



Swansea University
Prifysgol Abertawe



Cronfa - Swansea University Open Access Repository

This is an author produced version of a paper published in :

Data Visualization, Proceedings of the 6th Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (VisSym 2004)

Cronfa URL for this paper:

<http://cronfa.swan.ac.uk/Record/cronfa24630>

Conference contribution :

Laramee, R., Schneider, J. & Hauser, H. (2004). *Texture-Based Flow Visualization on Isosurfaces from Computational Fluid Dynamics*. *Data Visualization, Proceedings of the 6th Joint EUROGRAPHICS - IEEE TCVG Symposium on Visualization (VisSym 2004)*,

<http://dx.doi.org/10.2312/VisSym/VisSym04/085-090>

This article is brought to you by Swansea University. Any person downloading material is agreeing to abide by the terms of the repository licence. Authors are personally responsible for adhering to publisher restrictions or conditions. When uploading content they are required to comply with their publisher agreement and the SHERPA RoMEO database to judge whether or not it is copyright safe to add this version of the paper to this repository.

<http://www.swansea.ac.uk/iss/researchsupport/cronfa-support/>

Texture-Based Flow Visualization on Isosurfaces from Computational Fluid Dynamics

Robert S. Laramée,¹ Jürgen Schneider,² and Helwig Hauser¹

¹VRVis Research Center, Vienna, Austria, {Laramée,Hauser}@VRVis.at, <http://www.VRVis.at>

²AVL, Graz, Austria, Juergen.Schneider@avl.com, <http://www.avl.com>

Abstract

Isosurfacing, by itself, is a common visualization technique for investigating 3D vector fields. Applying texture-based flow visualization techniques to isosurfaces provides engineers with even more insight into the characteristics of 3D vector fields. We apply a method for producing dense, texture-based representations of flow on isosurfaces. It combines two well know scientific visualization techniques, namely iso-surfacing and texture-based flow visualization, into a useful hybrid approach. The method is fast and can generate dense representations of flow on isosurfaces with high spatio-temporal correlation at 60 frames per second. The method is applied in the context of CFD simulation data, namely, the investigation of a common swirl flow pattern and the visualization of blood flow.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation; I.3.7 [Computer Graphics] Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture; [Simulation and Modeling]: Simulation Output Analysis

1. Introduction -The Case of Swirl Flow

At the VRVis Research Center we collaborate with AVL (www.avl.com) in order to provide visualization solutions for analysis of their CFD simulation result data. AVL's own engineers as well as engineers at industry affiliates use visualization software to analyze and evaluate the results of their automotive design and simulation.

For many of the automotive components that undergo evaluation, there is an ideal pattern of flow the engineers are trying to create. Figure 1 illustrates the swirl motion of fluid flow in a combustion chamber from a diesel engine. In order to generate swirl motion, fluid enters the combustion chamber from the intake ports. Later on in the engine cycle, the kinetic energy associated with this swirl motion is used to generate turbulence for mixing of fresh oxygen into the fluid. The more turbulence generated, the better the mixture of air and diesel fuel, and thus the better the combustion itself. Ideally, enough turbulent mixing is generated such that 100% of the fuel is burned.

Since it is the swirling flow that is used to generate turbulence, the swirl should be maximized in order to maximize




Figure 1: The swirling motion of flow in the combustion chamber of a diesel engine (side view). The intake ports at the top provide the tangential component of the flow necessary for swirl.

turbulence. From the point of view of the mechanical engineers designing the intake ports, increased swirl flow leads to some beneficial conditions: (1) improved mixture preparation, i.e., more fuel contact with oxygen, (2) a higher EGR

(Exhaust Gas Ratio) which means a decrease in fuel consumption, and (3) lower emissions. However, too much swirl displaces the flame used to ignite the fuel. As such, a balance must be achieved between (1) generating enough swirl flow in order to create turbulence and (2) not displacing the flame used to ignite the flow.

Some routine questions that a mechanical engineer may ask when investigating swirl flow are: (1) Can visualization provide insight into or verify the characteristic shape(s) or behavior of the flow? (2) What tool(s) can help to visualize the swirl flow pattern? and (3) Where in the combustion chamber is the swirl flow pattern *not* being met?

