Pressing: Smooth Isosurfaces with Flats from Binary Grids

A. Chica, J. Williams[2], C. Andujar, P. Brunet, I. Navazo, J. Rossignac[2], A. Vinacua

Department of Software, Polytechnic University of Catalonia, Barcelona, Spain
[2]IRIS Cluster and GVU Center, College of Computing, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

We explore the automatic recovery of solids from their binary volumetric discretizations. In particular, we propose an approach, called Pressing, for smoothing isosurfaces extracted from binary volumes while recovering their large planar regions (flats). Pressing yields a surface that is guaranteed to contain the samples of the volume classified as interior and exclude those classified as exterior. It uses global optimization to identify flats and constrained bilaplacian smoothing to eliminate sharp features and high-frequencies from the rest of the isosurface. It recovers sharp edges between flat regions and between flat and smooth regions. Hence, the resulting isosurface is usually a very accurate approximation of the original solid. Furthermore, the segmentation of the isosurface into flat and curved faces and the sharp/smooth labelling of their edges may be valuable for shape recognition, simplification, compression, and various reverse engineering and manufacturing applications.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations, voxels, isosurfaces, smoothing, segmentation.

1. Introduction

Binary volumetric models have become more important over the last few years. Binary voxelizations are commonly generated as the result of segmentation algorithms working on volumetric medical data [TSH98] that must assign each voxel to a specific organ. Discrete volume objects are also created by the voxelization of different input models. Many reconstruction [EBV05] and model repair [BPK05] algorithms are based on volume binary models, and a number of volume operations like splitting produce binary information in the modified regions.

We can formally define the volume binary model as follows. Consider a solid model M whose boundary ∂M is smooth and contains large planar faces (called flats). The set of samples of a regular axis-aligned lattice in a box containing M may be divided into the set G of blue samples in M and the set R of red samples out of M. The collection of cubical voxels centered at the blue samples provides a rough approximation of M.

The *cells* of the grid are axis-aligned boxes having for vertices a $2 \times 2 \times 2$ arrangement of neighboring samples from the grid. Cells with vertices of different colors are said to be *mixed*. The axis-aligned edges of the lattice connecting adjacent nodes of different color are called *sticks*. Let the free space F be the union of the mixed cells, and let S be a triangulated isosurface in F that separates R from G and has as its vertices the midpoints of the sticks (Figure 1-a). We explore isotopies in F that will deform S into S' by sliding its vertices along their sticks. In particular, we strive to increase the smoothness of S and at the same time to reproduce in S' close approximations of the flats of ∂M, using only the information encoded by G and R. The resulting surface S' (Figure 1-c) is called the *pressed S*, and the process for computing it is called *Pressing*.

This is a hard problem, parts of which have been addressed by other authors previously [Gib98, Whi00, NGH*03, Nie04], also using only binary in–out data. However none of them solves both aspects of the problem (smoothing and detecting features) in a completely satisfactory way, as discussed in Section 4. In this paper we build on their contributions, and propose new algorithms that yield improved results.

In both two and three dimensions, Pressing starts by grouping the sticks into clusters that can each be stabbed by a flat separating the red and blue samples at the sticks' endpoints. Then Pressing slides the sticks' vertices to their intersections with the flat and *freezes* them. The other *fresh*

Figure 1: *(a) The aliased isosurface was extracted from a 128^3 binary voxelization. Its vertices are at the stick midpoints (as shown by the inset image). (b) The flats were identified and color-coded. The junction points along the boundaries between planar and non-planar regions are also identified and shown as orange dots. Note that a flat may be connected to other flats (yellow square) or to smooth faces through sharp edges (red square) and to smooth faces through smooth edges (magenta square). (c) The resulting pressed isosurface.*

(non-frozen) vertices are then adjusted along their sticks to smooth the isosurface.

We first illustrate this process with a 2D example. A region M with a boundary ∂M that contains several flats (line segments) (Figure 2-a) is rasterized (Figure 2-b) on a regular axis-aligned lattice by painting blue the lattice nodes in M and red the other ones.

Pressing starts from this red/blue labeling and reconstructs an isocontour S (Figure 2-c) approximating ∂M. The mixed cells are the lattice squares having both red and blue vertices. S is contained in their union. A mixed cell is bounded by two or four sticks. A cell with two sticks contributes a single edge to S. A cell with four sticks contributes two edges to S. There are two ways of constructing these two edges so that they do not intersect. The choice affects the topology of the result. We have dealt with this problem in three dimensions in [ABC*05], which is the solution we will adopt here. We will not dwell further in this issue, as the algorithm presented in this paper will work equally with any correct starting triangulation. For the purpose of this illustration, one can assume we have an oracle that decides for us the connectivity to use.

