for each k we define B_k as the internal node of the Reeb graph with the lowest value and belonging to the kth latitudinal cycle of the basis, and A_k as the internal node with the greatest value and belonging to the kth longitudinal cycle of the basis.

To construct the edges OA_k and OB_k of the quadrangulation, we continuously distort each generator of the constructed canonical set of generators so that O belongs to it (this can be done in time $O(g \cdot n)$, n being the number of vertices in the initial triangulation, since Dijkstra's algorithm computes for each vertex a path to the source point). Then we follow the distorted generators: the longitudinal ones for the edges OA_k and the latitudinal ones for the edges OB_k. Since O and C play the same role in the pre-quadrangulation (see Section 2.3), we can do the same to construct the edges CA_k and CB_k: edges CA_k correspond to the latitudinal generators and edges CB_k correspond to the longitudinal generators. See Fig. 11 for an example.

4. Results and discussion

We have implemented our algorithm in C++; the user interface was developed using the GLUT and Mesa libraries. The user can choose to visualize either the vertices of the triangulated surface which correspond to the Reeb nodes, or the Reeb graph embedded on the triangulated surface, or a canonical set of generators embedded on the triangulated surface, or the edges of the topological quadrangulation as computed by the algorithm. Some results are shown on Fig. 12.

Top figures show the Reeb graph computed on three surfaces: a helix, which is homeomorphic to a sphere, a surface with one hole, and a rocker arm (courtesy

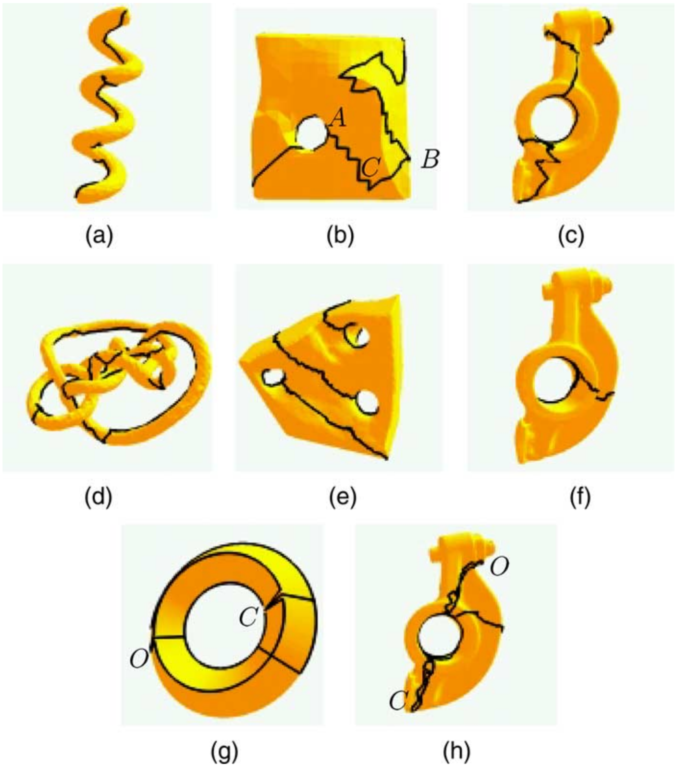

Fig. 12. Reeb graph (a,b,c), canonical set of generators (d,e,f), and topological quadrangulation (g,h) of some surfaces.

of Cyberware, http://www.cyberware.com), also homeomorphic to a torus. The Reeb graph for the helix is a tree, and we found six Reeb nodes: one global minimum, one global maximum, two internal Reeb nodes and two local maxima. These two maxima are close to the internal nodes, and can be seen on the middle of the helix. A local maximum has also been detected on the bottom right of the one-hole surface, and on the left of the rocker arm. We see, particularly on the one-hole surface, that a Reeb edge between two nodes A and B (with $f(A) < f(B)$) follows two different directions: starting from B, it follows first the gradient of f on the surface until it reaches a vertex C which belongs to the local level set of A, then it follows the shortest edge path from C to A.

Middle figures show a canonical set of generators computed on a knot-shaped torus, a surface with genus 3, and the rocker arm. The latitudinal generator of the knot is in the bottom left of the surface. All latitudinal generators embedded on a surface are almost parallel, and orthogonal to the gradient of f, since they follow local level sets. Consequently, they do not necessarily correspond to locally shortest closed curves.

Finally, bottom figures show the coarse topological quadrangulations of a simple torus and the rocker arm constructed with our algorithm (only edges of the quadrangulations are visualized). Edges starting from the vertex O (or C), using notations from the Section 2.3, are sometimes close to each other.

The computation time of our algorithm to construct the Reeb graph and the generators on the surface is $O(n^2)$. Time used by Dijkstra's algorithm to construct the Reeb nodes is $O(n.\log(n))$, but then to construct the Reeb edges we need to compute the local level sets of Reeb nodes. This takes $O(n^2)$ time. An open question would be to know if this computational time can be improved.

5. Conclusion and future work

We have presented an algorithm to construct a new discrete representation of a triangulated surface with strictly positive genus by a coarse topological quadrangulation, based on a quadrangulation of the canonical polygonal schema of the surface. Our method first constructs a Reeb graph embedded on the surface, using Dijkstra's algorithm, then deduces the generators of this surface.

