

Lagrangian-Eulerian Advection of Noise and Dye Textures for Unsteady Flow Visualization

Bruno Jobard, Gordon Erlebacher, and M. Yousuff Hussaini

Abstract—A new hybrid scheme (LEA) that combines the advantages of Eulerian and Lagrangian frameworks is applied to the visualization of dense representations of time-dependent vector fields. The algorithm encodes the particles into a texture that is then advected. By treating every particle equally, we can handle texture advection and dye advection within a single framework. High temporal and spatial correlation is achieved through the blending of successive frames. A combination of particle and dye advection enables the simultaneous visualization of streamlines, particle paths, and streaklines. We demonstrate various experimental techniques on several physical flow fields. The simplicity of both the resulting data structures and the implementation suggest that LEA could become a useful component of any scientific visualization toolkit concerned with the display of unsteady flows.

Index Terms—Flow visualization, texture advection, unsteady flow fields.

1 INTRODUCTION

TRADITIONALLY, unsteady flow fields are thought of as a collection of pathlines or streaklines that originate from user-defined seed points [9], [10]. These trajectories are often visualized in experimental laboratories through the injection of dye into the fluid [17]. Release of the dye can be continuous or pulsed, colored, locally concentrated, or spread out over an extended region. Proper use of dye makes it possible to extract quantitative information about the velocity magnitude, direction, rate of strain tensor, and to visualize pathlines, streaklines, and timelines [7], [11], [13], [18].

At the other end of the spectrum, the use of dense representations of the flow as a means of maximizing information content has been attracting a lot of attention [3], [6], [7], [8], [12], [13], [15], [18]. In a time-dependent context, the fundamental challenge faced by this class of algorithms is to produce smooth animations with good spatial and temporal correlation.

In this paper, we describe a new algorithm based on dense representations of time-dependent vector fields and apply it to noise and dye advection. The method combines the advantages of the Lagrangian and Eulerian formalisms. Briefly stated, a dense set of particles is stored as coordinates in a texture. Each iteration is defined by a Lagrangian step (backward time integration of the particles) and a Eulerian step (update of the image pixel colors). Texture advection and dye advection are differentiated chiefly by how the advected textures are defined. All information is stored in a few two-dimensional arrays. By its very nature, the algorithm takes advantage of spatial

locality and instruction pipelining and can generate animations at interactive frame rates.

The rest of the paper is organized as follows: Section 2 gives an overview of related work. The general approach is described in Section 3, while the algorithm is examined in Sections 4 (noise advection) and 5 (dye advection). Section 6 discusses parameter selection. Timing results are presented in Section 7. Conclusions are drawn in Section 8.

2 RELATED WORK

Several techniques have been advanced to produce dense representations of unsteady vector fields. Best known is, perhaps, UFLIC (Unsteady Flow LIC), developed by Shen and Kao [15] and based on the Line Integral Convolution (LIC) technique [2]. The algorithm achieves good spatial and temporal correlation. However, the images are difficult to interpret: The paths are blurred in regions of rapid change of direction and are thickest where the flow is almost uniform. The low performance of the algorithm is explained by the large number of particles (three to five times the number of pixels in the image) to process for each animation frame.

The spot noise technique, initially developed for the visualization of steady vector fields, has a natural extension to unsteady flows [3]. A sufficiently large collection of elliptic spots is chosen to entirely cover an image of the physical domain. The position of these spots is integrated along the flow, bent along the local pathline or streamline, and finally blended into the animation frame. The rendering speed of the algorithm can be increased by decreasing the number of spots in the image. The control of pixel coverage is done by assigning a fixed lifespan to each spot.

Max and Becker [12] propose a texture-based algorithm to represent steady and unsteady flow fields. The basic idea is to advect a texture along the flow either by advecting the vertices of a triangular mesh or by integrating the texture coordinates associated with each triangle backward in time. When texture coordinates or particles leave the physical

• B. Jobard is with the Swiss Center for Scientific Computing, Via Cantonale, 6928 Mano, Switzerland. E-mail: bjobard@cscs.ch.

• G. Erlebacher and M.Y. Hussaini are with the School of Computational Science and Information Technology, Dirac Science Library, Florida State University, Tallahassee, FL 32306-4120. E-mail: {erlebach, myh}@csit.fsu.edu.

Manuscript received 15 Feb. 2002; revised 15 Mar. 2002; accepted 2 Apr. 2002.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org, and reference IEEECS Log Number 116206.

domain, an external velocity field is linearly extrapolated from the boundary. This technique attains interactive frame rates by controlling the resolution of the underlying mesh.

A technique to display streaklines was developed by Rumpf and Becker [13]. They precompute a two-dimensional noise texture whose coordinates represent time and a boundary Lagrangian coordinate. Particles at any point in space and time that originate from an inflow boundary are mapped back to a point in this texture.

