Image Space Based Visualization of Unsteady Flow on Surfaces

Robert S. Laramee[†]   Bruno Jobard[‡]   Helwig Hauser[†]

[†]VRVis Research Center, Austria, `www.VRVis.at`, {`Laramee,Hauser`}`@VRVis.at`
[‡]University of Pau, France, `www.univ-pau.fr`, `bjobard@univ-pau.fr`

Figure 1: Visualization of flow on the surface of an intake manifold. The ideal intake manifold distributes flow evenly to the piston valves.

Figure 2: Visualization of flow at the complex surface of a cooling jacket -a composite of over 250,000 polygons.

Abstract

We present a novel technique for direct visualization of unsteady flow on surfaces from computational fluid dynamics. The method generates dense representations of time-dependent vector fields with high spatio-temporal correlation using both Lagrangian-Eulerian Advection and Image Based Flow Visualization as its foundation. While the 3D vector fields are associated with arbitrary triangular surface meshes, the generation and advection of texture properties is confined to image space. Frame rates of up to 20 frames per second are realized by exploiting graphics card hardware. We apply this algorithm to unsteady flow on boundary surfaces of, large, complex meshes from computational fluid dynamics composed of more than 250,000 polygons, dynamic meshes with time-dependent geometry and topology, as well as medical data.

CR Categories: 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

Keywords: Unsteady flow visualization, computational fluid dynamics (CFD), surface representation, texture mapping

1 Introduction

Dense, texture-based, unsteady flow visualization on surfaces has remained an elusive problem since the introduction of texture-based flow visualization algorithms themselves. The class of fluid flow visualization techniques that generate dense representations based on textures started with the Spot Noise [van Wijk 1991] and LIC [Cabral and Leedom 1993]. The main advantage of this class of algorithms is their *complete* depiction of the flow field while their primary drawback is, in general, the computational time required to generate the results.

Recently, two new algorithms, namely Lagrangian-Eulerian Advection (LEA) [Jobard et al. 2001] and Image Based Flow Visualization (IBFV) [van Wijk 2002], have been introduced that overcome the computation time hurdle by generating two-dimensional flow visualization at interactive frame rates, even for unsteady flow. This paves the way for the introduction of new algorithms that overcome the same problems on boundary surfaces and in three dimensions. In this paper we present a new algorithm that generates dense representations of arbitrary fluid flow on complex, non-parameterized surfaces, more specifically, surfaces from computational fluid dynamics (CFD). However, the algorithm is general enough to apply to other vector field data associated with a surface such as blood vessel flow.

Traditional visualization of boundary flow using texture mapping first maps one or more 2D textures to a surface geometry defined in 3D space. The textured geometry is then rendered to image space.

Here, we alter the classic order of operations. First we project the surface geometry to image space and then apply texturing. In other words, conceptually texture properties are advected on boundary surfaces in 3D but in fact our algorithm realizes texture advection solely in image space. The result is a versatile visualization technique with the following characteristics:

- generates a dense representation of unsteady flow on surfaces

- visualizes flow on complex surfaces composed of polygons whose number is on the order of 200,000 or more

- visualizes flow on dynamic meshes with time-dependent geometry and topology

- visualizes flow independent of the surface mesh's complexity and resolution

- supports user-interaction such as rotation, translation, and zooming always maintaining a constant, high spatial resolution

- the technique is fast, realizing up to 20 frames per second

The performance is due, among other reasons, to the exploitation of graphics hardware features and utilization of frame-to-frame coherency. The rest of the paper is organized as follows: in Section 2 we discuss related work, Section 3 details unsteady flow visualization on surfaces from CFD. Implementation details are described in Section 4 while results and conclusions are discussed in Section 5.

2 Related Work

Our work focuses on texture-based representations of unsteady flow on complex, non-parameterized surfaces. The challenge of visualizing time-dependent vector fields on surfaces at fast frame rates remains unsolved in surveys of the research literature [Post et al. 2002; Stalling 1997]. However, several techniques have been proposed to successfully resolve parts of the problem. In the next two sections we describe the two main categories of approaches for dense representations on surfaces and dense representations of unsteady 2D vector fields.

