High Quality Surface Mesh Generation for Multi-physics Bio-medical Simulations

Dominik Szczerba, Robert McGregor, and Gábor Székely

Computer Vision Lab, ETH, CH-8092 Zürich, Switzerland
domi@vision.ee.ethz.ch

Abstract. Manual surface reconstruction is still an everyday practice in applications involving complex irregular domains, necessary for modeling biological systems. Rapid development of biomedical imaging and simulation, however, requires automatic computations involving frequent re-meshing of (r)evolving domains that human-driven generation can simply no longer deliver. This bottleneck hinders the development of many applications of high social importance, like computational physiology or computer aided medicine. While many commercial packages offer mesh generation options, these depend on high quality input, which is rarely available when depending on image segmentation results. We propose a simple approach to automatically recover a high quality surface mesh from low-quality, oversampled and possibly non-consistent inputs that are often obtained via 3-D acquisition systems. As opposed to the majority of the established meshing techniques, our procedure is easy to implement and very robust against damaged or partially incomplete, inconsistent or discontinuous inputs.

1 Introduction

Generating a mesh is a necessary pre-condition when obtaining numerical solutions of partial differential equations. An adequate mesh highly impacts both the accuracy and the efficiency of numerical procedures. Since the inception of the finite element method dating back to the middle of the last century, automatic mesh generation with sufficient quality over an arbitrary domain has remained a central topic of intensive research, without being able to reach a fully satisfying solution up to now. Even though a tremendous number of approaches has been described, at the end none of them offers truly interaction-free, general-purpose processing. From the perspective of half a century's work it seems clear that mesh generation, despite its scientific context and origin, bears every sign of artistic sculpturing that escapes automation due to complex decision making and continuous adaptation during the creation process. Interactive meshing is therefore an everyday practice in applications involving complex static domains. For man-made objects, usually emerging from CAD applications, even a few almost-automatic volume meshing methods are available, like the advancing front technique [1], [2]. Such methods usually work very well for geometries represented as either a constructive solid geometry (CSG), non-uniform rational B-spline (NURB) patches, or by an already existing surface mesh of sufficient quality. The rapid development of biomedical imaging and subsequent simulation, however, requires computations involving frequent re-meshing of largely (r)evolving anatomical domains that human-driven

generation can simply no longer deliver. In these applications the previously listed requirements hardly ever hold: a constructive geometry description of an anatomical domain may not make any sense at all, NURB patches are obviously not available, and surface meshes obtained directly from segmentation are generally of very low quality, oversampled, often with broken topology. This serious bottleneck hinders development on many domains of high social importance like computational physiology or computer aided medicine. The computational pipeline required for such applications can be sketched as follows: data acquisition $\rightarrow$ domain segmentation $\rightarrow$ surface representation $\rightarrow$ volume meshing $\rightarrow$ discretization of governing equations $\rightarrow$ solution. The data, acquired via imaging techniques like MRI, CT, US, laser scanning, etc., becomes input to a segmentation procedure (see e.g. [3]). Quite rudimentary algorithms, like the marching cube method [4] are then used to represent the identified objects by volumetric or surface meshes, which are, however, too low quality to be directly used for numerical simulations. High quality surface reconstruction suitable as input to volume meshing algorithms has therefore generated a lot of interest in computer science.

2 Related Work

Several established mesh simplification algorithms are available [5], [6]. The techniques include merging coplanar facets, decimation, re-tiling, local re-meshing, vertex clustering, energy function optimization, wavelet-based methods and many other. They are very useful in downsizing of oversampled input meshes, such methods do not aim, however, to improve input mesh quality.

Automatic surface mesh generation from constructive solid geometry (CSG) or stereo-lithography (STL) representations is largely used in mechanical engineering (e.g., [7], [8], [9]). To derive a high quality surface mesh for such representations, information about the object's boundary elements (faces, vertices) must be first extracted. This works very well for geometries which can be decomposed into simple primitives, however, the required boundary evaluation using e.g. Bézier patches or algebraic surfaces does not always produce the expected results for inherently irregular and unpredictable physiological domains. Similarly, methods based on feature recognition (e.g., [10], [11]) or surface parameterization (e.g., [12], [13]) suffer from the same fundamental limitation.

A robust and almost automatic surface reconstruction can be achieved using knowledge based methods, as demonstrated by e.g. [14]. The approach to blood vessel meshing relies on medial axis extraction and subsequent insertion of boundary faces along the sweeping axes. Such methods can be attractive for specific cases where *a priori* assumptions about the domain's shape can be made, however, are far from what can be called a versatile approach.