Next, Pressing identifies the flats. It associates each flat with a subset of the sticks it stabs, so that each stick is associated with at most one flat. Then it snaps the vertices of these sticks to the flat that stabs them, by sliding these vertices along their stick (Figure 2-d). These vertices will remain locked in this position (we say that they are *frozen*). The remaining vertices are said to be fresh. Then Pressing identifies *junction* points among the fresh vertices adjacent to frozen vertices; these will correspond to sharp corners. Finally, Pressing perturbs the fresh vertices through a custom smoothing process that retains sharp corners and produces a polygonal curve S' (Figure 2-e) that is isotopic to S (i.e.,

may be continuously deformed into S without crossing any lattice point). Note that S' is a close approximation of ∂M and has straight lines, smooth curves, and sharp corners that match those of ∂M.

In a similar fashion, the 3D version of the Pressing algorithm works as follows:

- We cluster sticks [ABC*04] into flats that may be stabbed by a plane (Figure 1-b).
- We identify junction points: vertices between a flat and a curved region (Figure 1-b) [Section 3].
- We fair non-flat regions by an iterative process, which, at each step and for each fresh vertex combines arc-length re-sampling, bilaplacian smoothing, and snapping. These three operations are performed independently on each X, Y, and Z slice [NGH*03]. Their results are combined for each stick using a special filter at borders [Section 4].
- Sharp edges are recovered through the use of a modified Edge Sharpener algorithm [AFRS04, AFSR05] (Figure 1-c) [Section 5].

In summary, the technical contributions presented in this paper are:

- We propose to combine segmentation (for recovering flats) and smoothing (for joining them with smooth transition surfaces).
- We modify iso-surface smoothing to preserve portions frozen by equality constraints corresponding to detected flats.
- We present a new smoothing operator, which preserves the connectivity of the initial iso-surface representation and achieves smoothness in the presence of equality and inequality constraints.

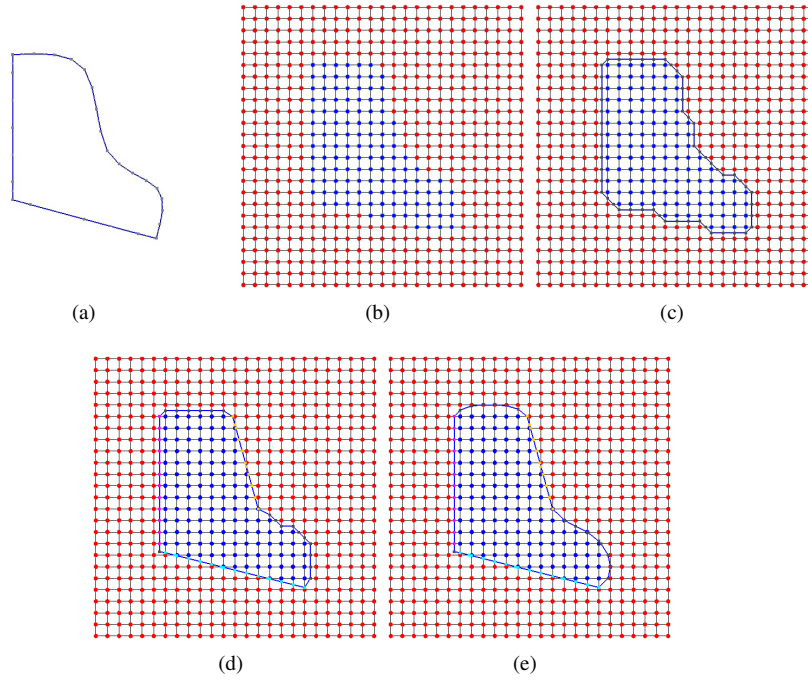The combination of these advances leads to new functionalities:

Figure 2: *2D region M bounded by straight and curved edges (a) Red/blue classification of grid-samples produced by rasterizing M. (b) Reconstructed isocontour S with vertices at stick midpoints. (c) Straight edges are recovered by snapping the vertices of each cluster to its flat. (d) A pressed version S' of S is obtained after 50 pressing iterations. (e) Its straight and curved edges correspond to those of M.*

- Flat regions, curved regions, and sharp edges can be automatically recovered from raw binary voxelizations even though scalar field and Hermite data are not available.
- The reconstruction error is therefore bounded, since the final isosurface is constrained to stab the initial sticks and is completely contained in the set of mixed cells.
- The isosurface is automatically segmented into flat and curved regions, which may facilitate shape identification, manufacturing and assembly planning.