2.1 Texture-Based Flow Visualization on Surfaces

Previous research with a focus on representations of the vector field on boundary surfaces is generally restricted to steady-state flow. This is mainly due to the prohibitive computational time required. An enhanced version of Spot Noise is applied to surfaces by de Leeuw and van Wijk [de Leeuw and van Wijk 1995]. Battke et al. [Battke et al. 1997] describe an extension of LIC for arbitrary surfaces in 3D. Some approaches are limited to curvilinear surfaces, i.e., surfaces that can be parameterized using 2D coordinates. Forssell and Cohen [Forssell and Cohen 1995] extend LIC to curvilinear surfaces with animation techniques and add magnitude and directional information. Mao et al. [Mao et al. 1997] present an algorithm for convolving solid white noise on triangle meshes in 3D space and extend LIC for visualizing a vector field on arbitrary surfaces in 3D. Stalling [Stalling 1997] provides a helpful overview of LIC techniques applied to surfaces. In particular, a useful comparison of parameterized vs. non-parameterized surfaces is given.

2.2 LEA and IBFV

The algorithm in this paper is a new approach that incorporates features of both LEA and IBFV. These very effective algorithms have recently been introduced to produce dense representations of unsteady, 2D vector fields.

Jobard et al. introduced a Lagrangian-Eulerian texture advection technique for 2D vector fields at interactive frame rates [Jobard


Figure 3: A wire frame view of the surface of two intake ports showing its 221,000 polygonal composition: (left) an overview from the top, note that many polygons are cover less than one pixel (right) a close-up view of the mesh between the two intake ports.

et al. 2001; Jobard et al. 2002]. The algorithm produces animations with high spatio-temporal correlation. Each still frame depicts the instantaneous structure of the flow, whereas an animated sequence of frames reveals the motion of a dense collection of particles when released into the flow. Particle paths are integrated as a function of time, referred to as the Lagrangian step, while the color distribution of the image pixels is updated in place (Eulerian step). The result represents a large step forward in bringing the visualization of unsteady flow to interactive frame rates.

Image Based Flow Visualization by van Wijk [van Wijk 2002] is the most recent algorithm for dense, 2D, unsteady vector field representations. It is based on the advection and decay of textures in image space. Each frame of the visualization is defined as a blend between the previous image, warped according to the flow direction, and a number of background images composed of filtered white noise textures. Performance times up to 50 frames per second are achieved through effective use of the graphics hardware.

3 Unsteady Flow Visualization on Surfaces

In this section we describe our technique in detail, starting with a discussion of those factors motivating the approach.

3.1 Physical Space vs. Parameter Space vs. Image Space

One approach to advecting texture properties on surfaces is via the use of a parameterization, a topic that has been studied *ad nauseam* (e.g., Levy et al. [Lévy et al. 2002]). According to Stalling [Stalling 1997], applying LIC to surfaces becomes particularly easy when the *whole* surface can be parameterized globally in two dimensions, e.g., in the manner of Forssell and Cohen [Forssell 1994; Forssell and Cohen 1995]. However, there are drawbacks to this approach. Texture distortions are introduced by the mapping between parameter space and physical space and, more importantly, for a large number of surfaces, no global parameterization is available such as isosurfaces from marching cubes and most unstructured surface meshes resulting from CFD. Surface meshes from CFD may consist of smoothly joined parametric patches, but can have a complex topology and therefore, in general, cannot be parameterized globally. Figures 2 and 3 are examples of surfaces for which a global parameterization is not easily derived.

Another approach to advecting texture properties on surfaces would be to immerse the mesh into a 3D texture, then the texture properties could be advected directly according to the 3D vector

132

field. This would have the advantages of simplifying the mapping between texture and physical space and would result in no distortion of the texture. However, this visualization would be limited to the maximum resolution of the 3D texture, thus causing problems with zooming. Also, this approach would not be very efficient in that most of the texels are not used. The amount of texture memory required would also exceed that available on our graphics card, e.g., we would need approximately 500MB of texture memory if we use 4 bytes per texel and a 512^3 resolution texture.

