

Edge Transformations for Improving Mesh Quality of Marching Cubes

Carlos A. Dietrich, Carlos E. Scheidegger, John Schreiner, João L.D. Comba, Luciana P. Nedel, and Cláudio T. Silva, *Member, IEEE*

Abstract—Marching Cubes is a popular choice for isosurface extraction from regular grids due to its simplicity, robustness, and efficiency. One of the key shortcomings of this approach is the quality of the resulting meshes, which tend to have many poorly shaped and degenerate triangles. This issue is often addressed through postprocessing operations such as smoothing. As we demonstrate in experiments with several data sets, while these improve the mesh, they do not remove all degeneracies and incur an increased and unbounded error between the resulting mesh and the original isosurface. Rather than modifying the resulting mesh, we propose a method to modify the grid on which Marching Cubes operates. This modification greatly increases the quality of the extracted mesh. In our experiments, our method did not create a single degenerate triangle, unlike any other method we experimented with. Our method incurs minimal computational overhead, requiring at most twice the execution time of the original Marching Cubes algorithm in our experiments. Most importantly, it can be readily integrated in existing Marching Cubes implementations and is orthogonal to many Marching Cubes enhancements (particularly, performance enhancements such as out-of-core and acceleration structures).

Index Terms—Meshing, marching cubes.

1 INTRODUCTION

ISOSURFACES are ubiquitous in visualization [14], where they are often the main computational component of important processing pipelines and are heavily used in practice. Some visualization applications require the ability to render an isosurface, and this can be done by first converting this implicit surface into a triangle mesh and then rendering it. Other applications, however, operate on the resulting mesh, such as finite-element method simulations, tetrahedral mesh generation, and inverse problems. These applications require meshes of good quality, which are often dictated by the quality of its worst triangle [30], *regardless* of any other triangle in the mesh.

The classic approach to compute isosurfaces is to apply the Marching Cubes (MC) algorithm [22] or one of its variants, such as [5] and [11]. Although robust and simple to implement, these generate surfaces that require additional processing steps to improve triangle quality and mesh size. All of these algorithms are cell based: they work by iteratively examining each cell of the grid on which the scalar function f is defined and by producing a triangulation for each cell

separately. These triangulations are created in such a way that when they are connected together, they produce a watertight manifold mesh [25]. The simplicity of these methods allows robust and efficient algorithms, which have been expanded and extended in significant ways [2], [7], [20], [36]. Among these techniques are optimization strategies that are orders of magnitude faster than the original algorithm and can work on data of arbitrarily large sizes. However, they still produce low-quality triangles.

In this paper, we address the issue of improving the quality of the worst triangle generated by MC. Fig. 1 illustrates a comparison of MC against several other strategies, showing triangle quality, as well as the histogram of triangle qualities and Hausdorff distances between the mesh generated by each approach and the original MC mesh (as measured by Metro [8]). The red line in the histograms shows the quality of the worst triangle. In essence, that is the measure we are interested in, since it often dictates the quality of the simulations performed on the mesh [30].

We show in this paper a novel method to improve the quality of the triangle mesh generated by MC, by modifying the MC sampling grid in a very simple and intuitive way. We call our method **Marching Cubes using Edge Transformations (Macet)**. Macet generates a mesh with identical connectivity to the MC mesh, which is very close to the MC mesh as measured by Hausdorff distance. Most importantly, it consistently generates meshes whose worst triangles are well above the current state-of-the-art techniques. Macet is simple to implement and very fast. Although we have no lower bound proof for the triangle quality in the general case, we provide extensive experimental evidence and a fully open-source implementation.

MC often generates bad triangles because the regular sampling structure of the grid leads to intersections being

- C.A. Dietrich, J.L.D. Comba, and L.P. Nedel are with the Universidade Federal do Rio Grande do Sul, Instituto de Informática, Av. Bento Gonçalves, 9500 - Campus do Vale - Bloco IV, Bairro Agronomia - Porto Alegre - RS - Brasil, CEP 91501-970 Caixa Postal: 15064. E-mail: {cadietrich, comba, nedel}@inf.ufrgs.br.
- C.E. Scheidegger is with the University of Utah, 1120 S. 1300 E., Salt Lake City, UT 84105. E-mail: cscheid@sci.utah.edu.
- J. Schreiner and C.T. Silva are with the School of Computing, University of Utah, 50 S. Central Campus Dr. RM 3190, Salt Lake City, UT 84112. E-mail: {jmschrei, csilva}@sci.utah.edu.

Manuscript received 28 June 2007; revised 14 Nov. 2007; accepted 12 Mar. 2008; published online 7 Apr. 2008.

Recommended for acceptance by M. Chen.

For information on obtaining reprints of this article, please send e-mail to: tvccg@computer.org, and reference IEEECS Log Number TVCG-2007-06-0078.

Digital Object Identifier no. 10.1109/TVCG.2008.60.