The paper is organized as follows. Section 2 discusses the previous work and alternative approaches to the problem. Section 3 describes how large planar regions are identified along sharp features. In Section 4, a modified bilaplacian filter algorithm is used to smooth the rest of the isosurface. Finally, sharp features are recovered as explained in Section 5. Sections 6 and 7 discuss the obtained results and its potential applications.

2. Previous work

Algorithms that extract isosurfaces from a discrete sampling of a scalar field on a regular grid strive to ensure topological consistency and geometric fidelity [LC87, NFHL91, Nie03, Lac96, MSS94, ABC*05]. Some approaches base topological decisions on scalar field values [CGMS00, Nie03, LB03]

or Hermite data [HWC*05], i.e. the estimates of surface normals at the vertices. Geometric fidelity often depends on the delicate ability to recover sharp features (which may follow smooth curved edges) and to smooth the isosurface away from these edges. The Extended Marching Cubes [KBUS01] detects cells containing features and recovers them by inserting an additional point in each one of these cells. The Dual Contouring algorithm [JLSW02] uses a quadratic error metric to compute a new point inside each of the four cells around each stick and generates a quad connecting these four new points. A similar approach is adopted in [VKSM04]. Several of these methods make use of Hermite data and thus cannot be applied to recover sharp features on binary grids where only the in/out classification of the grid nodes is stored.

Mesh smoothing algorithms strive to remove noise and high-frequency details from a general triangulated surface by the iterative application of a smoothing operator. Most approaches derive non-shrinking smoothing operators from discrete approximations of the Laplacian [Tau95, Kob97, DMSB99].

One of the first approaches that proposed an algorithm for smoothing Marching Cubes isosurfaces from binary (and not binary) volume models is [OB01]. A major concern of

smoothing techniques is the addition of constraints on the vertex placement to guarantee the separation of in/out grid nodes [Gib98,Whi00,NGH*03,Nie04]. Gibson [Gib98] proposed an algorithm for reducing the terracing artifacts in isosurfaces extracted from binary grids. It places one vertex in each mixed cell and then links the vertices in face-connected mixed cells to form a net. This net is relaxed to reduce the energy measure in the links, and this relaxation process minimizes edge lengths, similar to a Laplacian. A constraint is applied to keep each node in this original cell. After the relaxation a triangulated surface, which may not be a manifold, is generated in a straightforward way. Nielson et al. [NGH*03] proposed a closely related technique to our constrained smoothing approach. Like Pressing, mesh vertices are moved along the sticks and surface smoothing is obtained by combining two orthogonal polygonal smoothing operations. However, the displacement for each vertex in [NGH*03] is driven by a non-linear optimization algorithm that minimizes an energy function defined on each polygonal curve. A different approach is presented in [Nie04] which introduced a smoothing operator based on the dual of the dual surface of a Marching Cubes mesh, locating the vertices at the intersections of the dual's quads with the lattice edges. In [Fre04] a smoothing algorithm for biomedical data that does not shrink the model is presented, but it does not restrict the isosurface to remain inside the discrete band and does not recover features.

The computation of planar regions (flats) approximating a given geometric model is an important problem with wide applications in computer vision, modeling and impostor-based simplification. The technique we apply for detecting flats is related to superfaces [KT94] and face clustering [GWH01, She01], which group connected sets of nearly coplanar faces of a given triangulated surface. Superfaces uses a greedy algorithm to cluster triangles. Each triangle in a superface imposes constraints on the set of feasible approximating planes for the superface, most notably that the triangle's vertices must be within a fixed distance of the planes. Hierarchical face clustering [GWH01] uses quadric error metrics [GH97] to iteratively merge adjacent faces. Cohen-Steiner et al. [CSAD04] adopt a variational geometric partitioning approach to group faces into best-fitting regions according to a normal deviation error metric. Decoret et al. [DDSD03] use an optimization algorithm to find a set of approximating planes, using a discretization of a plane parameterization in spherical coordinates, and propose a greedy optimization algorithm of a density field in this plane parameterization. The main drawback is the time complexity of the plane optimization algorithm and the lack of uniformity in the parameterization of planes.

Unlike the approaches above, which require a triangulated surface, Andujar et al. [ABC*04] propose an efficient algorithm for the computation of the largest flat region (*tile*) in a discrete geometric model. The input of the algorithm is the set of sticks. Using a voting-based approach, the plane that slices the largest number of sticks is computed.

3. Detection of flats and junction points

The first step of Pressing identifies sufficiently large flat clusters of sticks and freezes their vertices on the best fit plane for each cluster. To do so, we use the approach proposed in [ABC*04] for computing maximum tiles (flat regions).