Can the problem be reduced to two dimensions? The surface patches can be packed into texture space via a triangle packing algorithm in the manner described by Stalling [Stalling 1997]. However, the packing problem becomes complex since our CFD meshes are composed of many scalene triangles as opposed to the equilateral and isosceles triangles often found in computational geometry. The problem of packing scalene triangles has been studied by Carr et al. [Carr and Hart 2002]. For CFD meshes, triangles generally have very disparate sizes. For a given texture resolution, many triangles would have to be packed that cover less than one texel. To by-pass this, the surfaces could be divided into several patches which could be stored into a texture atlas [Lévy et al. 2002]. In any case, computation time would be spent generating texels which cover polygons hidden from the current point of view. The preceding discussion leads us to an alternative solution that, ideally, has the following characteristics: works in image space, efficiently handles large numbers of surface polygons, spends no extra computation time on occluded polygons, does not spend computation time on polygons covering less than a pixel, and supports user interaction such as zooming, translation, and rotation.

3.2 Method Overview

The algorithm presented here simplifies the problem by confining the advection of texture properties to image space. We project the surface geometry to image space and then apply a series of textures. This order of operations eliminates portions of the surface hidden from the viewer. In short, our proposed method for visualization of flow on surfaces is comprised of the following procedure:

1. associate the 3D flow data with the polygons at the boundary surface i.e., a velocity vector is stored at each polygon vertex of the surface

2. project the surface and its vector field onto the image plane

3. identify geometric discontinuities

4. advect texture properties according to the vector field in image space

5. inject and blend noise

6. apply additional blending along the geometric discontinuities previously identified

7. overlay all optional visualization cues such as showing a semi-transparent representation of the surface with shading

These stages are depicted schematically in Figure 4. Each step of the pipeline is necessary for the dynamic cases of unsteady flow, time-dependent geometry, rotation, translation, and scaling, and only a subset is needed for the static cases involving steady-state flow and no changes to the view-point. We consider each of these stages in more detail in the sections that follow.

3.3 Vector Field Projection

In order to advect texture properties in image space, we must project the vector field associated with the surface to the image plane, taking into account that the velocity vectors are stored at the polygon vertices. We chose to take advantage of the graphics hardware to


Figure 4: Flow diagram of texture-based flow visualization on complex surfaces -*k* represents time as a frame number.

project the vector field to the image plane. We assign a color whose *R*, *G*, and *B* values encode the *x*, *y*, and *z* components of the local vectors to each vertex of the boundary surface respectively. The velocity-colored geometry is rendered to the framebuffer. We use the term *velocity image* to describe the result of encoding the velocity vectors at the mesh vertices into color values. The velocity image is interpreted as the vector field and is used for the texture advection in image space. More precisely, the color assignment can be done with a simple scaling operation. For each color component, $\mathbf{h}_{rgb}$, we assign a velocity, $\mathbf{v}_{xyz}$ component according to:

$$\mathbf{h}_r = \frac{\mathbf{v}_x - \mathbf{v}_{min\,x}}{\mathbf{v}_{max\,x} - \mathbf{v}_{min\,x}}$$
$$\mathbf{h}_g = \frac{\mathbf{v}_y - \mathbf{v}_{min\,y}}{\mathbf{v}_{max\,y} - \mathbf{v}_{min\,y}} \qquad (1)$$
$$\mathbf{h}_b = \frac{\mathbf{v}_z - \mathbf{v}_{min\,z}}{\mathbf{v}_{max\,z} - \mathbf{v}_{min\,z}}$$

The minimum velocity component is subtracted for each color component respectively, in an effort to minimize loss of accuracy.

The use of a velocity image yields the following benefits: (1) the advection computation and noise blending is simpler in image space, thus we inherit advantages from the LEA and IBFV, (2) the vector field and polygon mesh are decoupled, thereby freeing up expensive computation time dedicated to polygons smaller than a single pixel, (3) conceptually, this is performing hardware-accelerated occlusion culling, since all polygons hidden from the viewer, are immediately eliminated from any further processing, and (4) this operation is supported by the graphics hardware. Saving the velocity image to main memory is simple, fast, and easy. A sample velocity image is shown in Figure 5 (top, left).