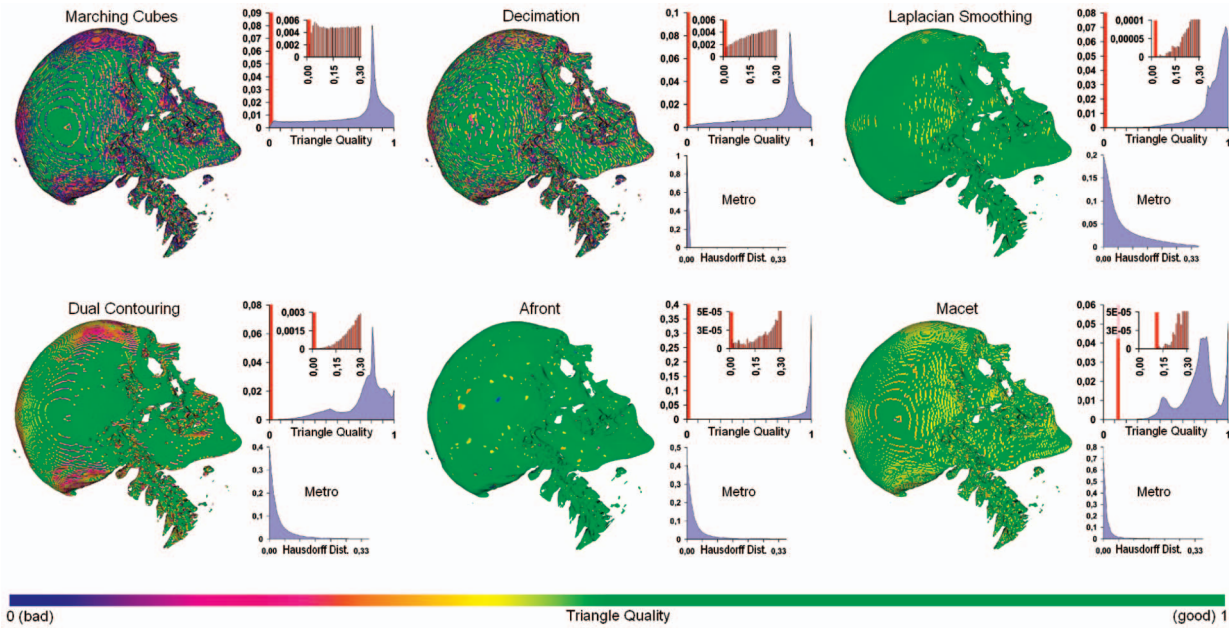


Fig. 1. Improving mesh quality in MC. From top to bottom, left to right: Visible human data set, mesh generated by the classic MC (base mesh used in the comparisons), and results obtained with a postprocessing using Decimation to 90 percent, Laplacian smoothing, Dual Contouring, afront, and our modified **Marching Cubes** using **Edge Transformations** (Macet). Quality is color-coded on a per-vertex basis, defined from the minimum radii ratio of all incident triangles, and color-coded from 0 (degenerate triangle) to 1 (equilateral triangle). Triangle quality and Metro histograms show the quality of the worst mesh triangles (indicated by the red line) and the forward distance between the mesh generated by MC in each alternative, respectively. We are interested in the quality of the worst triangle (and not on the overall histogram shape), and thus, each histogram is scaled to emphasize the histogram bin resulting from the worst triangle of the mesh. Even though smoothing methods greatly improve the overall quality of the mesh, they still generate a considerable amount of badly shaped triangles. Macet, on the other hand, did not produce a single degenerate triangle during our experiments.

computed at inconvenient locations on a subset of the active edges (often close to one of the endpoints of the active edge). By allowing active edges to be repositioned (i.e., moving their endpoints) to more adequate locations, we improve the quality of the triangles that use the intersections generated. Although the initial sampling grid of the MC is modified, we provide conditions that preserve the topology of the mesh. Our algorithm is easy to understand and implement. It keeps most of the MC structure intact, being implemented as a new stage between the detection of active edges and the intersection calculation. Given this, most of the suggested MC optimizations (like Span Space [29]) in the literature still apply. The main contributions introduced in this paper are given as follows:

- A novel approach for understanding the triangle quality distribution of meshes generated by MC. This work motivated the ideas behind the edge transformations that are shown to eliminate the bad triangles generated by MC.
- A revised MC algorithm called Macet that combines edge transformations to generate an output mesh with the same connectivity and small error (as measured by Metro) to MC and also with improved quality of its worst triangle.
- A detailed experimental study comparing the triangle quality of existing techniques.

In Section 2, we present a discussion on related techniques proposed to solve several issues raised in this work. Then, in Section 3, we study the issues related to the mesh quality generated by MC. This leads us to revisit the

inner computation in those techniques and to propose an alternative scheme, described in Section 4. The latter part of this paper presents our experimental results, discussion, and future work.

2 RELATED WORK

Isosurface polygonization methods are efficient tools for extraction and visualization of isosurfaces since the pioneering work in the early 1980s [1], [6], [15]. Methods based on surface tracking place *seed* sampling points on the isosurface and perform an iterative refinement search for optimal positions or generation of new seeds. Examples include pseudo-physical algorithms [10], [37] and advancing front (*afront*) algorithms [28]. Spatial decomposition methods, described first by Herman and Udupa [15], rely on the assumption that inside a smaller cell of a grid we could assume that the underlying scalar function is locally linear [35], and thus, the isosurface can be represented by a plane. Using this simple assumption, Lorensen and Cline [22] proposed the MC algorithm, which is arguably the most important isosurface polygonization algorithm due to its simplicity, efficiency, and robustness.

However, MC is also known for the poor quality of the resulting triangle mesh. Several strategies for measuring triangle quality are discussed in [26]. The metric most commonly used for measuring a single triangle quality is the radii ratio: the ratio between the triangle’s incircle and circumcircle. The radii ratio is a *fair* metric: every degenerate triangle has a radii ratio of zero. It penalizes both small and large angles, which makes it suitable as a “first-order” metric: although some applications might not have




Fig. 2. Several possible plane orientations colored by triangle quality.

problems with small or large angles, a fair metric will favor meshes that are suitable across many application domains. For measuring the quality of meshes, Shewchuk [30] states that the quality of the worst triangle gives a better estimate of mesh quality than other aggregate metrics, such as average or median quality.

In order to overcome such shortcomings of MC, several extensions are discussed in the literature [3], [4], [5], [16], [21], [23], [24], [33], [34], [35]. To obtain a high-quality triangle mesh from MC, postprocessing steps are typically applied directly to the triangle mesh [10]. For instance, one might apply a standard smoothing algorithm such as Laplacian Smoothing (LS) to improve the triangle qualities or a decimation procedure to remove degenerate triangles. Observe in Fig. 1 that although LS improves the quality, it also moves the surface away from the original MC mesh (as can be seen on the Hausdorff error histogram). QEM-based decimation [12] is more faithful to the original mesh but fails to remove many degenerate triangles. Several other polygonization methods have been proposed to generate better quality meshes [9], [10], [13], [28]. These techniques often take a more global approach that try to optimize vertex sampling over the complete isosurface. For instance, the recently developed *afront* algorithm of Schreiner et al. [28] works by first creating an initial seed point on the isosurface and iteratively growing the triangulation over the entire isosurface, while guaranteeing a global grading constraint over most triangles. It can produce very high quality adaptive triangle meshes, albeit at a computational cost that is much higher than MC. Additionally, the algorithm itself is more involved, leading to a higher complexity of the implementation, and still prone to generating a few degenerate triangles (Fig. 1). Dual methods such as Dual Contouring (DC) [16] modify the sampling grid and produce better meshes than MC (as can be seen in Fig. 1). The purpose of DC, however, is to generate adaptive meshes while preserving the sharp features of the isosurface, instead of high-quality meshes. Although the quality of the resultant mesh is visually impressive, many degenerate triangles are still generated. The DC method was later improved by Schaefer et al. [27], which propose an extension that constructs adaptive *manifold* surfaces. Although a manifold mesh is better suited for subsequent mesh smoothing methods, their extension produces meshes with the same quality of DC meshes.

Following recent work [3], [16], [34], we show that the triangle quality can be significantly improved by modifying the grid *before* the mesh generation. We start from the observation that MC cells generate well-shaped triangles in many, but not all, of the possible intersections with a planar




Fig. 3. A point p is associated with an isosurface orientation given by $H(p)$ that passes through p and has normal $p/|p|$. This plane is used to create the function f by setting the value of f at the vertices of the cell to be the signed distance from $H(p)$. The marching method can then be run on the cell to produce a set of triangles $T(p)$.

isosurface. Modifying a cell to produce good triangles would then result in a high-quality mesh that accurately reproduces the isosurface, reducing the need of postprocessing steps.