1.1. Isosurfaces

Engineers often start an analysis of CFD simulation data using techniques to visualize the flow at the surface in order to get a first impression of the simulation results. The next logical step is to investigate the properties of the flow *inside* the volume. Slices are commonly used, but visualizing 3D characteristics of the flow like swirl can be difficult with 2D slices. Engineers are interested in visualization techniques that provide insight into the spatial dimension orthogonal to the slice as well.

Isosurfaces are a visualization tool used routinely by mechanical engineers to investigate the properties of the flow inside a 3D volume. The shape of an isosurface can give the engineer insight into its 3D characteristics. One reason engineers use isosurfaces, as opposed to say streamsurfaces, is because they are so common. They feel very comfortable with isosurfaces because, like isolines, they are very familiar. The mechanical engineers we spoke to are not as familiar with the notion of a streamsurface and even less its interpretation.

Figure 2 shows an isosurface in the combustion chamber of the data set in Figure 1. The engineer can see that the flow has some of the swirling orientation that they are looking for. However, what is missing from Figure 2 is a clear indication of flow direction, e.g., the upstream and downstream nature of the flow. In particular, it is not obvious where the flow does *not* follow the ideal swirl pattern that the combustion chamber should encapsulate.

1.2. Applying Texture-Based Flow Visualization

Applying texture-based flow visualization techniques to such isosurfaces provides engineers even more insight into the characteristics of 3D vector fields. And this has become a feasible option only recently. We apply the method of Laramée et al. [5] for producing dense, texture-based representations of flow on isosurfaces. The result is a combination of two well know scientific visualization techniques, namely iso-surfacing and texture-based flow visualization, into a useful hybrid approach. Our application is a versatile visualization technique with the following characteristics:




Figure 2: *The swirling motion of flow in the combustion chamber of a diesel engine (side view) as illustrated by an isosurface. This is a velocity isosurface with an isovalue of 5.0 m/s. Any CFD attribute can be mapped to hue.*

(1) generates a dense representation of flow on adaptive resolution isosurfaces, (2) visualizes flow on complex isosurfaces composed of polygons whose number is on the order of 200,000 or more, (3) visualizes flow independent of the isosurface mesh's complexity and resolution, (4) supports user-interaction such as rotation, translation, and zooming always maintaining a constant, high spatial resolution, and (5) produces fast animations, realizing up to 60 frames per second.

We note that the method must be applicable to adaptive resolution isosurfaces like that of Figure 4 (color plate). Note that many of the polygons in Figure 4 cover less than one pixel. The isosurface algorithm used here is an extension of the Marching Cubes [6] and Marching Tetrahedra [11] algorithms that takes into account more cell types. It handles adaptive resolution meshes in the same spirit as Laramée and Bergeron [4]. The rest of the paper is organized as follows: in Section 2 we discuss related work, Section 3 reviews the research that the work here is built upon. Section 4 details texture-based flow visualization on isosurfaces from CFD. Results and conclusions, including a discussion of the user questions, are presented in Section 5.

2. Related Work

The dense, texture-based category of fluid flow visualization techniques generally started out with Spot Noise [12] and LIC [2]. The main advantage of the texture-based class of algorithms is their *complete* depiction of the flow field while their primary drawback is, in general, the computational time required to produce the results. Previous research with a focus on representations of the vector field on boundary surfaces is generally restricted to steady-state flow on

simple boundary geometries such as spheres. This is mainly due to the prohibitive computational time required.

Some approaches are limited to curvilinear surfaces. Spot Noise is applied to parameterized surfaces by Van Wijk [12]. Forssell and Cohen [3] extend LIC to curvilinear surfaces with animation techniques and add magnitude and directional information. Stalling [10] provides a helpful overview of LIC techniques applied to surfaces. In particular, a useful comparison of parameterized vs. non-parameterized surfaces is given. However, isosurfaces, like those generated by the Marching Cubes algorithm [6] are not, in general, parameterizable.

Battke et al. [1] describe an extension of LIC for arbitrary surfaces in 3D. The method works by tessellating a given surface with triangles. The triangles are packed (or tiled) into texture memory and a local LIC texture is computed for each triangle. The results are limited to relatively small, (1,600-4,000 polygons) simple surfaces composed of equilateral triangles however. Mao et al. [7] extend LIC for visualizing a vector field on an arbitrary surface in 3D. The convolution of a 3D white noise image, with filter kernels defined along the local streamlines, is performed only at visible ray-surface intersections. However ray tracing is, in general, costly.