The construction of the velocity image allows us to take advantage of hardware-accelerated flow field reconstruction. During the construction of the velocity image, we enable Gouraud Shading, also supported by the graphics hardware. Since each velocity component is stored as hue at each polygon vertex of the surface, the smooth interpolation of hue amounts to hardware-accelerated vector field reconstruction. This is important for a minimum of two reasons. First, the polygonal primitive we choose at image advection time is independent of the original mesh polygons (more in Section 3.4). In other words, the vertices of the mesh we use to dis-

133

Figure 5: The 5 component images, plus a 6^{th} composite image, used for the visualization of surface flow on a ring: (top, left) the velocity image, (top, right) the geometric edge boundaries, (middle, left) the advected and blended textures, (middle, right) a sample noise image, (bottom, left) an image overlay, (bottom, right) the result of the composited images with an optional velocity color map. The geometric edge boundaries are drawn in black for illustration.

tort the image are not the same vertices where the original velocity vectors are stored. Second, interpolation is essential for flow field reconstruction. When the surface is rendered with velocity encoded as hue, the vertices of some polygons are clipped during the projection process. However, we still need to access the vector field values inside those polygons, and not just at the vertices, hence the need for reconstruction. We also note that we are not necessarily limited to linear interpolation for reconstruction. Higher order interpolation schemes can be supported by graphics hardware [Hadwiger et al. 2001].

The velocity vectors are de-coded from the velocity image according to:

$$
\begin{aligned}
\mathbf{v}_x &= \mathbf{h}_r \cdot (\mathbf{v}_{maxx} - \mathbf{v}_{minx}) + \mathbf{v}_{minx} \\
\mathbf{v}_y &= \mathbf{h}_g \cdot (\mathbf{v}_{maxy} - \mathbf{v}_{miny}) + \mathbf{v}_{miny} \\
\mathbf{v}_z &= \mathbf{h}_b \cdot (\mathbf{v}_{maxz} - \mathbf{v}_{minz}) + \mathbf{v}_{minz}
\end{aligned}
\tag{2}
$$

The de-coded velocity vectors used to compute the advection mesh (Sec 3.4) are then projected onto the image plane.

The magnitude of the velocity vectors at those parts of the surface orthogonal to the image plane may be shortened as a result of perspective projection, i.e., if the dot product between the image plane normal and the 3D surface normal is zero or close to zero. This can reduce the visual clarity of the vector field's direction during animation. In our implementation, we added an option that allows the user to apply a bias to the velocity vectors that are shortened close to zero due to the projection. We can use this bias to reduce the scaling effect for curved surfaces. Conceptually it is like applying a reverse velocity clamp.

The projection of the vectors to the image plane is done after velocity image construction for 2 reasons: (1) not *all* of the vectors have to be projected (Sec. 3.4), thus saving computation time and (2) we use the original 3D vectors for the velocity mask (Sec. 4.2).

3.4 Advection Mesh Computation and Boundary Treatment

After the projection of the vector field we compute the mesh used to advect the textures similar to IBFV. We distort a regular, rectilinear mesh according to the velocity vectors stored at mesh grid intersections. The distorted mesh vertices can then be computed by advecting each mesh grid intersection according to the discretized Euler approximation of a pathline, **p**, (the same as a streamline for steady flow) expressed as:

$$
\mathbf{p}_{k+1} = \mathbf{p}_k + \mathbf{v}^p(\mathbf{p}_k; t)\, \Delta t
\tag{3}
$$

where $\mathbf{v}^p$ represents a magnitude and direction after projection to the image plane. The texture coordinates located at the regular, rectilinear mesh vertices are then mapped to the (forward) distorted mesh positions. The distorted mesh positions are stored for fast advection of texture properties for static scenes.

Special attention must be paid in order to handle flow at geometric boundaries of the surface. Figure 6 shows an overview of the original IBFV process. During the visualization, each frame is advected, rendered, and blended in with a background image. If we look carefully at the *distort* phase of the algorithm, we notice that there is nothing to stop the image from being advected outside of the physical boundary of the geometry. While this is not a problem when the geometry covers the entire screen, this can lead to artifacts for geometries from CFD, especially in the case of boundaries with a non-zero outbound flow, e.g., flow outlets.