Unlike our approach, the Warping Cubes approach of Tzeng [34] allows topological changes in the mesh but requires user parameters to decide which triangles should be removed. Volume Warping [3] also deforms the grid where isosurfaces are extracted, but the goal there is to increase sampling in a specific user-defined region, instead of improvements in mesh quality. Recently, Labelle and Shewchuk have proposed a related algorithm to generate tetrahedral meshes that warps a different lattice, the body-centered cubic (BCC) [18]. The approach used in that paper to prove bounds in the dihedral angles might be fruitful to prove bounds in our case.

3 STUDYING THE MESH QUALITY OF MARCHING CUBES

In this section, we show that MC generates bad triangles only in some particular configurations. This basic insight leads to the transformations we propose later. Our analysis is centered on a simple scenario—computing all possible intersections of a planar isosurface against a single grid cell centered at the origin. Observe in Fig. 2 some examples of such intersections, with each intersection color-coded using the radii-ratio triangle quality as done in Fig. 1.

One way to analyze all plane orientations intersecting a cell is to uniquely associate each of them with a single point p in the embedding space. For each plane H , we associate it with point p on the plane that is closest to the origin of the grid cell. We denote this association by $H(p)$. Given $H(p)$, we reconstruct the function f at the cell vertices as the signed distance to $H(p)$. After defining the value of the function f at each vertex, we run MC on the cell to extract the set of triangles $T(p)$ representing the isosurface (Fig. 3). Finally, we compute the average quality of these triangles to create a single scalar quality value for the isosurface parameterized by point p .

We use this mapping from point p to plane $H(p)$, to the set of triangles $T(p)$, and finally to the quality to define a scalar field $Q(p)$ over \mathbb{R}^n . This scalar field can then be visualized to get an estimate on the triangle qualities generated by MC (Fig. 4 shows two slices of Q). A clear structure can be seen in the visualization of the quality field Q . The best triangles are created with points p near the corners of the cell, where the plane cuts off the corner with a




Fig. 4. Two slices of the quality field. Worst quality is obtained when points are closer to the center of the edges.

nearly equilateral triangle (Fig. 2a). These high-quality triangles make up the majority of triangles that are generated, as can be seen in Fig. 4. We observe that the worst quality triangles are generated with points p from the center of the cell out toward the edge centers (Fig. 2b). These points correspond to isosurfaces that are nearly parallel to some of the edges of the cell. This pattern of poor quality triangles is independent of cell shape, being observed with cubes, tetrahedra, and octahedra. This is the key insight: *cell edges nearly parallel to the isosurface are correlated with poor triangles in MC*. Our modification to MC directly addresses this issue.

4 MARCHING CUBES USING EDGE TRANSFORMATIONS

From now on, we focus our discussion only on MC, since extensions to other cell types are straightforward. MC is conceptually a simple algorithm designed to process an implicit function, defined over a 3D grid, and given by its samples $f(i, j, k)$. It operates by *marching* over the *cubes* implicitly defined by the 3D sampling grid, looking at each cell of eight vertices independently, and computing small patches of the overall isosurface by considering the intersection of the function f with a cell.

One way to see how MC operates on a given 3D grid cell is to consider its eight vertices v_1, v_2, \dots, v_8 , and how their scalar values $f(v_1), f(v_2), \dots, f(v_8)$ compare to the desired isovalue λ . For simplicity of presentation, it is convenient to work with the function $f^*(x) = f(x) - \lambda$. Clearly, if all $f^*(v_i)$ share the same sign, the isosurface does not intersect the cell (using trilinear reconstruction). Not taking into account symmetries, there are 2^8 possible configurations. A lookup table of template topologies for these configurations can be created in such a way that the triangles from each cell are joined into a conforming mesh.

For each configuration, the set of *active edges*, i.e., edges where the function f^* has different signs at the endpoints, independently determines both the geometry and topology of the mesh. The topology is determined by the configuration of active edges for each cell. The geometry is determined by the location of the isovalue along each active edge. This independence leaves room to change the geometry of the mesh while keeping the topology intact.

Our modification to MC explores this independence. Some edge transformations do not change the underlying topology: they only change the geometry of the isosurface. In particular, consider an active edge (v, w) of the regular




Fig. 5. Standard MC is depicted on (a). We propose two methods for improving the shapes of the elements by transforming the active edges. In (b), we move the edge endpoints along ∇f , and on (c), we move the endpoints parallel to the isosurface. The intersection between the edge and the isosurface is recomputed after the edge transformation, which guarantees that the resulting mesh adheres to the isosurface defined by the interpolant. Both methods improve the quality of the resulting mesh.

sampling grid illustrated in the left side of Fig. 5. If we move each of its endpoints continuously, with the constraint that $f^*(v)$ and $f^*(w)$ do not change signs, this edge will remain active (two different strategies are shown in Fig. 5). For each active edge, then, we create a “detached” edge (v', w') , which allows shared endpoints to move independently. We will move the detached edges to improve their intersection points with the isosurface. To guarantee that the geometry of the isosurface stays valid (i.e., the mesh remains oriented, does not self-intersect, and is consistent with the initial topology), we need two *conditions* on the transformations:

- An active edge must keep only one intersection with the isosurface throughout the motion of its endpoints. Since each active edge is shared by several cells, we must enforce that each intersection computation, which is done independently by MC, generates the same intersection point. This is trivially done in MC since this computation is performed only once for each active edge.
- The intersection induced by the detached edge (v', w') with the isosurface must not result in triangle flips.

Now, with these conditions in place, we can return to our goal of improving triangle shape. In Section 3, we traced the poor triangle quality to those triangles created when the active edges of a cell are nearly parallel to the isosurface. Here, we will use this observation, together with the necessary conditions discussed above to transform detached versions of the active edges of a cell to make them orthogonal to the isosurface. Below, we show two possible methods for achieving this, followed by an algorithm proposal that combines both approaches.

4.1 Gradient Transformation

The intuition behind our first method is based on moving the vertices of the cell along the direction ∇f away from the isosurface. Since ∇f is orthogonal to the isosurface at points near the isosurface, this will increase the distance between the endpoints, but will not significantly move them tangentially. This results in edges that are more perpendicular to the isosurface and, thus, produce higher quality triangles.

Given an active edge (v, w) , we first generate a detached edge (v', w') . Each vertex of the detached edge will be moved independently. Each vertex is moved away from the




Fig. 6. Edge track computation. A track for a vertex is defined by the projection of the edge to be orthogonal to ∇f . The displacement added to the vertex is half of the track, and the intersection with the isosurface is recomputed.

surface in the direction of the gradient, until one of the two previously described conditions is violated. To avoid long computations, we limit the displacement to half of the original edge length. In other words, we choose a vertex p' that maximizes α in

$$p' = p + \alpha \tilde{\nabla} f(p),$$

where $\tilde{\nabla} f$ denotes the direction of the gradient but pointing away from the isosurface.