The family of algorithms based on computational geometry (e.g., [15], [16], [17]), Delaunay refinement (e.g., [18], [19]) or optimization principles (e.g., [20]) is demonstrated to often provide high quality outputs, in many cases with guarantees on convergence and for lower limits on element quality. These sophisticated methods are generally very sensitive to the consistency, continuity and differentiability of the input. In practice they often fail on real-life biomedical data due to precondition violations. In addition they are usually quite difficult to implement.

Fig. 1. Fragment of a low quality, inconsistent and excessively sized surface mesh obtained from the segmentation of a human lateral ventricle (left), its high quality uniform (middle) and curvature-adapted reconstruction (right). The upper and middle row shows patches of the reconstructed surface with different magnifications, while histograms of the corresponding triangle quality (normalized radii-ratios) are presented by the lower figures.

The number of approaches based on smoothing is too large to provide a comprehensive summary in this overview. The basic idea behind them is to relax the nodal positions such that the elements eventually become evenly sized. Some more advanced versions attempt to eliminate the shrinking effect (e.g. [21]) by e.g. applying a band pass filter or simply projecting the relaxed nodes back to the original boundary. This often works sufficiently well if a topologically correct input mesh is available for the relaxation. In general, these methods are easy to implement but suffer from loss of high frequency features. In addition, they do not offer adaptive refinement options, which are necessary for efficient numerical solution procedures.

The work reported in [22] demonstrates the ability of mass-spring dumper models to generate high quality tetrahedral meshes by relaxing pre-refined and pre-compressed mass-spring connections, which, while expanding, naturally fill the volume constrained by the boundary. The method is efficient and robust, if a (very expensive) consistent volumetric representation of the input is available. Shortcomings include discontinuous refinement and resulting tetrahedralizations not being of the Delaunay type. This method is similar to ours in that it relies on physical and not on mathematical principles. The important differences are that 1) we use a very cheap, possibly inconsistent input represented as a polygonal mesh and not volumetric data; 2) we allow for smooth refinement of the re-generated mesh and 3) we produce topology conforming to the Delaunay condition.

To complete the survey, mesh-free approaches have also to be mentioned, which eliminate the discretization problems inherent to meshes by fundamentally not relying on them. Even though it sounds very attractive, these techniques are at their infancy while still relying on some local topology. There are also somewhat exotic approaches based on fractal geometries, chaos theories or neural networks that will not be discussed here.

3 Method

We propose an iterative procedure as follows: 1. initialization; 2. removal of outside elements (if any); 3. calculation of reaction forces; 4. time integration; 5. correction at boundaries; 6. convergence check, eventual termination. 7. loop back to 2. The procedure is initialized by a set of points filling the bounding box of the input surface (1). The loop begins by the removal of the elements found outside the object (2). In subsequent steps, a physical equilibrium between the remaining points is sought by enforcing a prescribed distance between points (3,4):

$$M \frac{\partial^2 \mathbf{r}}{\partial t^2} + D \frac{\partial \mathbf{r}}{\partial t} = \mathbf{F}_r,$$

with $\mathbf{F}_r$ being the reaction forces and incorporating an internal energy dissipation term regulated by D for stability. Special care is needed next to the object's boundary: the points there must be magnetically snapped to it, allowed only to slide along the surface but not to detach (5). We have tested both mass-spring dumpers and charged particle clouds for the actual implementation and did not observe any qualitative differences. It is only important to note that the structure must be initially preloaded with some potential energy such that expansion forces act already during the early iteration stages. Also, we observed that somewhat better results are achieved using magnetic and not Hookean reaction forces. After some iterations (7), this procedure results in evenly distributed points that very closely follow the surface features (6). Topology provided by a robust Delaunay connectivity algorithm [23] produces a tetrahedral mesh of surprisingly good quality, but not free from some badly conditioned elements (slivers). Their removal in general is far from trivial, but not impossible [24]. In fig. 2 we demonstrate that with an additional effort it is possible to repair and re-use those tetrahedral meshes. A detailed description of how we deal with this issue is, however, beyond the scope of this paper.

Fig. 2. Meshes generated by the algorithm proposed by this paper. **Top:** a surface mesh of a human abdominal aortic bifurcation resulting directly from the segmentation of the underlying MRI data (**left**) and the boundary layer generated by an adaptive volumetric mesh refinement (**right**). **Bottom, left:** the quality histogram (normalized radii-ratios) for the resulting mesh. **Bottom, right:** the convergence rates of subsequent fluid simulations performed on different meshes. There are 3 simulations for different mesh qualities and there are 4 curves for each of them: three momentum residues (solid lines) and a mass residue (dashed line).