To address this problem we borrow a notion from LEA that treats non-rectangular flow domains, namely, the use of backward coordinate integration (also proposed by Max and Becker [Max and Becker 1999]). Using backward integration, equation 3 becomes:

$$
\mathbf{p}_{k-1} = \mathbf{p}_k - \mathbf{v}^p(\mathbf{p}_{k-1}; t)\, \Delta t
\tag{4}
$$

In this case the texture coordinates located at the (backward) distorted mesh positions are mapped to the regular, rectilinear mesh vertices. Backward integration does not allow advection of image properties past the geometric boundaries.


image k   distort   render   blend

k = k + 1

Figure 6: An overview of the original image based flow visualization

134

3.5 Edge Detection and Blending

While we gain many advantages by decoupling the image advection process with the 3D surface geometry, artifacts can result, especially in the case of geometries with sharp edges. If we look carefully at the result of advecting texture properties in image space, we notice that in some cases a visual flow continuity is introduced where it may be undesirable. A sample case is shown in Figure 7.


Figure 7: When a 3D surface geometry (left) is projected, continuity is created in image space (right). If the flow aligned texture properties are advected across this edge, an artificial flow continuity may result.

A portion of the 3D geometry, shown colored, is much less visible after the projection onto the image plane. If the flow texture properties are advected across this edge in image space, also shown colored, an artificial continuity results. To handle this, we incorporate a geometric edge detection process into the algorithm. During the image integration computation, we compare spatially adjacent depth values during one integration and advection step. We compare the associated depth values, z_{k-1} and z_k in world space of $\mathbf{p}_{k-1}$ and $\mathbf{p}_k$ from equation 4, respectively. If

$$|z_{k-1} - z_k| > \varepsilon \cdot |\mathbf{p}_{k-1} - \mathbf{p}_k| \qquad (5)$$

where ε is a threshold value, then we identify an edge. All positions, $\mathbf{p}$, for which equation 5 is true, are classified as edge crossing start points. Special treatment must be given when advecting texture properties from these points. This process does not detect *all* geometric edges, only those edges across which flow texture properties should not be advected.

Figure 5 top, right shows one set of edges from the detection process. The geometric edges are identified and stored during the dynamic visualization case and additional blending is applied (depicted schematically in Figure 4). During the edge blending phase of the algorithm we introduce discontinuities in the texture aligned with the geometric discontinuities from the surface, i.e., gray values are blended in at the edges. This has the effect of adding a gray scale phase shift to the pixel values already blended. This could obviously be handled in different ways, e.g., choosing a random noise value to advect or inverting the noise value already present. Some results of the edge detection and blending phase are illustrated in Figure 8. In our data sets an ε of 1-2% of depth buffer is practical. However, the users may set their own value if fine tuning of the visualization is needed.

The same edge detection and blending benefits incoming boundary flow. Also an artifact of the IBFV algorithm, geometric boundaries with incoming flow may appear dimmer than the rest of the geometry. This is a result of the noise injection and blending process described in Section 3.6. In short, the background color shows through more in areas of incoming flow because not as much noise has been blended in these areas. Figure 9, top, shows a 2D slice through a 3D mesh from a CFD simulation with incoming boundary flow coming in through the narrow inlet from the right. Note that the edge of the inlet appears dim. Figure 9, bottom, shows the same slice with edge blending turned on. The boundary artifacts


Figure 8: A close-up example of geometric edge detection: on the left side, geometric edge detection is disabled, on the right side enabled.

of the noise injection and blending process are no longer a distraction. Edge detection and blending also plays in important role while an object is rotating. Without special treatment, contours in image space become blurred when different portions of a surface geometry overlap, such as when blood vessels in Figure 12 overlap during rotation.

3.6 Noise Blending

By reducing the image generation process back to two dimensions, the noise injection and blending phase falls in line with the original IBFV technique, namely, an image, F, is related to a previous image, G, by [van Wijk 2002]:

$$F(\mathbf{p};k) = \alpha \sum_{i=0}^{k-1} (1-\alpha)^i G(\mathbf{p}_{k-i};k-i) \qquad (6)$$

where $\mathbf{p}$ represents a pathline, α defines a blending coefficient, and k represents time as a frame number. Thus a point, $\mathbf{p}_k$, of an image F_k, is the result of a convolution of a series of previous images, $G(\mathbf{x};i)$, along the pathline through $\mathbf{p}_k$, with an decay filter defined by $\alpha(1-\alpha)^i$. The blended noise images have both spatial and temporal characteristics. In the spatial domain, a single noise image, $g(x)$, is described as a linearly interpolated sequence of n random values, G_i, in the range $[0, n-1]$, i.e.,