As the vertices move, we need to update the induced intersection points. Fig. 7 shows three steps while applying a gradient transform on a spherical isosurface. Even for small displacements, the triangle quality is significantly better. Complete results are shown in Section 5.

4.2 Tangential Transformation

The other way to make the active edges more perpendicular to the isosurface is to move the endpoints orthogonally to ∇f . This method for transforming edges works by moving the edge endpoints parallel to the isosurface. This tangential transform allows tracking of complex isosurface behavior by moving the vertices of active edges in such way that they become *nearly* perpendicular to the isosurface.

The main step of this transform consists in calculating, for each edge vertex, a *vertex track*, tangent to the isosurface and aligned to the edge. We will optimize the vertex positions along the tracks so that the active edges become more orthogonal to the isosurface.

More formally, a vertex track is defined as the projection of the active edge on the isosurface, displaced to start at each edge vertex. If the isosurface is planar, the track is a line segment whose direction is given by

$$\begin{aligned} d(v) &= (I - \hat{\nabla} f(p) \hat{\nabla} f(p)^T)(v - w) \\ d(w) &= (I - \hat{\nabla} f(p) \hat{\nabla} f(p)^T)(w - v), \end{aligned}$$

where v and w are the active edge vertices, and $I - vv^T$ is the matrix that projects a vector onto the orthogonal complement of v . In our implementation, we simplify the track by using a set of line segments that approximates the track. The vertex positions are, then, the one that maximizes $\langle \nabla f(t), v' - w' \rangle$, where

$$\begin{aligned} v' &= v + \alpha d(v) \\ w' &= w + \alpha d(w) \end{aligned}$$

and t is the intersection of the edge (v', w') and the isosurface. If the gradients were all constant, it is easy to show that the optimal vertex position would be with

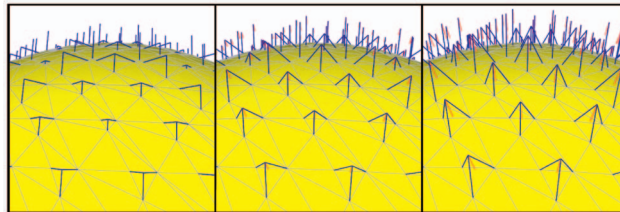


Fig. 7. Three steps of the gradient transform on active edges (blue lines). The procedure iteratively moves endpoints along gradient vectors (dotted red lines), improving the underlying mesh.

$\alpha = 0.5$. In practice, the gradients are not constant, so this might not be true. For the curved case, then, we find the midpoint of a piecewise linear approximation of the track. Fig. 6 illustrates the transformation, while Fig. 8 shows three steps while applying a tangential transform on a spherical isosurface.

4.3 Combining Edge Transformations

Although gradient and tangential transforms are based on different approaches to make edges nearly perpendicular to the isosurface, they generate the same transformed edge if and only if the conditions we impose for the transformations to operate (discussed in the previous section) are satisfied. However, due to practical difficulties in moving points inside the scalar field, we have found that conditions are not simultaneously satisfied for both transformations, and therefore, each transformation tends to improve the overall triangle quality in different situations. That is, in areas where the gradient transform does not significantly change the triangle quality, the tangential transform may exhibit a large improvement, and vice versa. This observation suggested to us that combining the transformations would further improve the overall triangle quality when compared to MC.

We use an iterative approach for combining edge transformations. Since both transformations only alter vertex positions, while keeping intact the mesh topology, we can choose which of the two positions produces a better average quality for the triangles connected to that vertex. Since the choice for a particular vertex might affect its neighbors, this process might require several iterations. For all of our tests, we proceed with this iterative selection without explicitly building the mesh until the overall mesh quality stabilizes (often less than four steps in all our experiments). The results we obtain show that the output mesh has improved quality over each individual mesh produced by using a single edge transformation.




Fig. 8. Three steps of the tangential transform applied to the active edges (blue lines). The procedure iteratively moves the active edge endpoints tangentially to the isosurface (dotted red lines), effectively improving the underlying triangle mesh (gray lines).

TABLE 1

Results Comparing Several Isosurface Extraction Methods: MC, MC Followed by LS, MC Followed by Decimation, DC, the afront Method of Schreiner et al. [28], and Our Method (Macet)

Dataset	Method	radii ratio		Metro frac of bb diagonal
		min	avg	
Engine 256x256x128 isovalue 49.5	MC	0.000448	0.701	-
	LS	0.000546	0.868	0.00168
	Dec	9.50e-7	0.723	4.20e-5
	DC	0.00528	0.829	0.00246
	afront	0.000995	0.938	0.00832
	Macet	0.236	0.775	0.00129
Lobster 301x324x56 isovalue 30.5	MC	0.000970	0.673	-
	LS	0.0314	0.865	0.00245
	Dec	6.80e-7	0.700	8.10e-5
	DC	0.00226	0.833	0.00204
	afront	0.000396	0.939	0.00915
	Macet	0.0757	0.763	0.00163
Bonsai 256x256x256 isovalue 49.5	MC	0.000192	0.680	-
	LS	0.000792	0.862	0.00182
	Dec	7.53e-7	0.709	7.50e-5
	DC	4.52e-5	0.818	0.00265
	afront	1.49e-5	0.935	0.0824
	Macet	0.100	0.767	0.00140
Silicium 98x34x34 isovalue 140.5	MC	0.00298	0.666	-
	LS	0.0847	0.872	0.00685
	Dec	8.96e-5	0.697	0.000370
	DC	0.000336	0.832	0.00758
	afront	0.0283	0.936	0.0321
	Macet	0.0885	0.764	0.00561

Higher radii ratios indicate higher quality meshes. The Metro results show the Hausdorff distance of each mesh to the mesh generated by the original MC.

4.4 Implementation Details

The implementation of edge transformations requires only minor changes to MC. The transformations are performed after finding active edges, as an intermediate stage before the intersection calculation. However, there are important issues related to each transform parameters and the way the interpolator is used. MC assumes a linear interpolation along sampling edges of the grid. Even though this is a hard constraint that might lead to topological inconsistencies in the polygonization, it is also a reasonable assumption for fine sampling granularities. This assumption is no longer valid when we move edge endpoints freely inside the volume, because small movements of a grid vertex will result in nonlinear variations of the scalar field along its incident grid edges (even if we prevent active edges from changing their state upon vertex movements). Therefore, we need a more robust intersection calculation procedure to track the intersection of the edge against the isosurface. We use standard bisection-based root finding procedures, which do not assume linearity along the edge.