More recently, Jobard et al. [6], [7] extend the work of Heidrich et al. [4] to animate unsteady two-dimensional vector fields. The algorithm relies heavily on extensions to OpenGL proposed by SGI, in particular, pixel textures, additive and subtractive blending, and color transformation matrices. They pay particular attention to the flow entering and leaving the physical domain, leading to smooth animations of arbitrary duration. Excessive discretization errors associated with 12 bit textures are addressed by a tiling mechanism [6]. Unfortunately, the graphics hardware extension this algorithm relies on most, the pixel texture extension, was not adopted by other graphics card manufacturers. As a result, the algorithm only runs on the SGI Maximum Impact and the SGI Octane with the MXE graphics card. Application of LEA to flows with shocks is considered in Hussaini et al. [5]. Recently, a new algorithm based on the Nvidia GeForce3 graphics card has been developed for texture advection [18].

3 LAGRANGIAN-EULERIAN APPROACH

We wish to track a collection of particles, p_i , along a prescribed time-dependent velocity field, that densely covers a rectangular region. If we assign a property $P(p_i)$ to the i th particle p_i , the property remains constant as the particle follows its pathline. At any given instant t , each spatial location \mathbf{x} has an associated particle, labeled $p^t(\mathbf{x})$. One expresses that the particle property is invariant along a pathline by

$$\frac{\partial P(p^t(\mathbf{x}))}{\partial t} + \mathbf{v}^t(\mathbf{x}) \cdot \nabla P(p^t(\mathbf{x})) = 0. \quad (1)$$

The property attached to each particle takes on the role of a passive scalar. Its value is therefore not affected by diffusion or source terms (associated with chemical or other processes). This equation has two interpretations. In the first, the trajectory of a single particle, denoted by $\mathbf{x}^t(p)$ where p tags the particle, satisfies

$$\frac{d\mathbf{x}^t(p)}{dt} = \mathbf{v}^t(\mathbf{x}^t, p). \quad (2)$$

In this *Lagrangian* approach, the trajectory of each particle is computed separately. The time evolution of a collection of particles is displayed by rendering each particle by a glyph (point, texture spot [3], arrows). Except for the recent work of Jobard et al. [5], [6], [7], [18], current time-dependent algorithms are all based on particle tracking, e.g., [1], [3], [10], [12], [15]. While Lagrangian tracking is well-suited to the task of understanding how dense groups of particles evolve in time, it suffers from several shortcomings. In regions of flow convergence, particles may




Fig. 1. Rotation of bitmap image about the lower left corner. (a) Original image. (b) Image rotated with Lagrangian scheme. (c) Image rotated with Eulerian scheme.

accumulate into small clusters that follow almost identical trajectories, leaving regions of flow divergence with a low density of particles. To maintain dense coverage of the domain, the data structures must support dynamic insertion and deletion of particles [15] or track more particles than needed [3], which decreases the efficiency of any implementation.

Alternatively, a *Eulerian* approach solves (1) directly. Particles lose their identity. However, the particle property, viewed as a field, is known for all time at any spatial coordinate. Unfortunately, any explicit discretization of (1) is subject to a *Courant condition*¹ so that, in practice, the numerical integration step is limited to at most one to two cell widths. In turn, this imposes a maximum rate at which flow structures can evolve.

In our approach, we choose a hybrid solution. Between two successive time steps, coordinates of a dense collection of particles are updated with a Lagrangian scheme, whereas the advection of the particle property is achieved with a Eulerian method. At the beginning of each iteration, a new dense collection of particles is chosen and assigned the property computed at the end of the previous iteration. We refer to the hybrid nature of this approach as a Lagrangian-Eulerian Advection (LEA) method.

To illustrate the idea, consider the advection of the bitmap image shown in Fig. 1a by a circular vector field centered at the lower left corner of the image. With a pure Lagrangian scheme, a dense collection of particles (one per pixel) is first assigned the color of the corresponding underlying pixel. Each particle advects along the vector field and deposits its color property in the corresponding pixel in a new bitmap image. This technique does not ensure that every pixel of the new image is updated. Indeed, holes usually appear in the resulting image (Fig. 1b).

A better scheme considers each pixel of the new image as a particle whose position is integrated backward in time. The particle position in the initial bitmap determines its color. There are no longer any holes in the new image (Fig. 1c). Repeating the process at each iteration, any property can be advected while maintaining a dense coverage of the domain.

The core of the advection process is thus the composition of two basic operations: *coordinate integration* and *property advection*.

1. If the discrete time step exceeds some maximum value, severe numerical instabilities result.

Given the position $\mathbf{x}^0(i, j) = (i, j)$ of each particle in the new image, backward integration of (2) over a time interval h determines its position

$$\mathbf{x}^{-h}(i, j) = \mathbf{x}^0(i, j) + \int_0^{-h} \mathbf{v}^{t-\tau}(\mathbf{x}^\tau(i, j)) d\tau \quad (3)$$

at a previous time step. h is the integration step, $\mathbf{x}^\tau(i, j)$ represents intermediary positions along the pathline passing through $\mathbf{x}^t(i, j)$ and \mathbf{v}^τ is the vector field at time τ .