Our review of the literature [8, 9] indicates that only recently have dense, texture-based flow visualization techniques on surfaces become more feasible. Perhaps this is one reason why we have not seen them applied to isosurfaces. Laramee et al. [5] and Van Wijk [13] both present texture-based flow visualization for boundary surfaces. In both methods, texture is advected in the direction of the flow at fast frame rates. Also, both methods are suited for the visualization of unsteady flow on surfaces. The application we present here builds on the work of Laramee et al. [5]. In what follows, we describe why this method is well suited for the case of isosurfaces and discuss its utility.

3. Method Background

In order to understand how and why we apply texture-based flow visualization to isosurfaces, we briefly outline the method background. In brief, the algorithm presented by Laramee et al. [5] simplifies the problem by confining the advection of texture properties to image space. The surface geometry is projected to image space and then a series of textures are mapped, blended, and advected. This order of operations eliminates portions of the surface hidden from the viewer. The previous method for visualization of flow on surfaces is comprised of the following procedure: (1) project the vector field to the image plane, (2) detect geometric edge discontinuities, (3) compute advected texture coordinates, (4) advect the image, (5) inject and blend in noise, (6) blend additional noise along geometric edge discontinuities, and (7) apply shading and other additional graphics. Steps 1-7 of the pipeline are necessary for the dynamic cases of changes

to the isovalue, time-dependent geometry, rotation, translation, and scaling, and only a subset is needed for the static cases (Steps 4-7) involving no changes to the view-point or isovalue. Each stage is described in more detail in previous research [5].

4. Texture-Based Flow Visualization on Isosurfaces

Here we describe the way in which we apply the method described in Section 3 to isosurfaces. Specifically, we describe ways to address: (1) the normal component of the flow to the isosurface, (2) perceptual challenges associated with viewing flow on isosurfaces, (3) issues related to resampling the vector field, and (4) some implementation details.

4.1. Applying a Normal Mask

When visualizing flow on normal boundary surfaces the direction of the flow generally coincides with the surface itself. As the flow approaches the boundary, it is not allowed to pass through and is pushed in a tangential direction, i.e., it can be described as surface aligned flow. However, in the case of isosurfaces this is no longer true. The flow at an isosurface can sometimes exhibit a strong flow that is normal to the surface, e.g., cross-surface flow. The same also holds true for the case of arbitrary clipping geometries. Simply advecting texture properties according to the vector field projected onto the isosurface could be considered misleading. Is there a way in which this cross-surface component of the flow can be incorporated into the result visualization?

Battke et al [1], who applied LIC to surfaces, address this problem by varying the length of the convolution filter according to the magnitude of the vector component tangential to the surface. In areas where the vector field is oriented almost perpendicular to the surface only very little smearing of the texture occurs, i.e., the input noise is visible instead of a convolved texture. Our approach is required to be consistent with the visualization of flow on boundary surfaces. When we apply texture-based flow visualization to boundary surfaces, the amount of texture-smearing indicates velocity magnitude, i.e., texture is smeared into longer streaks in areas of higher velocity magnitude. We don't want to change the semantic interpretation of smearing for isosurfaces.

We propose an idea inspired by the well known velocity mask, namely, a *normal mask*. A velocity mask can be used to dim or highlight high frequency noise in low velocity regions. Whereas, a normal mask can be used to dim regions of the vector field that have strong cross-flow component to the isosurface. We define the normal mask as: $\beta = (\bar{\mathbf{v}} \cdot \mathbf{n})^m$ where β increases as a function of the product of the velocity, \mathbf{v} , and normal vector to the surface, \mathbf{n} , at that point. Here, m is an arbitrary number. In practice, it is typically around unity giving the opacity a nearly linear behavior. In our case, the image overlay becomes more opaque in regions with a strong cross-flow component and more transparent in areas of highly tangential velocity. With the normal mask enabled,

the viewer's attention is drawn away from areas of strong cross-flow component, and towards areas of high tangential velocity. However, the texture properties are still advected according to the velocity vectors projected onto the isosurface. Note that if we encode the normal mask as opacity, another simulation attribute can be mapped to hue. In our application this is a requirement for consistency.