The vertex track used in the edge transformations also demands the continuity of the gradient field (C^1 continuity). Since the trilinear interpolation often used with MC is not differentiable at the cell boundaries, there may be unpredictable results. In this work, we use a cubic spline interpolation to reconstruct the derivatives of the scalar field. We also use the corresponding cubic spline to reconstruct the scalar field itself at the sampling grid vertices. This means there might be more than one root in active edges. However, since the topology has already been determined by the lookup into the MC tables and the intersection calculation returns only one vertex, this presents no practical problems.

TABLE 2

Performance Results of MC Compared to Macet

Dataset	MC	Macet	Overhead(%)
Engine	44.4s	81.5s	183
Lobster	27.4s	46.5s	169
Bonsai	86.9s	125.1s	143
Silicium	0.77s	2.6s	341

Notice how the overhead decreases as the data set size increases.

The combination of edge transforms is implemented by keeping two buffers with vertex locations, one for the gradient transforms and another for the tangential transforms. After the transformations have been performed, we use a greedy algorithm to pick which of the two vertices should be used in the final improved version. The vertex pairs are processed one at a time, and the previously computed vertices are used as soon as they are available, in an iterative fashion. We stop the iterative procedure when the mesh quality (the quality of the worst triangle) is no longer improved (maximized) or the number of iterations is greater than the maximum number of iterations allowed. In our experiments, this final pass has small overhead when compared to the full surface extraction.

5 RESULTS

The techniques described in Section 4 have been implemented in C++. Tables 1 and 2 show a summary of our experimental results. All the timings reported were obtained on a Pentium 4 3.0-GHz PC with 2-Gbyte RAM. The edge transforms require additional accesses to the scalar field, in order to determine the vertex displacement. As the data set complexity increases (high curvature or sharp features), the constraints force the vertex movement to happen in smaller steps, which increase the processing time. Notice, however, that the overhead for Macet decreases as the data set size increases, and that for bigger data sets this overhead is at most a factor of two.

Fig. 11 shows the forward Hausdorff distance results produced by Metro, which measure the accuracy of the mesh in relation to the original MC mesh. Although QEM-based decimation produces a mesh very close to MC, it creates more degenerate triangles, as can be seen in Table 1. Fig. 12 shows visual comparisons and triangle quality histograms for the same data. In particular, we highlight a zoomed version of the complete histogram to allow a better evaluation on the distribution of degenerate triangles. Since Fig. 4 provided the motivation for the development of edge transformations, it is important to check the quality of our triangulations under that metric. Fig. 9 shows two slices of the quality field after applying Macet, with a clear reduction of degenerate triangles. Similar improvement is also present in real data sets, as can be seen in Table 1.

6 APPLICATION: CDT FOR MESH GENERATION

One way to illustrate the importance of removing degenerate triangles is given below. `tetgen` is a suite of geometric algorithms that allows computing a Constrained Delaunay Tetrahedralization (CDT) from a piecewise linear complex [31], [32]. Further processing of CDTs often require quality bounds for the generated mesh. `tetgen` allows imposing




Fig. 9. Visualization of two slices of the *Macet* quality field. Notice the reduction (or extinction) of cases where bad-shaped triangles are generated inside the cell. Compare the inset to Fig. 4.

TABLE 3


Comparison of the Tetrahedral Meshes that *tetgen* Creates When Given Both MC and *Macet* Surface Meshes as Input

Dataset	Method	Time	Steiner Points	Output		
				Points	Faces	Tets
Silicium	MC	152.9s	169K	192K	1.45M	646K
	<i>Macet</i>	62.6s	66.8K	87.4K	598K	257K
Lobster	MC	3130.5s	1.72M	1.98M	16.0M	7.27M
	<i>Macet</i>	767.9s	693K	912K	6.72M	2.97M
Engine	<i>Macet</i>	1362.97s	1.27M	1.69M	10.7M	4.57M

tetgen was not able to create a mesh when using the MC input for the engine data set.

quality constraints to the CDT construction, such as a minimum radius edge ratio, and such constraints are enforced by adding Steiner points to the CDT. This process is prone to numerical problems and very sensitive to the triangle quality of the input mesh.

Table 3 shows *tetgen* results, setting minimum radius edge ratio to be 4 (Silicium and Lobster) and 6 (Engine). *tetgen* crashes using the MC mesh for the engine data set but works for *Macet*. Note that to compensate for bad quality triangles, *tetgen* generates many more Steiner points



(a) (b)

Fig. 10. Boundary of the CDT generated by *tetgen* for the (a) MC and (b) *Macet* meshes. Note several degenerate tetrahedra in the MC mesh.

for MC, which increases all output elements generated, as well as processing time. Fig. 10 compares the *tetgen* CDT generated for MC and *Macet* meshes.

7 CONCLUSIONS AND FUTURE WORK

In this work, we have presented edge transforms as a simple and effective way to generate higher quality meshes. We incorporated them into an existing MC implementation, which we called *Macet* that has a relatively modest performance penalty (especially when compared to non-marching algorithms). Since we are able to produce better results, there is much less of a dependence on postprocessing operations to improve the quality of the triangles. Through detailed analysis and comparison against competing strategies, we show how *Macet* generates meshes that improve the