$$g(x) = \sum h_s(x - is) G_{i \bmod n} \qquad (7)$$

where the spacing, s, between noise samples is generally greater than or equal to the distance traversed by an image property in one advection step and h_s represents a triangular black and white pulse function. Here x represents a location in the flow domain. In practice, we give the user control of s, resulting in multi-frequency texture resolutions in the spacial domain. The background textures used for blending also vary in time. In the temporal domain, each point, G_i in the background texture, periodically increases and decays according to a profile, $w(t)$, e.g.,

$$G_{i;k} = w((k/M + \phi_i) \bmod 1) \qquad (8)$$

where ϕ_i represents a random phase, drawn from the interval $[0,1)$, M is the total number of background noise images used, and where $w(t)$ is defined for all time steps. We use a square wave profile, i.e., $w(t) = 1$ if $t < 1/2$ and 0 otherwise. In our application, the user has the option of varying M. Smaller values of M result in higher frequency noise in the temporal domain whereas higher values M result in a lower temporal frequency. Figure 5 (middle, left) shows a sample blended image and Figure 5 (middle, right) shows a sample noise image.

135

Figure 9: Here we see a 2D slice through a 3D geometry from a CFD simulation. (top) With no edge blending, the background color shows through boundary areas with incoming flow. (bottom). With edge blending, these artifacts are no longer a distraction.

3.7 Image Overlay Application

The rendering of the advected image and the noise blending may be followed by an optional image overlay. An overlay enhances the resulting texture-based representation of surface flow by applying color, shading, or any attribute mapped to color (Fig. 5, bottom, left). In implementation, we generate the image overlay following the construction of the velocity image. This overlay may render any CFD simulation attribute mapped to hue. The overlay is constructed once for each static scene and applied after the image advection, edge blending, and noise blending phases. Since the image advection exploits frame-to-frame coherency, the overlay must be applied after the advection in order to prevent the surface itself from being smeared. Also worthy of mention, is that the opacity value of the image overlay is a free parameter we provide to the user.

4 Implementation

In this section we consider some aspects of the algorithm not previously discussed which are important for implementation. Our implementation is based on the highly portable OpenGL 1.1 (www.opengl.org) library.

4.1 Texture Clipping

In our application, the resolution of the quadrilateral mesh used to warp the image can be specified by the user. The user may specify a coarse resolution mesh, e.g., 128×128, for faster performance or a fine resolution mesh, e.g., 512×512, for higher accuracy. However, if the resolution of the advection mesh is too coarse in image space, artifacts begin to appear. Figure 10, left, illustrates these artifacts zoomed in on the edge of a surface. In order to minimize the jagged edges created by coarse resolution texture quadrilaterals, we apply a texture clipping function. Subsets of textured quadrilateral that do not cover the surface are clipped from the visualization as

shown in Figure 10, right. This can be implemented simply with the image overlay by maximizing the opacity wherever the depth buffer value is maximized, i.e., wherever there is a great depth.


Figure 10: The result of, left, a coarse resolution advection mesh with artifacts and, right, the application of texture clipping. The resolution of the advection mesh shown on the left is 32×32 for illustration.

4.2 Velocity Mask

In order to dim high frequency noise in low velocity regions, the user also has the option of applying a velocity mask. We adopt the velocity mask of Jobard et al. [Jobard et al. 2001] for our purposes here, namely:

$$\alpha = 1 - (1 - \mathbf{v})^m \qquad (9)$$

where α decreases as a function of velocity magnitude. In our case, the image overlay becomes more opaque in regions of low velocity and more transparent in areas of high velocity. With the velocity mask enabled, the viewer's attention is drawn away from areas of stagnant flow, and towards areas of high flow velocity. We note that in the context of CFD simulation data, engineers are often very concerned about areas of stagnant flow. In the case of a cooling jacket, stagnant flow may represent a region of the geometry where the temperature is too high, possibly leading to boiling conditions thus reducing the effectiveness of the cooling jacket itself. Therefore, in our case the engineers may disable the velocity mask or use the velocity mask to *highlight* areas of flow, e.g., make the hue brighter in areas of low velocity.