A large flat is characterized by the fact that there is a plane that separates black nodes from white ones. Thus, the problem of finding a large planar region is transformed to that of looking for the plane which intersects the maximum number of sticks. Every stick votes for all the planes on a given discretization which intersect them. After this process, the plane with the largest number of votes is the largest stabbing plane. This procedure may be repeated by removing from the vote process all those sticks which stab previously obtained planes.

As shown in Figure 1-b, the boundary of a flat F may include three types of edges. (1) Edges connecting F to another flat. (2) Sharp edges connecting F to a curved surface. (3) Smooth edges connecting F to a curved surface with normal continuity. Edges of type (1) and (2) will be identified as *sharp edges* and will be preserved. Edges of type (3) will be faired by our smoothing algorithm.

To detect cases where (2) or (3) apply, and before proceeding to the smoothing step, we consider every edge of S joining a frozen vertex V_p on a planar region and a fresh vertex V_s on a smooth region. Let N_p be the normal of the tile stabbing V_p. Let N_s be normal to the plane produced by a least square fit to the fresh vertices on the 2-ring neighborhood around V_s. When the angle between N_p and N_s exceeds 30 degrees, we decide that the edge is a chamfer-edge cutting through a sharp edge of S and label V_s as a junction point. After experimenting with various angles, we have empirically concluded that a 30 degree threshold leads to the best compromise limiting the false positives and false negatives.

4. Smoothing and snapping

In the next step, we seek to smooth out the surface obtained thus far, with a special consideration for sharp edges between smooth and planar portions. To this end, we iterate the following two steps of our approach: Smoothing and Snapping. Together, they iteratively modify the positions of the fresh vertices.

The current literature contains different algorithms that address similar situations, but which fail here for different reasons. Those which are not prepared to deal with non-uniform sample spacing do not yield an acceptable result ([Tau95,Gib98]). Given that the snapping step, which maintains vertices on their sticks, preserves non uniformity, it also affects the smoothing step adversely. In order to compensate

Figure 3: *The smoothing operator (drawn in blue) is computed on the corresponding two isocurves (in cyan). Each isocurve is on an axis-aligned slice. The resulting combination is shown in orange.*


(a)

(b)

Figure 4: *The applied bilaplacian smoothing $L^2(C)$ can be computed as the difference $L(C) - L(C')$ of two Laplacian displacements. $L(C)$ moves C to the average C' of its neighbors. Now assume that B and D have also been moved to the averages B' and D' of their neighbors. $-L(C')$ moves C' to C''.*

for this effect, a filtering method like cotangent weights may be applied. The problem remains though, because vertices' movement may cause some edges to collapse, which in turn can produce spikes on the surface. To fix this problem, we have implemented a constrained modified version of the bilaplacian filter.

The 3D smoothing step computes for each fresh vertex C of S a displacement vector w along the stick I of C. w is obtained as a linear combination of the two displacements, computed in each of the two different axis-aligned slices of the grid that contain I, as shown in Figure 3.

Consider one such slice. Assume that C lies on a curve where the slice intersects S. To compute the corresponding displacement, we have developed a variant of the bilaplacian smoothing, which uses two points at a fixed arc-length distance from C along the curve on each side of C. The construction is explained below. The resulting displacement is projected onto the line supporting the stick I.

Consider five consecutive vertices (A, B, C, D, E) along a polygonal curve. As shown in Figure 4, the bilaplacian smoothing displacement vector $L^2(C)$ associated with vertex C may be computed as:

$$L^2(C) = \frac{-A + 4B - 6C + 4D - E}{4} \qquad (1)$$

We first precompute and store $L^2(C)$ for each vertex C. Then, we move each vertex C to $C'' = C + s^2 L^2(C)$, where s represents the step size of the filter at each iteration. Small values of s need more iterations to converge, but a too large

value destabilizes the smoothing process. We have found that a value of $s = 0.85$ yields good results.

Next, the snapping step projects each displaced vertex C'' to the closest point on the line supporting the stick it came from. Then, it is constrained to its stick to ensure that the isosurface will not cross any nodes. As a consequence, the node classification is not altered. Anyway this naïve approach does not converge to the smoothest possible curve subject to the constraints (Figure 5). We found that, rather than approaching zero, some displacement vectors eventually become orthogonal to the sticks, so that projecting a displaced vertex simply returned it to its previous position. Hence, the curve is stuck in a suboptimal shape. Furthermore, when the curve converges to a node of the grid, the node imposes inequality constraints on the displacements of vertices of incident sticks. Thus as two or three of these vertices converge towards the node, the arc-length parameterization of the samples along the curve is no longer uniform. Because the formula in equation 1 was developed for converging to a uniform parameterization, it performs poorly near these nodes, creating sharp discontinuities.