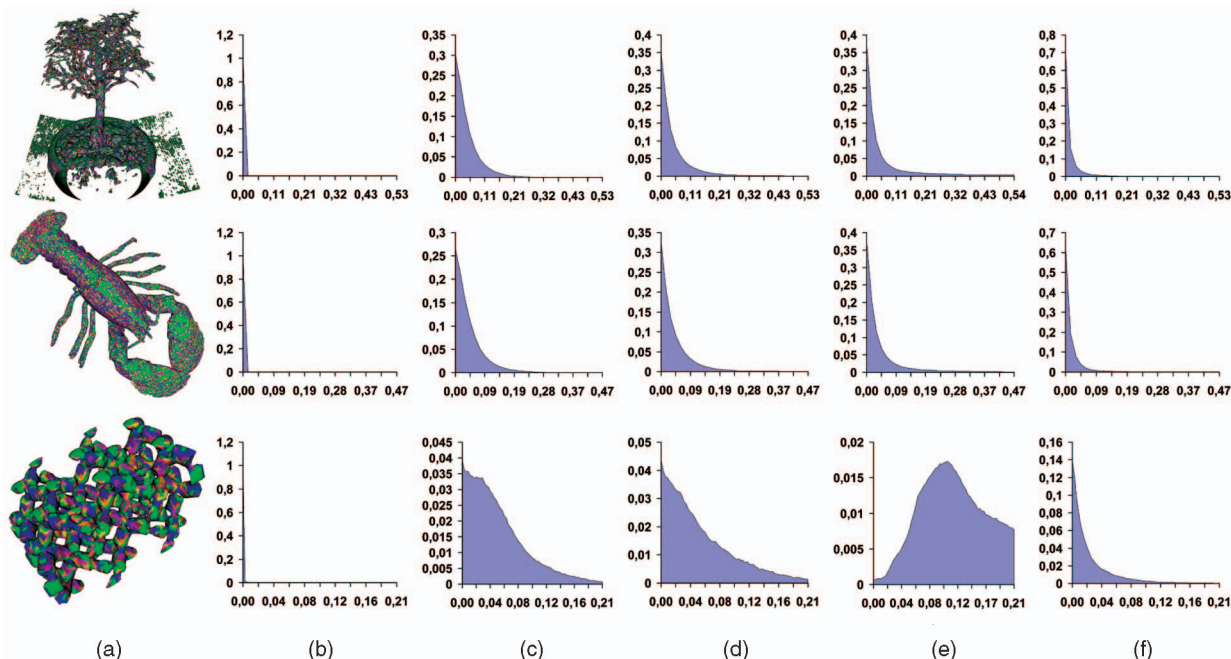


Fig. 11. Metro results: Histogram shows the forward distance between mesh generated by each approach and the original MC mesh. The error histograms for *afront* show worse results than expected because *afront* triangulates a cubic spline surface, while MC uses trilinear interpolation. (a) Data set. (b) Decimation. (c) Laplacian. (d) DC. (e) *afront*. (f) *Macet*.

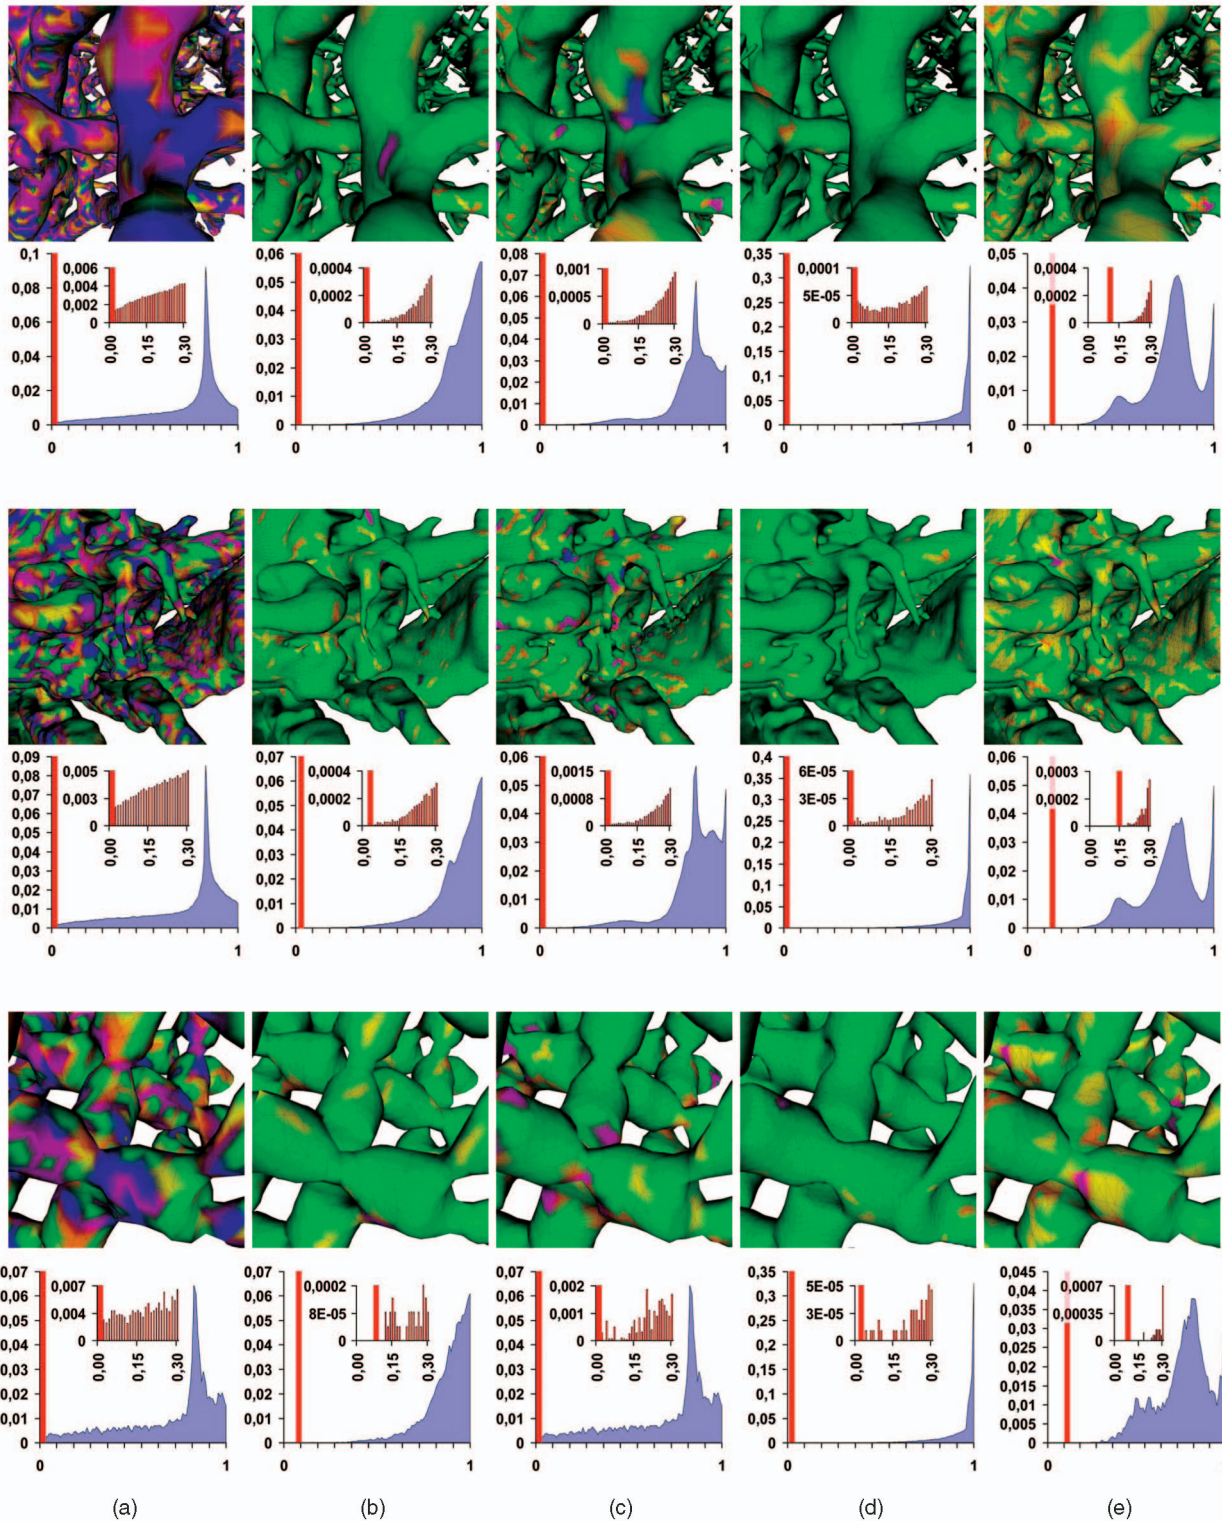


Fig. 12. Comparisons on enlarged sections of Bonsai (rows 1 and 2), Lobster (rows 3 and 4), and Silicium (rows 5 and 6). Color-coded visualization shows that most methods improve the overall quality of triangles. The quality histogram shown below each screenshot shows how Macet consistently reduces degenerate triangles—pay special attention to the zoomed version of the histogram, and compare to the minimum quality results in Table 1. (a) Decimation. (b) Laplacian. (c) DC. (d) afront. (e) Macet.