4.2. Normal Mask Implementation

We can integrate the implementation of the normal mask with very little overhead by taking advantage of the graphics hardware. If we look closely at the construction of the velocity image, we note that the R , G , and B image channels are used to encode the x , y , and z values of the vector components at each vertex defining the surface. This leaves the alpha channel as a free parameter in the velocity image construction. In order to implement the normal mask, we simply store β into the alpha channel of the framebuffer at the same time we are storing the x , y , and z vector components -when rendering the velocity image. And when reading back the image buffer, reading the alpha component in addition to the R , G , and B component comes at very little overhead. Some results of applying this normal mask to an isosurface are shown in Figure 3. We can see that the flow at the isosurface just below the intake port in the foreground (in white) has a strong normal component to the isosurface. The higher frequency texture in this region is difficult to see. Figure 6 (color plate) shows a close-up for increased clarity of exposition. Note also that we have chosen a simpler color scale in this case to reduce the visual complexity of the result. We find that using a full range of hue for the color mapping (like in Figure 2) in combination with variable opacity for the normal mask is visually complex. So we provide the option of trading off some complexity in the color map while applying the normal mask.

4.3. Perceptual Issues

Figure 6 (color plate) shows a close-up view of a velocity isosurface with texture-based flow visualization applied. One perceptual problem with the result is that of occlusion. There is more structure to this isosurface than we can see. Perceptual problems such as occlusion and visual complexity are common to generally all 3D visualizations. One way we addressed this is by implementing an interactive clipping plane. The clipping plane allows the user to see occluded parts of the isosurface by removing sub-sets of the geometry on one side of the plane, in this example, the side closer to the viewer. Again, the users are interested in cutting planes because of their familiarity. Figure 5 (color plate) shows another view of the isosurface with a clipping plane being used. New structures in the isosurface are revealed, namely the structure resulting from flow around an intake port valve.

Of course another alternative is for the engineer to take a 2D slice through the volume, rather than creating an isosurface. This is essentially trading off dimensionality in order to

reduce perceptual problems. In our application, the user has both options. Another option for the user in our application is the use of arbitrary clipping geometries. For example, the user can define a clipping geometry in the shape of a sphere or cylinder and apply the texture-based flow visualization. Again however, this is a trade-off. We may gain by lowering visual complexity and occlusion, but we lose some information about the behavior of the flow, namely, that visualized by the isosurface.

4.4. A Resampling Point of View

The reason this texture-advection method is faster than previous texture-based methods on surfaces is because the injection, blending, and advection of noise textures is done in image space. The key to transforming the three-dimensional nature of textures on surfaces to a two-dimensional problem is via the projection of the vector field to image space.

The vector field from the isosurface is projected onto image space via the velocity image described in Section 3. Then, the image is warped according to a regular, rectilinear mesh. By distorting the image according to the projected velocity vectors located at the grid intersections, we are implicitly resampling the vector field. This resampling implies some consequences. Unlike the nature of boundary surfaces, isosurfaces may contain very small, disconnected pieces. In some cases these pieces may only cover a few pixels. This implies that we need a high vector field resampling rate when advecting the textures in image space. In other words, the sampling-to-pixel ratio should not be too small, e.g., sampling at every pixel or every other pixel. In order to handle this, we give the user the option of different advection grid resolutions. In our implementation, the highest grid resolution samples the vector field at every pixel, while the second highest advection grid resolution samples the vector field at every other pixel.

Another reason we give the user control over the resolution of the advection grid is because we want to retain the advantages obtained by decoupling the original surface mesh with the mesh used to advect the textures. This decoupling prevents computation on those polygons whose area covers less than one pixel. And in the case of Figure 4 there are thousands of such polygons. We note also that zooming in on a surface implicitly increases the sampling rate of the vector field because more of the image is spread out while the resolution of the advection grid at the same time remains the same.

The fact that an isosurface may contain many small, disconnected pieces also implies that we need a high frequency texture in the spatial domain. In our implementation, we give the user control over the spatial frequency of the noise injection. Using a high spatial frequency allows for the visualization of flow on even very small, disconnected pieces of an isosurface.




Figure 3: (left) a velocity isosurface of value 5.0 m/s with a CFD simulation attribute mapped to hue and texture-based flow visualization (right) texture-based flow visualization on the isosurface combined with a normal mask. A close-up is shown in Figure 6.