To solve this problem, we have explored a variety of alternatives, including using a more general cubic fit to non-uniformly distributed samples, whose parameters in the parametric cubic expression are estimated from the arc-length distances between vertices. We have concluded that

Figure 5: *The color-coding of the nodes was obtained by rasterizing a circle. Computing the bilaplacian from neighboring vertices and snapping does not converge to an acceptable approximation of the circle (center). Applying arc-length resampling yields a much better fit (right).*


Figure 6: *Instead of using the vertices B and D, two virtual neighbors (CR and CL) are computed on each side of C by moving a fixed amount along the curve .*

the resampling approach described below leads to the most effective smoothing. Hence when computing the bilaplacian $L^2(C)$, we do not use A, B, D, and E — the neighboring vertices of C on the curve. Instead, we compute new samples at a fixed distance d along the curve in both directions (Figure 6). This arc-length resampling prevents the formation of unwanted corners and yields satisfactory results when used with the snapping. In any case, d has to be chosen large enough to result in big steps (which translate into less iterations), while not being so large that it lets the curve twist around too much. We select d to be 75% of the length of a cell's side. As the bilaplacian filter is computed in two passes, we only need to add 2 points at distance d at each pass, rather than adding four at distances d and $2d$, which is equivalent.

Next we turn the combination of these results into a displacement w. We could compute the displacement one vertex at a time. For example, consider the vertex V on a stick I that is parallel to the Z axis. We could compute its displacements in the $X - Z$ section of the grid that contains I and on the $Y - Z$ section. Then we would combine these displacements ensuring that the resulting vertex remains on I. Instead, for simplicity and implementation efficiency, we perform the smoothing and line-snapping steps on each $X - Y, X - Z$, and $Y - Z$ slice and then collect the results, combining two

displacements for each stick and clamping the result to the stick. Both approaches decompose the bivariate surface bilaplacian into the equivalent combination of two univariate curve bilaplacians, so we choose the second one which results in a simpler implementation.

Because the stick of each vertex belongs to exactly two slices [NFHL91], we have two suggested displacements for each vertex. We average the two displacement vectors to obtain the vertex's displacement and then snap the displaced vertex to the closest point onto its stick. When averaging the two displacements different weighting techniques can be applied. We have experimented with three methods: (1) equal weights, (2) weights are proportional to the dot products of the stick tangent with the normal to the curve, and (3) weights are proportional to the dot products of the stick tangent with the tangent to the curve. There are special cases in which (2) yields markedly worse results, and cases where (3) does that. We found that the simpler approach in case (1) uniformly yields good results that are close to those of the best of the other two.

In order to avoid artifacts at junction areas we have modified the smoothing filter for vertices that lie close enough to a feature edge. If a vertex is at a distance less than d of a feature edge, the sampled neighbors of our arc-length resampling method cannot be reliably computed. Those vertices do not use the bilaplacian. Instead, they align themselves with the vertices on the smooth side (Figure 7). In the absence of any information on this junction, we have adopted this unclamped approach, which has been used in generating all the pictures.

5. Sharp edge recovering

The errors between the original shape and the iso-surface recovered thus far are usually concentrated near the features, which were not appropriately captured by the regular sam-

Figure 7: *The "smooth" vertices on a feature edge move themselves along their stick to be aligned with the neighbors on their "smooth" side.*

pling. To improve the accuracy of the recovered surface, we sharpen the boundaries of planar regions using a variant of EdgeSharpener [AFRS04, AFSR05].

The vertices belong to either a tile (planar face) or to a curved face. We use the term chamfer edges for those mesh edges with vertices on two different faces. A triangle with one or more chamfer edges is called chamfer triangle.

In order to recover the sharp features we apply three steps. First, the chamfer triangles are identified. Then, we subdivide them appropriately by inserting new vertices. Finally, we position the new vertices to better recover the sharp features.


Figure 8: *Subdivision of a chamfer triangle with one (a) two (b) or three (c) chamfer edges.*

Three cases arise when subdividing the chamfer triangles (Figure 8):

1. Triangles with one chamfer edge are split into two triangles.
2. Triangles with two chamfer edges result in three triangles.
3. When three different faces meet at a triangle we have three chamfer edges. After subdividing we will have six triangles and one interior vertex to represent the corner.