quality of degenerate triangles and are faithful to the isosurface generated by MC (as measured by Metro). We show how this improvement on quality generated by our method is crucial for further mesh processing, such as the generation of CDTs using tetgen.

For future work, we plan to extend this approach for working with other variants of MC, in particular, variants that provide topological guarantees, are able to keep sharp corners, and adaptive versions of MC [17], [19]. In addition, we want to investigate the effect of warping the cells and

active edges on the reconstructed surface on a theoretical level and, in particular, try to provide formal theoretical bounds on the quality of the triangles.

7.1 Reproducibility

We have made the techniques presented in this paper publicly available, so that the reader can accurately assess the validity of the experiments and algorithms. Every result generated for this paper can be reproduced with open-source software and publicly available data sets. Here is a link to the source code with instructions: <http://www.vistrails.org/index.php/ImprovingMeshQualityOfMarchingCubes>.

ACKNOWLEDGMENTS

The work of Carlos Dietrich is supported by a CAPES scholarship. The works of João Comba and Luciana Nedel are partially supported by CNPq Grants 478445/2004-0, 305947/2005-2, and 306099/2004-7. Carlos Scheidegger, John Schreiner, and Cláudio Silva are partially funded by the US National Science Foundation (NSF) (Grants CCF-0401498, EIA-0323604, OISE-0405402, IIS-0513692, CCF-0528201, and CNS-0751152), the Department of Energy, IBM Faculty Awards (2005, 2006, and 2007), Sandia National Laboratories, and Lawrence Livermore National Laboratory. Carlos Scheidegger is also funded by an IBM PhD fellowship. All data sets used in this work are available at <http://www.volvis.org>. The authors would also like to thank the authors of Metro and tetgen for making their code available as an open source.

REFERENCES

- [1] E. Artzy, G. Frieder, and G.T. Herman, "The Theory, Design, Implementation and Evaluation of a Three-Dimensional Surface Detection Algorithm," *Proc. ACM SIGGRAPH '80*, pp. 2-9, 1980.
- [2] C.L. Bajaj, V. Pascucci, and D.R. Schikore, "Fast Isocontouring for Improved Interactivity," *Proc. Symp. Volume Visualization (VolVis '96)*, pp. 39-46, 1996.
- [3] L. Balmelli, C.J. Morris, G. Taubin, and F. Bernardini, "Volume Warping for Adaptive Isosurface Extraction," *Proc. IEEE Visualization (VIS '02)*, pp. 467-474, 2002.
- [4] H. Carr, T. Möller, and J. Snoeyink, "Artifacts Caused by Simplicial Subdivision," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 2, pp. 231-242, Mar./Apr. 2006.
- [5] H. Carr, T. Theussl, and T. Möller, "Isosurfaces on Optimal Regular Samples," *Proc. Symp. Data Visualisation (VISSYM '03)*, pp. 39-48, 2003.
- [6] L.-S. Chen, G.T. Herman, R.A. Reynolds, and J.K. Udupa, "Surface Shading in the Cuberille Environment," *IEEE Computer Graphics and Applications*, vol. 5, no. 12, pp. 33-43, 1985.
- [7] P. Cignoni, C. Montani, E. Puppo, and R. Scopigno, "Optimal Isosurface Extraction from Irregular Volume Data," *Proc. Symp. Volume Visualization (VolVis '96)*, pp. 31-38, 1996.
- [8] P. Cignoni, C. Rocchini, and R. Scopigno, "Metro: Measuring Error on Simplified Surfaces," *Computer Graphics Forum*, vol. 17, no. 2, pp. 167-174, 1998.
- [9] P. Crossno and E. Angel, "Isosurface Extraction Using Particle Systems," *Proc. IEEE Visualization (VIS '97)*, pp. 495-498, 1997.
- [10] L.H. de Figueiredo, J. de Miranda Gomes, D. Terzopoulos, and L. Velho, "Physically-Based Methods for Polygonization of Implicit Surfaces," *Proc. Graphics Interface (GI '92)*, pp. 250-257, 1992.
- [11] A. Doi and A. Koide, "An Efficient Method of Triangulating Equivalued Surfaces by Using Tetrahedral Cells," *IEICE Trans. Comm. and Electronics Information Systems*, vol. E74, no. 1, pp. 214-224, Jan. 1991.
- [12] M. Garland and P.S. Heckbert, "Surface Simplification Using Quadric Error Metrics," *Proc. ACM SIGGRAPH '97*, pp. 209-216, 1997.
- [13] M. Gavriliu, J. Carranza, D. Breen, and A. Barr, "Fast Extraction of Adaptive Multiresolution Meshes with Guaranteed Properties from Volumetric Data," *Proc. IEEE Visualization (VIS '01)*, pp. 295-302, 2001.
- [14] *The Visualization Handbook*, C.D. Hansen and C.R. Johnson, eds., Academic Press, 2004.
- [15] G.T. Herman and J.K. Udupa, "Display of 3D Digital Images: Computational Foundations and Medical Applications," *IEEE Computer Graphics and Applications*, vol. 3, pp. 39-46, 1983.
- [16] T. Ju, F. Losasso, S. Schaefer, and J. Warren, "Dual Contouring of Hermite Data," *Proc. ACM SIGGRAPH '02*, pp. 339-346, 2002.
- [17] L.P. Kobbelt, M. Botsch, U. Schwanecke, and H.-P. Seidel, "Feature-Sensitive Surface Extraction from Volume Data," *Proc. ACM SIGGRAPH '01*, pp. 57-66, 2001.
- [18] F. Labelle and J. Shewchuk, "Isosurface Stuffing: Fast Tetrahedral Meshes with Good Dihedral Angles," *ACM Trans. Graphics*, vol. 26, 2007.
- [19] T. Lewiner, H. Lopes, A.W. Vieira, and G. Tavares, "Efficient Implementation of Marching Cubes' Cases with Topological Guarantees," *J. Graphics Tools*, vol. 8, no. 2, pp. 1-15, 2003.
- [20] Y. Livnat and X. Tricoche, "Interactive Point-Based Isosurface Extraction," *Proc. IEEE Visualization (VIS '04)*, pp. 457-464, 2004.
- [21] A. Lopes and K. Brodlie, "Improving the Robustness and Accuracy of the Marching Cubes Algorithm for Isosurfacing," *IEEE Trans. Visualization and Computer Graphics*, vol. 9, no. 1, pp. 16-29, Jan.-Mar. 2003.
- [22] W. Lorensen and H. Cline, "Marching Cubes: A High-Resolution 3D Surface Construction Algorithm," *Proc. ACM SIGGRAPH '87*, pp. 163-169, 1987.
- [23] S.V. Matveyev, "Approximation of Isosurface in the Marching Cube: Ambiguity Problem," *Proc. IEEE Visualization (VIS '94)*, pp. 288-292, 1994.
- [24] G.M. Nielson, "Dual Marching Cubes," *Proc. IEEE Visualization (VIS '04)*, pp. 489-496, 2004.
- [25] G.M. Nielson and B. Hamann, "The Asymptotic Decider: Resolving the Ambiguity in Marching Cubes," *Proc. IEEE Visualization (VIS '91)*, pp. 83-91, 1991.
- [26] P.P. Pebay and T.J. Baker, "A Comparison of Triangle Quality Measures," *Proc. 10th Int'l Meshing Roundtable (IMR '01)*, pp. 327-340, 2001.
- [27] S. Schaefer, T. Ju, and J. Warren, "Manifold Dual Contouring," *IEEE Trans. Visualization and Computer Graphics*, vol. 13, no. 3, pp. 610-619, May/June 2007.
- [28] J. Schreiner, C. Scheidegger, and C.T. Silva, "High-Quality Extraction of Isosurfaces from Regular and Irregular Grids," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1205-1212, Sept./Oct. 2006.
- [29] H.-W. Shen, C.D. Hansen, Y. Livnat, and C.R. Johnson, "Isosurfacing in Span Space with Utmost Efficiency (ISSUE)," *Proc. IEEE Visualization (VIS '96)*, pp. 287-294, 1996.
- [30] J.R. Shewchuk, "What is a Good Linear Element? Interpolation, Conditioning, and Quality Measures," *Proc. 11th Int'l Meshing Roundtable (IMR '02)*, pp. 115-126, 2002.
- [31] H. Si, "On Refinement of Constrained Delaunay Tetrahedralizations," *Proc. 15th Int'l Meshing Roundtable (IMR '06)*, pp. 509-528, 2006.
- [32] H. Si and K. Gaertner, "Meshing Piecewise Linear Complexes by Constrained Delaunay Tetrahedralizations," *Proc. 14th Int'l Meshing Roundtable (IMR '05)*, pp. 147-163, 2005.
- [33] G.M. Treece, R.W. Prager, and A.H. Gee, "Regularised Marching Tetrahedra: Improved Iso-Surface Extraction," *Computers and Graphics*, vol. 23, pp. 583-598, 1999.
- [34] L. Tzeng, "Warping Cubes: Better Triangles from Marching Cubes," *Proc. 20th European Workshop Computational Geometry (EWCG '04)*, 2004.
- [35] A. van Gelder and J. Wilhelms, "Topological Considerations in Isosurface Generation," *ACM Trans. Graphics*, vol. 13, no. 4, pp. 337-375, 1994.
- [36] J. Wilhelms and A.V. Gelder, "Octrees for Faster Isosurface Generation," *ACM Trans. Graphics*, vol. 11, no. 3, pp. 201-227, 1992.
- [37] A.P. Witkin and P.S. Heckbert, "Using Particles to Sample and Control Implicit Surfaces," *Proc. ACM SIGGRAPH '94*, pp. 269-277, 1994.