However, often only a surface mesh is needed, e.g., for conversion to NURB patches or simply to generate other volume elements then tetrahedrons. In such cases, a surface mesh can be easily extracted by preserving only the triangles having one adjacent tetrahedron and removing all others (fig 1).

The procedure sketched above results in a high quality uniform mesh closely following features of the input surface. However, to simulate large domains, adaptive refinement is often sought to reduce the solution accuracy in the regions out of interest, and this way decrease the load on the underlying PDE solver by reducing the number of finite elements used. Such adaptation can be achieved by modifying point-point interactions to make the regional equilibrium radius correspond to the locally requested element size instead of being a global constant. The refinement map can be arbitrarily

specified by the user or automatically derived from some selected features of the input, like the distance from the boundary or the local curvature. If curvature or any other feature is used, which is only defined on the object's surface, its effect must be propagated into the interior elements e.g. by diffusion:

$$\frac{\partial c}{\partial t} = \text{div}(\text{grad}(c)) + R_c,$$

with c the surface scalar to be distributed. The reaction term, R_c , is necessary to control the magnitude of the resulting gradient of c inside the volume to ensure smooth distribution of the refinement map. There are two possible choices to integrate this process into the discussed procedure: a *multi-grid* or a *coupled* method. In the first case a desired refinement map is provided on an existing, previously generated, uniform mesh and will be continuously interpolated onto a converging refined mesh. This gives very good control of the adaptive refinement process, but requires additional implementation efforts and computing time. The second possible option is to formulate a *coupled* problem, where the distance/curvature based refinement function is evolving simultaneously with the formation of the adaptive mesh, while using its actual state for the numerical representation of the map. This is much easier to implement, much faster to compute and surprisingly stable, but does not allow one to check the refinement map beforehand. Either of these methods will result in a smoothly refined mesh, with an adaptive element size. Note that the technique does not require any pre-existing expensive voxel-based (e.g. level set) description of the input surface. As a consequence, large inputs with very fine details can be resolved on a standard PC, whereas volumetric storage of all necessary transient information about a modest 512^3 domain results in several gigabytes of memory usage.

This method is nearly automatic with minor manual corrections eventually needed to enforce specific information about complex features to be followed for adaptive meshing. The disadvantage of the technique in the first place is its inferior speed inherent to any methods based on solving differential equations. In addition, we are not aware of either mathematically proven convergence behavior or lower limits on element quality. In practice, however, the algorithm always generated meshes with the desired properties and quality, as demonstrated by the examples on the figures.

4 Conclusions

We have presented an automatic technique for high quality reconstruction of oversampled, low-grade and possibly topologically inconsistent surface meshes that can be an attractive alternative to existing approaches. Our procedure is easy to implement, and is very robust against damaged or partially incomplete, inconsistent or discontinuous input. Compared to smoothing-based approaches our technique does not result in any feature loss and naturally offers refinement options (distance to selected regions, curvature, etc.). The method is automatic, manual input may only be necessary if the map used to govern adaptive refinement cannot be calculated from the input data (including the original, non-segmented acquisition) alone. The obtained high quality triangular meshes

can become input to other applications like NURB patching or other volumetric meshing procedures. If tetrahedral meshes are sought, however, the volumetric mesh generated as a by-product of our method could be rectified by a sliver removal procedure.

Fig. 2 demonstrates the effect of volumetric mesh quality on the performance of the subsequent numerical simulations. The human abdominal bifurcation was segmented out of raw image data and three meshes with different average qualities were produced using the presented method. The meshes were subsequently used in numerical flow computations using a momentum-pressure PDE solver. Using tetrahedrons of average quality around 0.9 resulted in 3 times faster convergence when compared to average quality around 0.65. In case when only a surface mesh is sought for further processing, it is crucial to note that its quality will strongly influence the quality of the results, as is the case with procedures like the advancing front method.

Obviously, the presented method is not limited to bio-medical applications. Due to its strong feature preserving nature it can be used in general engineering applications where e.g. sharp edges or point-like singularities need to be preserved. The major disadvantages of our technique are inferior speed and missing theoretical bounds on element quality. However, we did not detect any quality defects on the numerous meshes we have generated up to now.

Acknowledgments

This work is a part of the Swiss National Center of Competence in Research on Computer Aided and Image Guided Medical Interventions (NCCR Co-Me), supported by the Swiss National Science Foundation.