The process is presented in Figures 9 and 10. To find the position of a new vertex V inserted in a chamfer edge E, we consider the two original vertices, A and B, of E. We compute a normal N_A for the vertex A using its face, and define a plane P that is orthogonal to N_A and passes through A. Similarly, we compute a normal N_B for the vertex B using its face, and define a plane Q that is orthogonal to N_B and passes through B. Finally, we move V to the closest point on the line of intersection between planes P and Q. When one of


Figure 9: *Inserting a new vertex on a triangle with its vertices on three different faces.*


Figure 10: *Edges with vertices on two different faces are subdivided to recover a feature.*

these vertices belongs to a curved face, its normal is taken to be the normal to a plane estimated as in the computation of junction points (the minimum square fit to the free neighbors in a 2-ring).

To find the position of an interior vertex W, we consider the vertices A, B and C of the corner triangle. We compute normals N_A, N_B and N_C using their respective faces. Using them we define planes P, Q and R which are orthogonal to N_A, N_B and N_C, and pass through A, B and C, respectively. Then, W is moved to the intersection point of P, Q and R.

Junction points estimated by the Edge Sharpener are based on inaccurate nodes, and therefore are noisy. Thus, we recognize features by arranging the new vertices into chains that follow the newly inserted edges. The curves defined by these chains are smoothed using the bilaplacian algorithm described previously. Vertices at any of the two endpoints of a chain are left untouched. This includes vertices generated by the intersection of three tiles.

6. Results and discussion

In this section we present and discuss several examples. They are shown in Figures 11, 12, 13 and 14. All isosurfaces are rendered using flat shading with surface normals calculated directly from each triangle in order to emphasize the underlying geometry. The initial surfaces of these models have been converted into a binary voxel representation, and these voxelizations have been used as the input to Pressing.

Figure 11: *A* $128 \times 128 \times 128$ *voxelization of a pump model, and the final pressed isosurface.*


Figure 12: *A* $128 \times 128 \times 128$ *voxelization of a mechanical part (referenced as MECH PART in Tables 1 and 2), the final pressed isosurface, and a closer look to the recognized features.*

In order to show the performance of Pressing in the most general case, we have intentionally applied a random rotation to the models to ensure that the main faces are not axis-aligned. The leftmost part of Figure 13 displays the results of just smoothing the Fandisk model, without processing flats separately. It also displays the improvements achieved with Pressing by showing the results at resolutions of 128 and 256.

All these models show several sharp edge features between flat regions. Although the initial information is only a binary voxelization, our algorithm is successful in detecting and reconstructing flats and sharp edges. In Figures 11 and 12, Pressing also recovers the smooth regions of the model and the curved features between flat and smooth pieces. Tables 1 and 2 summarize the performance of the algorithm to obtain the results depicted in Figures 15, 11, 12, 13 and 14. Notice —by comparing Figure 11 with Figure 1-c (which uses 300 iterations)— that a larger number of iterations can achieve still smoother results, at the expense of time. The smoothing algorithm may be also applied to med-

ical models and the results may be seen in Figure 15. These display the good behavior of our constrained smoothing algorithm even when there are no features.

Table 1 presents the running time of our algorithm, running on a Pentium 4 at 3.4 GHz and 1 GB of RAM. The Max Tiles step for the detection and reconstruction of flat regions is based on a previous work and is not presented as a contribution in this paper. The times for detecting sharp features and edge sharpening are not significant in front of the time complexity of the smoothing part of the algorithm. This smoothing step is however below 10 seconds in all $128 \times 128 \times 128$ CAD models. The higher times in the second mechanical part are a consequence of the finer $256 \times 256 \times 256$ voxelization.

We have observed that our algorithm achieves excellent approximations to the original models in the examples tested. Table 2 presents the evolution of the reconstruction errors, which are computed as the average of the unsigned distances between a vertex of the isosurface and the intersection of its stick with the original model. Voxels are con-

Figure 13: *From left to right, the result of smoothing a* $128 \times 128 \times 128$ *voxelization of the Fandisk model, its corresponding pressed isosurface, and the pressed isosurface obtained from a* $256 \times 256 \times 256$ *discretization. Notice that the result at* $256 \times 256 \times 256$ *contains fewer artifacts, because at that resolution features are sufficiently separated by binary samples.*


Figure 14: *A* $256 \times 256 \times 256$ *voxelization of a second mechanical part (referenced as MECH PART 2 in Tables 1 and 2), and the final pressed isosurface.*

sidered to be of size $1.0 \times 1.0 \times 1.0$. The medical models are not included in Table 2, because they are the result of the binary segmentation of a volumetric model.

The first row measures the error between the scanned isosurface and the midpoint isosurface, the second one adds the planar regions, while the third and the fourth ones include the smoothing with 100 and 1000 iterations each. The errors monotonically decrease at each Pressing step, and reach small values after only 100 smoothing operations. The increase of accuracy when the number of iterations goes from 100 to 1000 is not significant.