Carlos A. Dietrich received the BS degree in computer science from the Federal University of Santa Maria, Santa Maria, Brazil and the MS degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil. He is currently a fourth-year PhD student working in the Computer Graphics Group at the Federal University of Rio Grande do Sul. His research interests include graphics, visualization, and the use of GPUs as general-purpose processors.



Carlos E. Scheidegger received the BSC degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil. He is a graduate student at the Scientific Computing and Imaging Institute, University of Utah. He is currently an IBM PhD student fellow. His research interests include geometric processing, scientific visualization, and visualization systems.



John Schreiner received the BS degree in computer science and mathematics from the University of Wisconsin-Madison, where his undergraduate thesis was advised by Dr. Michael Gleicher. He is currently a PhD candidate at the University of Utah, studying under Dr. Cláudio Silva. His research interests include meshing algorithms, geometry processing, and scientific visualization.



João L.D. Comba received the BS degree in computer science from the Federal University of Rio Grande do Sul, Porto Alegre, Brazil, the MS degree in computer science from the Federal University of Rio de Janeiro, Rio de Janeiro, and the PhD degree in computer science from Stanford University. He is an associate professor of computer science at the Federal University of Rio Grande do Sul. His main research interests are in graphics, visualization, spatial data structures, and applied computational geometry. His current projects include the development of algorithms for large-scale scientific visualization, data structures for point-based modeling and rendering, and general-purpose computing using graphics hardware. He is a member of the ACM SIGGRAPH.



Luciana P. Nedel received the PhD degree in computer science from the Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland, under the supervision of Prof. Daniel Thalmann. She is with the Federal University of Rio Grande do Sul (UFRGS), where she has been an assistant professor since 2002 and does research in the Computer Graphics Group, Informatics Institute. Her main interests include interactive visualization, virtual humans, computer animation, path planning, tiled displays, and nonconventional interaction techniques and devices. In these fields, she published more than 60 scientific articles in journals, conference proceedings, and edited books.



Cláudio T. Silva received the BS degree in mathematics from the Federal University of Ceara, Fortaleza, Brazil, in 1990 and the PhD degree in computer science from the State University of New York at Stony Brook in 1996. He is an associate professor of computer science and a faculty member of the Scientific Computing and Imaging (SCI) Institute, University of Utah. Before joining the University of Utah in 2003, he worked in industry (IBM and AT&T), government (Sandia and LLNL), and academia (Stony Brook and OGI). He coauthored more than 100 technical papers and eight US patents, primarily in visualization, geometric computing, and related areas. He is an active member of the visualization, graphics, and geometric computing research communities, having served on more than 50 program committees. He is a coeditor of the Visualization Corner of the *IEEE Computing in Science and Engineering*. Previously, he was on the editorial board of the *IEEE Transactions on Visualization and Computer Graphics*. He was a paper cochair for the IEEE Visualization Conference in 2005 and 2006. He received IBM Faculty Awards in 2005, 2006, and 2007 and a Best Paper Award at the IEEE Visualization 2007. He is a member of the ACM, the Eurographics, and the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.