5. Results and Discussion

If we take a closer look at Figures 3 and 5 (color plate), we can see that the texture-based flow visualization provides additional insight into the behavior of the flow. One of the questions that the engineer poses is: Where in the volume is the ideal swirl flow pattern *not* being met? Within the texture, we can see that the ideal swirl flow pattern is not being met in just below the intake ports themselves. Namely, we can see that two areas of the flow are working *against* each other just beneath the actual boundary surface of the combustion chamber. This is shown more clearly in a close-up in Figure 6 (color plate). The normal mask in Figure 6 highlights the boundary between this destructive flow pattern. This is contrasted with only the isosurface itself (Figure 2) where area destructive flow is not obvious. The destructive flow pattern is made even more obvious in an animation of the flow. Supplementary MPEG animations and images of the results can be found at: <http://www.vrvis.at/ar3/pr2/VisSym04/>

Figure 7, top (color plate), shows the intersection of three blood vessels. The larger vessel on the right brings in blood and distributes it to two small vessels on the left. An abnormal pocket, e.g., an aneurysm, has developed at the junction of the three vessels. The observer may be interested to investigate the inside of the pocket to see the resulting blood flow pattern. If we look at the blood flow at the surface, as in Figure 7, top, we see mostly noise. The velocity of the blood flow at the surface of the pocket is moving very slow relative to the vessel surfaces. Figure 7, bottom, shows a velocity isosurface (of 0.04 m/s) inside the volume with texture-

based flow visualization applied. Shown clearly is the recirculation zone in the pocket with blood flowing upstream (in the opposing direction). This second example was chosen in an effort to support our claim that the hybrid approach of texture-based flow visualization on isosurfaces can be useful, not only in the automotive domain.

6. Performance

Performance was evaluated on a PC with an Nvidia 980XGL Quadro graphics card, with a 2.8 GHz dual-processor and 1 GB of RAM. The performance times reported in Table 1 support interactive exploration of flow on isosurfaces. This is important for the case of changing isovalues. When the user changes the isovalue, texture updates only require a fraction of a second. And the transition is generally coherent because each frame is blended with the previous frame [5].

The first time reported in the FPS column is for the static case, i.e., the absence of changes to the view point. The times shown in parenthesis indicate the dynamic cases of interactive zooming, rotation, and translation of the view point. The reported times are about three times faster than those reported by Laramee et al. [5]. This is mainly due to the updated hardware used for the evaluation and improvements to the implementation.

Normal mask construction does not introduce significant overhead since it is easily built into the advection process itself. For example, the isosurface shown in Figure 6 is composed of 243K polygons. In the static case, the normal mask has no effect on frame rates. They are the same as those listed in Table 1. In the dynamic case using a 128^2 advection mesh, the frame rate drops from 3.3 to 3.0 FPS with the addition of the normal mask.

num. polygons	advection mesh resolution	FPS
10K	64 ²	64 (35)
	128 ²	64 (18)
	256 ²	32 (8)
	512 ²	15 (2.3)
48K	64 ²	64 (13)
	128 ²	64 (10)
	256 ²	32 (6)
	512 ²	15 (2)
221K	64 ²	64 (4)
	128 ²	64 (4)
	256 ²	32 (2.9)
	512 ²	13 (1.5)

Table 1: Sample frame rates for the visualization algorithm.

7. Conclusions and Future Work

We have shown how the image space based texture-advection technique of Laramée et al. [5] can be applied to isosurfaces. Isosurfaces provide information into the 3D characteristics of flow that 2D slices and boundary surfaces alone cannot. We have shown that adding the texture-based representation of flow to isosurfaces can give the engineer new insight into the behaviour of swirl flow when examining CFD simulation data. We have also applied the method to the visualization of blood flow. Furthermore, the method is fast and supports user interaction such as zooming, rotation, and translation.

Future work can go in many directions including visualization of texture-based flow visualization on time-dependent isosurfaces, streamsurfaces, and unsteady 3D flow. Challenges will include both interactive performance time and perceptual issues. Future work also includes the application of more specialized programmable graphics hardware features in the manner of Weiskopf et al. [14]

The author(s) thank all those who have contributed to this research including AVL (www.avl.com) and Kplus (www.kplus.at). We extend a special thanks to Markus Hadwiger of VRVis (www.vrvis.at) for his technical support. We also thank Jarke J. van Wijk for his encouraging dialog and Helmut Doleisch for helping to prepare the final manuscript. All CFD simulation data shown in this paper has been provided courtesy of AVL.