7. Conclusions

We have proposed a novel smoothing approach for the automatic recovery of solids from binary volumetric discretizations. Our approach uses global optimization to identify flats and a constrained smoothing algorithm to recover the shape of non-planar regions. The proposed smoothing algorithm involves a snapping step after each bilaplacian smoothing step to guarantee that final vertices remain in the initial sticks of the voxelization.

Pressing works on general binary voxelizations and can recover flat and curved regions, and is hence particularly useful when no scalar field data or Hermite data are available. The isosurface is automatically segmented by sequences of junction points and it is constrained to stab the initial sticks. The reconstruction error is therefore bounded and the topology is preserved.

We use a three dimensional implementation of the constrained smoothing, which combines two two-dimensional smoothing steps for each vertex, one in each axis-aligned plane containing the vertex's stick, followed by snapping. A special version of the filter for vertices at the borders of the curved regions has been also developed.

Results on a variety of models have been reported and discussed. Pressing achieves small reconstruction errors and

Figure 15: *After segmenting a volume, the extracted isosurface may be smoothed using our algorithm. These pictures show the results of doing so, on the $256 \times 256 \times 256$ discretizations of a skull and a jaw.*

Stats (times in seconds)	PUMP	MECH PART	MECH PART 2	FANDISK	SKULL	JAW
TIME (MaxTiles)	73.39	15.99	94.38	58.89	0	0
TIME (Detect Junctions)	0.223	0.175	0.745	0.272	0	0
TIME (Smoothing)	8.296	5.314	29.207	6.036	113.619	61.563
TIME (Edge Sharpening)	0.221	0.248	0.693	0.244	0	0
# Tiles	8	8	10	11	0	0
# Iterations	52	68	98	65	100	100

Table 1: *Running time, number of reconstructed flat regions and number of iterations for each of the presented models.*

successfully recovers flats and sharp features in a reasonable amount of time.

Potential applications include shape recognition, simplification, compression and various reverse engineering and manufacturing problems.

8. Acknowledgements

References

[ABC*04] ANDÚJAR C., BRUNET P., CHICA A., NAVAZO I., ROSSIGNAC J., VINACUA A.: Computing maximal tiles and application to impostor-based simplification. *Computer Graphics Forum 23*, 3 (2004). Proceedings of Eurographics'04.

[ABC*05] ANDÚJAR C., BRUNET P., CHICA A., NAVAZO I., ROSSIGNAC J., VINACUA À.: Optimizing the topological and combinational complexity of isosurfaces. *Computer-Aided Design 37*, 8 (2005), 847–857.

[AFRS04] ATTENE M., FALCIDIANO B., ROSSIGNAC J., SPAGNUOLO M.: Edge-sharpener: Recovering sharp features in triangulations of non-adaptively re-meshed surfaces. In *Proc. of EG/ACM SIGGRAPH Symposium on Geometry Processing 2004* (2004), pp. 62–69.

[AFSR05] ATTENE M., FALCIDIANO B., SPAGNUOLO

Error	PUMP	MECH PART	MECH PART 2	FANDISK
Midpoint	0,2887	0,2837	0,2874	0,2882
Midpoint + Tiles	0,1254	0,2267	0,1706	0,1835
MidPoint + Tiles + Smoothing (100)	0,0462	0,0856	0,0738	0,0892
MidPoint + Tiles + Smoothing (1000)	0.0427	0,0850	0,0728	0,0857

Table 2: *Square root of the average square error after the different steps of the Pressing algorithm for each of the presented models.*

M., ROSSIGNAC J.: Sharpen&blend: Recovering curved edges in triangle meshes produced by feature-insensitive sampling. *IEEE Transactions on Visualization and Computer Graphics 11*, 2 (2005), 83–91.

[BPK05] BISCHOFF S., PAVIC D., KOBBELT L.: Automatic restoration of polygon models. *ACM Transactions on Graphics 24*, 4 (2005), 1332–1352.

[CGMS00] CIGNONI P., GANOVELLI F., MONTANI C., SCOPIGNO R.: Reconstruction of topologically correct and adaptive trilinear isosurfaces. *Computers and Graphics 24*, 3 (2000), 399–418.

[CSAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. *ACM Trans. Graph. 23*, 3 (2004), 905–914.

[DDSD03] DÉCORET X., DURAND F., SILLION F. X., DORSEY J.: Billboard clouds for extreme model simplification. *ACM Transactions on Graphics 22*, 3 (July 2003), 689–696.

[DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 317–324.

[EBV05] ESTEVE J., BRUNET P., VINACUA A.: Approximation of a variable density cloud of points by shrinking a discrete membrane. *Computer Graphics Forum 24*, 4 (2005), 791–808.