5 Performance and Results

Our visualization technique is applied primarily to large, highly irregular, adaptive resolution meshes commonly resulting from computational fluid dynamics simulations. [1] The ideal intake manifold (Fig. 1) supplies an equal amount of fluid flow to each piston valve. Visualizing the flow at the surface gives the engineer insight into any imbalances between the inlet pipes, in this case, the 3 long narrow pipes of the geometry. Figure 13 shows our method applied to a surface of an intake port mesh (from Fig. 3) composed of 221K polygons. The intake port mesh is composed of highly adaptive resolution surface polygons and for which no global parameterization is readily available. The method described here allows the user to zoom in at arbitrary view points always maintaining a high spatial resolution visualization. The algorithm applies equally well to meshes with time-dependent geometry and topology. Figure 11 shows the surface of a piston cylinder with the piston head (not shown) defining the bottom of the surface. The method here enables the visualization of fuel intake as the piston head slides down the cylinder. The resulting flow visualization has a smooth spatio-temporal coherency. Our algorithm also has applications in the field

---

[1]Supplementary video available at
http://www.VRVis.at/ar3/pr2/vis03/

136

Figure 11: Snapshots from the visualization of a time-dependent surface mesh composed of a 79K polygons with dynamic geometry and topology. This intake valve and piston cylinder can also be used to analyse the formation of *wall film*, the term used to describe the liquid buildup on surfaces.

of medicine. Figure 12 shows the circulation of blood at the junction of 3 blood vessels. An abnormal cavity has developed that may hinder the optimal distribution of blood.


Figure 12: Visualization of blood flow at the surface of the junction of 3 blood vessels. Stagnant blood flow may occur within the abnormal pocket at the junction.

Performance was evaluated on an HP Visualize workstation with an HP *fx* graphics card, running Red Hat Linux 7.2 with a 1 GHz Pentium III dual processor and 1 GB of RAM. The performance times reported in Table 1 support interactive exploration of unsteady flow on surfaces. The first time reported in the FPS column is for the static cases of steady-state visualization and the absence of changes to the view point. The times shown in parenthesis indicate the dynamic cases of unsteady flow and interactive zooming and rotation. More specifically, the dynamic cases require the construction of a velocity image, image overlay, as well as geometric edge detection. We include geometric edge detection in the frame rates reported in Table 1. It does not introduce significant overhead since it is easily built into the advection process itself.

The performance time of our algorithm depends on the resolution of the mesh used to perform the advection and the number of polygons in the original surface mesh. In the case of steady-state flow, the algorithm no longer depends on the number of polygons in

the surface mesh, but on the area covered in image space. The data set shown in Figure 1, left, does not cover as much image space, so its performance times are somewhat higher in the static case.

data set	number of polygons	advection mesh resolution	frames per second
ring (Fig 5)	10K	128×128	18 (5)
		256×256	9 (3)
		512×512	3 (1)
intake manifold (Fig 1)	48K	128×128	22 (2)
		256×256	14 (2)
		512×512	6 (1)
combustion chamber (Fig 11)	79K	128×128	17 (2)
		256×256	10 (2)
		512×512	4 (1)
intake port (Fig 13)	221K	128×128	17 (0.5)
		256×256	7 (0.5)
		512×512	2 (0.3)

Table 1: Sample frame rates for the visualization algorithm.

6 Conclusions and Future Work

We have presented a novel technique for dense representations of unsteady flow on boundary surfaces from CFD. The algorithm supports visualization of flow on arbitrary surfaces at up to 20 FPS via the careful use of graphics hardware. It supports exploration and visualization of flow on large, unstructured polygonal meshes, and on time-dependent meshes with dynamic geometry and topology. The method generates dense representations of time-dependent vector fields building on both the LEA and IBFV algorithms. It also does not waste computation time on occluded polygons or polygons covering less than one pixel. While the vector fields are defined in 3D and associated with arbitrary triangular surface meshes, the generation and advection of texture properties is confined to image space.

Future work can go in many directions including visualization of unsteady 3D flow, something we expect to see soon. Challenges will include both interactive performance time and perceptual issues. Future work also includes the application of more specialized graphics hardware features like programmable per-pixel operations in the manner of Weiskopf et al. [Weiskopf et al. 2002; Weiskopf et al. 2001] and the use of pixel textures like Heidrich et al. [Heidrich et al. 1999].

IEEE COMPUTER SOCIETY

Figure 13: A view of the surface of an 221K polygonal intake port mesh show in Figure 3. Texture-based flow visualization is applied to the surface.