Our work has mainly focused on combinatorial properties of the topological quadrangulation. Further work includes quadrangulation refinement to take surface geometrical singularities (to be defined) into account, and optimization (in terms of angles in quadrangles, extraordinary vertices, ...). The choice of the function used to define the Reeb graph can also be improved to fulfill these goals.

Acknowledgments

The authors would like to thank Francis Lazarus for fruitful discussions and helpful comments, especially concerning the number of cycles in the Reeb graph.

References

[1] U. Axen, H. Edelsbrunner, in: Auditory Morse Analysis of Triangulated Manifolds. Mathematical Visualization, Springer-Verlag, Berlin, 1998, pp. 223–236.

[2] S. Biasotti, B. Falcidieno, M. Spagnuolo, Extended Reeb Graphs for Surface Understanding and Description, in: Proceedings of DGCI'00, Lecture Notes in Computer Science, vol. 1953, 2000, pp. 185–197.

[3] H. Carr, J. Snoeyink, U. Axen, Computing Contour Trees in All Dimensions, in: Proceedings of ACM 11th Symposium on Discrete Algorithms, San Francisco, California, USA, Jan. 2000, pp. 918–926.

[4] T. Dey, H. Schipper, A new technique to compute polygonal schema for 2-manifolds with application to null-homotopy detection, Discrete Comput. Geom. 14 (1995) 93–110.

[5] T. Dey, S. Guha, Computing homology groups of simplicial complexes in $\mathbb{R}^3$, J. ACM 45 (2) (1998) 266–287.

[6] M. Eck, H. Hoppe, Automatic Reconstruction of B-Spline Surfaces Of Arbitrary Topological Type, in: Proceedings of SIGGRAPH'96, August 1996, pp. 325–334.

[7] H. Edelsbrunner, E.P. Mücke, Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms, ACM Trans. Graph. 9 (1) (1990) 66–104.

[8] A.T. Fomenko, T.L. Kunii, Topological Modeling for Visualization, Springer, Berlin, 1997.

[9] M. Gondran, M. Minoux, Graphs and Algorithms, Wiley, NewYork, 1995.

[10] N. Hartsfield, G. Ringel, Minimal quadrangulations of orientable surfaces, J. Comb. Theory, Series B 46 (1) (1989) 84–95.

[11] P.S. Heckbert, M. Garland, Survey of Polygonal Surface Simplification Algorithms, Multiresolution Surface Modeling Course, SIGGRAPH'97, 1997.

[12] M. Hilaga, Y. Shinigawa, T. Kohmura, T.L. Kunii, Topology Matching for Fully Automatic Similarity Estimation of 3D Shapes, in: Proceedings of SIGGRAPH'01, August 2001.

[13] M. van Kreveld, R. van Ostrum, C. Bajaj, V. Pascucci, D. Schikore, Contour Trees and Small Seed Sets for Isosurface Traversal, in: Proceedings of ACM 13th Symposium on Computational Geometry, Nice, France, June 1997, pp. 212–220.

[14] V. Krishnamurthy, M. Levoy, Fitting Smooth Surfaces to Dense Polygon Meshes, in: Proceedings of SIGGRAPH'96, August 1996, pp. 313–324.

[15] M.J. Lai, L. Schumaker, Scattered data interpolation using C^2 supersplines of degree six, SIAM J. Numer. Anal. 34 (1997) 905–921.

[16] F. Lazarus, A. Verroust, Level Set Diagrams of Polyhedral Objects, in: Proceedings of ACM 5th Symposium on Solid Modeling and Applications, Ann Arbor, Michigan, USA, June 1999, pp. 130–140.

[17] F. Lazarus, M. Pocchiola, G. Vegter, A. Verroust, Computing a Canonical Polygonal Schema of an Orientable Triangulated Surface, in: Proceedings of ACM 17th Symposium on Computational Geometry, Tufts University, Medford, USA, June 2001, pp. 80–89.

[18] S. Owen, A Survey of Unstructured Mesh Generation Technology, in: Proceedings of the 7th International Meshing Roundtable, Sandia National labs, Dearborn, USA, Michigan, October 1998, pp. 239–267.

[19] J. Peters, C^1 Surface splines, SIAM J. Numer. Anal. 32-2 (April) (1995) 645–666.

[20] G. Reeb, Sur les Points Singuliers d'une Forme de Pfaff Complètement Intégrable ou d'une Fonction Numérique, CR Acad. Sci., Paris, France 222 (1946) 847–849.

[21] Y. Shinagawa, T.L. Kunii, Surface coding based on Morse theory, IEEE Comput. Graph. Appl. (September) (1991) 66–78.

[22] S. Takahashi, Y. Shinagawa, T.L. Kunii, A Feature-based Approach for Smooth Surfaces, in: Proceedings of ACM 4th Symposium on Solid Modeling and Applications, Atlanta, Georgia, USA, May 1997, pp. 97–110.

[23] G. Vegter, C.K. Yap, Computational Complexity of Combinatorial Surfaces, in: Proceedings of ACM 6th Symposium on Computational Geometry, Berkeley, California, USA, June 1990, pp. 102–111.

[24] Z. Wood, M. Desbrun, P. Schröder, D. Breen, Semi-Regular Mesh Extraction from Volumes, in: Proceedings of IEEE Visualization 2000, October 2000, pp. 275–282.