[Fre04] FREY P. J.: Generation and adaptation of computational surface meshes from discrete anatomical data. *International Journal for Numerical Methods in Engineering 60*, 6 (2004), 1049–1074.

[GH97] GARLAND M., HECKBERT P. S.: Surface simplification using quadric error metrics. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co., pp. 209–216.

[Gib98] GIBSON S.: Constrained elastic surface nets: generating smooth surfaces from binary segmented data. In *MICCAI'98, Medical Image Computation and Computer Assisted Surgery* (1998).

[GWH01] GARLAND M., WILLMOTT A., HECKBERT P. S.: Hierarchical face clustering on polygonal surfaces. In *Proceedings of ACM Symposium on Interactive 3D Graphics* (Mar. 2001), ACM Press.

[HWC*05] HO C.-C., WU F.-C., CHEN B.-Y., CHUANG Y.-Y., OUHYOUNG M.: Cubical marching squares: Adaptive feature preserving surface extraction from volume data. *Computer Graphics Forum (Eurographics 2005) 24*, 3 (2005), 537–546.

[JLSW02] JU T., LOSASSO F., SCHAEFER S., WARREN J.: Dual countouring of hermite data. *ACM Transactions on Graphics 21*, 3 (2002), 339–346. Proc of Siggraph'02.

[KBUS01] KOBBELT L. P., BOTSCH M., U. SCHWANECKE H. P. S.: Feature sensitive surface extraction from volume data. *ACM Computer Graphics (Siggraph 2001)* (2001), 57–66.

[Kob97] KOBBELT L.: Discrete fairing. In *Proceedings of the Seventh IMA Conference on the Mathematics of Surfaces* (1997), pp. 101–131.

[KT94] KALVIN A. D., TAYLOR R. H.: Superfaces: Polyhedral approximation with bounded error. In *Medical Imaging: Image Capture, Formatting, and Display* (Feb. 1994), vol. 2164, SPIE, pp. 2–13. (Also IBM Watson Research Center tech report RC 19135).

[Lac96] LACHAUD J.-O.: Topologically defined isosurfaces. In *Proc. 6th Discrete Geometry for Computer Imagery (DGCI'96), Lyon, France* (1996), Springer-Verlag, Berlin, pp. 245–256.

[LB03] LOPES A., BRODLIE K.: Improving the robustness and accuracy of the marching cubes algorithm for isosurfacing. *IEEE Transactions on Visualization and Computer Graphics 9*, 1 (2003), 16–29.

[LC87] LORENSEN W., CLINE H.: Marching cubes: A high resolution 3D surface construction algorithm. *Computer Graphics 21*, 4 (1987), 163–169.

[MSS94] MONTANI C., SCATENI R., , SCOPIGNO R.: Discretized marching cubes. In *IEEE Visualization* (1994), pp. 281–287.

[NFHL91] NIELSON G., FOLEY T., HAMANN B., LANE D.: Visualizing and modeling scattered multivariate data. *IEEE Computer Graphics and Applications 11*, 3 (1991), 47–55.

[NGH*03] NIELSON G., GRAF G., HOLMES R., HUANG A., PHIELIPP M.: Shrouds: Optimal separating surfaces for enumerated volumes. In *EG-IEEE TCVG Symposium on Visualization 2003* (2003), pp. 75–84.

[Nie03] NIELSON G.: On marching cubes. *IEEE Transactions on Visualization and Computer Graphics 9*, 3 (2003), 283–297.

[Nie04] NIELSON G.: Dual marching cubes. In *IEEE Visualization 2004* (2004), pp. 489–496.

[OB01] OHTAKE Y., BELYAEV A.: Mesh optimization for polygonized isosurfaces. *Computer Graphics Forum 20*, 3 (2001).

[She01] SHEFFER A.: Model simplification for meshing using face clustering. *Computer Aided design 33*, 13 (2001), 925–934.

[Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1995), ACM Press, pp. 351–358.

[TSH98] TIEDE U., SHIEMANN T., HOEHNE K.: High quality rendering of attributed volume data. In *Proceedings of the IEEE Visualization'98 Conference* (1998), IEEE, pp. 255–262.

[VKSM04] VARADHAN G., KRISHNAN S., SRIRAM T., MANOCHA D.: Topology preserving surface extraction using adaptive subdivision. In *Proc. of EG/ACM SIGGRAPH Symposium on Geometry Processing 2004* (2004), pp. 241–250.

[Whi00] WHITAKER R. T.: Reducing aliasing artifacts in iso-surfaces of binary volumes. In *VVS '00: Proceedings of the 2000 IEEE symposium on Volume visualization* (New York, NY, USA, 2000), ACM Press, pp. 23–32.