Figure 14: A close-up view of the surface of the intake port mesh show in Figure 13. Here we illustrate user-supported zooming with automatic, on the fly recalculation of the flow texture.

References

BATTKE, H., STALLING, D., AND HEGE, H. 1997. Fast Line Integral Convolution for Arbitrary Surfaces in 3D. In *Visualization and Mathematics*, Springer-Verlag, Heidelberg, 181–195.

CABRAL, B., AND LEEDOM, L. C. 1993. Imaging Vector Fields Using Line Integral Convolution. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, vol. 27, 263–272.

CARR, N. A., AND HART, J. C. 2002. Meshed Atlases for Real-Time Procedural Solid Texturing. *ACM Transactions on Graphics 21*, 2, 106–131.

DE LEEUW, W., AND VAN WIJK, J. 1995. Enhanced Spot Noise for Vector Field Visualization. In *IEEE Visualization '95 Proceedings*, IEEE Computer Society, 233–239.

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

FORSSELL, L. K. 1994. Visualizing Flow over Curvilinear Grid Surfaces Using Line Integral Convolution. In *Proceedings of the Conference on Visualization*, IEEE Computer Society Press, Los Alamitos, CA, 240–247.

HADWIGER, M., THEUSSL, T., HAUSER, H., AND GRÖLLER, E. 2001. Hardware-Accelerated High-Quality Reconstruction on PC Hardware. In *Proceedings of the Vision Modeling and Visualization Conference 2001 (VMV-01)*, 105–112.

HEIDRICH, W., WESTERMANN, R., SEIDEL, H.-P., AND ERTL, T. 1999. Applications of Pixel Textures in Visualization and Realistic Image Synthesis. In *ACM Symposium on Interactive 3D Graphics*, ACM/Siggraph.

JOBARD, B., ERLEBACHER, G., AND HUSSAINI, M. Y. 2001. Lagrangian-Eulerian Advection for Unsteady Flow Visualization. In *IEEE Visualization*, IEEE.

JOBARD, B., ERLEBACHER, G., AND HUSSAINI, Y. 2002. Lagrangian-Eulerian Advection of Noise and Dye Textures for Unsteady Flow Visualization. In *IEEE Transactions on Visualization and Computer Graphics*, vol. 8(3), 211–222.

LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *SIGGRAPH 2002 Conference Proceedings*, Annual Conference Series, 362–371.

MAO, X., KIKUKAWA, M., FUJITA, N., AND IMAMIYA, A. 1997. Line Integral Convolution for 3D Surfaces. In *Visualization in Scientific Computing '97. Proceedings of the Eurographics Workshop in Boulogne-sur-Mer, France*, Eurographics, 57–70.

MAX, N., AND BECKER, B. 1999. Flow Visualization Using Moving Textures. In *Data Visualization Techniques*, 99–105.

POST, F. H., VROLIJK, B., HAUSER, H., LARAMEE, R. S., AND DOLEISCH, H. 2002. Feature Extraction and Visualization of Flow Fields. In *Eurographics 2002 State-of-the-Art Reports*, The Eurographics Association, Saarbrücken Germany, 69–100.

STALLING, D. 1997. LIC on Surfaces. In *Texture Synthesis with Line Integral Convolution*, SIGGRAPH '97, Int. Conf. Computer Graphics and Interactive Techniques, 51–64.

VAN WIJK, J. 1991. Spot noise-texture synthesis for data visualization. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, vol. 25, 309–318.

VAN WIJK, J. J. 2002. Image Based Flow Visualization. In *SIGGRAPH 2002 Conference Proceedings*, Annual Conference Series, 745–754.

WEISKOPF, D., HOPF, M., AND ERTL, T. 2001. Hardware-Accelerated Visualization of Time-Varying 2D and 3D Vector Fields by Texture Advection via Programmable Per-Pixel Operations. In *Proceedings of the Vision Modeling and Visualization Conference 2001 (VMV-01)*, 439–446.

WEISKOPF, D., ERLEBACHER, G., HOPF, M., AND ERTL, T. 2002. Hardware-Accelerated Lagrangian-Eulerian Texture Advection for 2D Flow Visualizations. In *Proceedings of the Vision Modeling and Visualization Conference 2002 (VMV-01)*, 439–446.

IEEE
COMPUTER
SOCIETY