References

1. Schöberl, J.: Netgen an advancing front 2d/3d-mesh generator based on abstract rules. *Computing and Visualization in Science* **V1**(1) (1997) 41–52
2. Tristano, J., Owen, S., Canann, S.: Advancing front surface mesh generation in parametric space using a reimannian surface definition (1998)
3. Yushkevich, P.A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage* (2006)
4. Lorensen, W.E., Cline, H.E.: Marching cubes: a high resolution 3d surface construction algorithm. *Computer Graphics (ACM)* **21**(4) (1987) 163–169
5. Kim, S.J., Kim, C.H., Levin, D.: Surface simplification using a discrete curvature norm. *Computers & Graphics* **26**(5) (2002) 657–663
6. Balmelli, L., Liebling, T., Vetterli, M.: Computational analysis of mesh simplification using global error. *Computational Geometry* **25**(3) (2003) 171–196
7. Boender, E., Bronsvort, W.F., Post, F.H.: Finite-element mesh generation from constructive-solid-geometry models. *Computer-Aided Design* **26**(5) (1994) 379–392
8. Bechet, E., Cuilliere, J.C., Trochu, F.: Generation of a finite element mesh from stereolithography (stl) files. *Computer-Aided Design* **34**(1) (2002) 1–17
9. Ryppl, D., Bittnar, Z.: Generation of computational surface meshes of stl models. *Journal of Computational and Applied Mathematics* **192**(1) (2006) 148–151

10. Cuilliere, J.C., Maranzana, R.: Automatic and a priori refinement of three-dimensional meshes based on feature recognition techniques. *Advances in Engineering Software* **30**(8) (1999) 563–573
11. Chappuis, C., Rassineux, A., Bretkopf, P., Villon, P.: Improving surface meshing from discrete data by feature recognition. *Engineering with Computers* **V20**(3) (2004) 202–209
12. Lee, C.K., Hobbs, R.E.: Automatic adaptive finite element mesh generation over rational b-spline surfaces. *Computers & Structures* **69**(5) (1998) 577–608
13. Hormann, K., Labsik, U., Greiner, G.: Remeshing triangulated surfaces with optimal parameterizations. *Computer-Aided Design* **33**(11) (2001) 779–788
14. Antiga, L., Ene-Iordache, B., Caverni, L., Paolo Cornalba, G., Remuzzi, A.: Geometric reconstruction for computational mesh generation of arterial bifurcations from ct angiography. *Computerized Medical Imaging and Graphics* **26**(4) (2002) 227–235
15. Frey, P.J., Borouchaki, H.: Geometric surface mesh optimization. *Computing and Visualization in Science* **V1**(3) (1998) 113–121
16. Garimella, R.V., Shashkov, M.J., Knupp, P.M.: Triangular and quadrilateral surface mesh quality optimization using local parametrization. *Computer Methods in Applied Mechanics and Engineering* **193**(9–11) (2004) 913–928
17. Montenegro, R., Escobar, J., Montero, G., Rodríguez, E.: Quality Improvement of Surface Triangulations. (2006)
18. Oudot, S., Rineau, L., Yvinec, M.: Meshing volumes bounded by smooth surfaces. In: *Proc. 14th International Meshing Roundtable*. (2005) 203–219 meshing the volume bounded by the surfaces.
19. Dey, T., Li, G., Ray, T.: Polygonal Surface Remeshing with Delaunay Refinement. (2006)
20. Escobar, J.M., Rodríguez, E., Montenegro, R., Montero, G., Gonzalez-Yuste, J.M.: Simultaneous untangling and smoothing of tetrahedral meshes. *Computer Methods in Applied Mechanics and Engineering* **192**(25) (2003) 2775–2787
21. Cebal, J.R., Löhner, R.: From medical images to anatomically accurate finite element grids. *International Journal for Numerical Methods in Engineering* **51**(8) (2001) 985–1008
22. Marroquim, R., Cavalcanti, P.R., Esperança, C., Velho, L.: Adaptive multi-resolution triangulations based on physical compression. *Communications in Numerical Methods in Engineering* **21**(10) (2005) 571–580
23. Barber, C.B., Dobkin, D.P., Huhdanpaa, H.: The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software* **22**(4) (1996) 469–483
24. Labelle, F.: Sliver removal by lattice refinement. In: *SCG '06: Proceedings of the twenty-second annual symposium on Computational geometry*, New York, NY, USA, ACM Press (2006) 347–356