References

[1] H. Battke, D. Stalling, and H.C. Hege. Fast Line Integral Convolution for Arbitrary Surfaces in 3D. In *Visualization and Mathematics*, pages 181–195. Springer-Verlag, Heidelberg, 1997.

[2] B. Cabral and L. C. Leedom. Imaging Vector Fields Using Line Integral Convolution. In *Proceedings of ACM SIGGRAPH*

1993, Annual Conference Series, pages 263–272. ACM Press / ACM SIGGRAPH, 1993.

[3] L. K. Forssell and S. D. Cohen. Using Line Integral Convolution for Flow Visualization: Curvilinear Grids, Variable-Speed Animation, and Unsteady Flows. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):133–141, June 1995.

[4] R. S. Laramée and R. D. Bergeron. An Isosurface Continuity Algorithm for Super Adaptive Resolution Data. In John Vince and Rae Earnshaw, editors, *Advances in Modelling, Animation, and Rendering: Computer Graphics International (CGI 2002)*, pages 215–237, Bradford, UK, July 1-5 2002. Computer Graphics Society, Springer.

[5] R. S. Laramée, B. Jobard., and H. Hauser. Image Space Based Visualization of Unsteady Flow on Surfaces. In *Proceedings IEEE Visualization '03*, pages 131–138. IEEE Computer Society, 2003.

[6] W. E. Lorensen and H. E. Cline. Marching Cubes: a High Resolution 3D Surface Construction Algorithm. In *Computer Graphics (Proceedings of ACM SIGGRAPH 87)*, pages 163–170. ACM, July 27–31 1987.

[7] X. Mao, M. Kikukawa, N. Fujita, and A. Imamiya. Line Integral Convolution for 3D Surfaces. In *Visualization in Scientific Computing '97. Proceedings of the Eurographics Workshop*, pages 57–70. Eurographics, 1997.

[8] F. H. Post, B. Vrolijk, H. Hauser, R. S. Laramée, and H. Doleisch. Feature Extraction and Visualization of Flow Fields. In *Eurographics 2002 State-of-the-Art Reports*, pages 69–100, Saarbrücken Germany, 2–6 September 2002. The Eurographics Association.

[9] F. H. Post, B. Vrolijk, H. Hauser, R. S. Laramée, and H. Doleisch. The State of the Art in Flow Visualization: Feature Extraction and Tracking. *Computer Graphics Forum*, 22(4):775–792, 2003.

[10] D. Stalling. LIC on Surfaces. In *Texture Synthesis with Line Integral Convolution*, pages 51–64. ACM SIGGRAPH 97, International Conference on Computer Graphics and Interactive Techniques, 1997.

[11] G. M. Treece, R. W. Prager, and A. H. Gee. Regularised Marching Tetrahedra: Improved Isosurface Extraction. *Computers and Graphics*, 23(4):583–598, August 1999.

[12] J. J. van Wijk. Spot noise-Texture Synthesis for Data Visualization. In Thomas W. Sederberg, editor, *Computer Graphics (Proceedings of ACM SIGGRAPH 91)*, volume 25, pages 309–318. ACM, 1991.

[13] J. J. van Wijk. Image Based Flow Visualization. *ACM Transactions on Graphics*, 21(3):745–754, 2002.

[14] D. Weiskopf, G. Erlebacher, M. Hopf, and T. Ertl. Hardware-Accelerated Lagrangian-Eulerian Texture Advection for 2D Flow Visualizations. In *Proceedings of the Vision Modeling and Visualization Conference 2002 (VMV-01)*, pages 439–446, November 21–23 2002.




Figure 4: (color plate) A close-up, wire-frame view of the isosurface from Figure 2. The algorithm we describe must be applicable to adaptive resolution isosurface meshes.




Figure 6: (color plate) A close-up of a velocity isosurface from Figure 3: with texture-based flow visualization and a normal mask applied. With the texture advection on the isosurface, it is clear that the ideal swirl flow pattern is not exhibited in this region.




Figure 5: (color plate) A close up view of the same isosurface shown in Figure 3 using a clipping plane tangent to the view-point in order to reveal occluded isosurface structures.




Figure 7: (color plate) (top) the intersection of three blood vessels. An abnormal pocket has formed at the junction. (bottom) a velocity isosurface of value 0.04 m/s with texture-based flow visualization applied. The recirculation zone where blood flows in the opposing direction becomes clear.