An image of resolution $W \times H$, defined at a previous time $t - h$, is advected to time t through the indirection operation

$$\mathbf{I}^t(i, j) = \begin{cases} \mathbf{I}^{t-h}(\mathbf{x}^{-h}(i, j)) & \forall \mathbf{x}^{-h} \in [0, W] \times [0, H] \\ \text{user-specified value} & \text{otherwise,} \end{cases} \quad (4)$$

which allows the image at time t to be computed from the image at any prior time $t - h$. This technique was used by Max and Becker [12]. However, instead of integrating back to the initial time to advect the same initial texture [12], we choose h to be the interval between two successive displayed images and always advect the last computed image. This minimizes the need to access coordinate values outside the physical domain. Notice that at least a linear interpolation of \mathbf{I}^{t-h} pixels at the positions \mathbf{x}^{-h} is necessary to obtain an image of acceptable quality.

In the next two sections, we describe noise-based and dye-based advection methods.

4 NOISE-BASED ADVECTION

With our Lagrangian-Eulerian approach, a full per-pixel advection requires manipulating exactly $W \times H$ particles. All information concerning any particle is stored in two-dimensional arrays with resolution $W \times H$ at the corresponding location (i, j) . Thus, we store the initial coordinates (x, y) of those particles in two arrays $\mathbf{C}_x(i, j)$ and $\mathbf{C}_y(i, j)$. Two arrays \mathbf{C}'_x and \mathbf{C}'_y contain their x and y coordinates after integration. A first order integration method requires two arrays \mathbf{V}_x and \mathbf{V}_y that store the velocity field at the current time. Similarly to LIC, we choose to advect noise images. Four noise arrays \mathbf{N} , \mathbf{N}' , \mathbf{N}_a , and \mathbf{N}_b contain, respectively, the noise to advect, two advected noise images, and the final blended image.

Fig. 2 shows a flowchart of the algorithm. After the initialization of the coordinate and noise arrays (Section 4.2), the coordinates are integrated (Section 4.3) and the initial noise array \mathbf{N} is advected (Section 4.4). The first advected noise array, \mathbf{N}' is then prepared for the next iteration by subjecting it to a series of treatments (left column in Fig. 2). Care is first taken to ensure that no spurious artifacts appear at boundaries where flow enters the domain (Section 4.5). This is followed by an optional masking process to allow for nonrectangular domains (Section 4.6). A low percentage of random noise is then injected into the flow to compensate for the effects of pixel duplication and flow divergence (Section 4.7). Finally, the coordinate arrays are reinitialized to ready them for the next iteration (Section 4.8). The right column in the flowchart describes the sequence of steps that transform the second advected

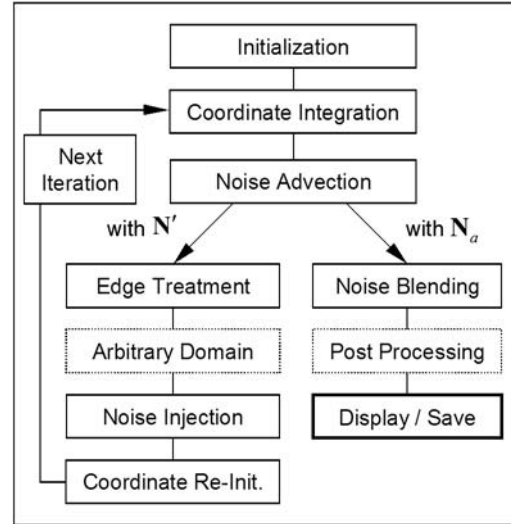


Fig. 2. Flowchart of LEA algorithm.

noise array \mathbf{N}_a into the final image. \mathbf{N}_a is first accumulated into \mathbf{N}_b via a blending operation to create the necessary spatio-temporal correlation (Section 4.9). Three optional postprocessing phases are then applied to \mathbf{N}_b before its final display: A line integral convolution filter removes aliasing effects (Section 4.10.1) and features of interest are emphasized via an opacity mask (Section 4.10.2).

4.1 Notation

Array cell values are referenced by the notation $\mathbf{A}(i, j)$ with i and j integers in $\{0, \dots, W - 1\} \times \{0, \dots, H - 1\}$. We adopt the convention that an array $\mathbf{A}(x, y)$ with real arguments is evaluated from information in the four neighboring cells using bilinear interpolation. A constant interpolation is explicitly noted $\mathbf{A}(\lfloor x \rfloor, \lfloor y \rfloor)$, where $\lfloor x \rfloor$ is the largest integer smaller than or equal to x . To simplify the notation, array operations such as $\mathbf{A} = \mathbf{B}$ apply to the entire domain of (i, j) .

The indirection operation $\mathbf{A}(i, j) = \mathbf{B}(r\mathbf{C}(i, j), s\mathbf{D}(i, j))$, where $\mathbf{C}(i, j)$ and $\mathbf{D}(i, j)$ lie in the range $[0, W - 1]$ and $[0, H - 1]$, respectively, and r and s are scalars, is denoted by $\mathbf{A} = \mathbf{B}(r\mathbf{C}, s\mathbf{D})$.

4.2 Coordinate and Noise Initialization

We first initialize the coordinate arrays $\mathbf{C}_x, \mathbf{C}_y$ and the noise arrays \mathbf{N} and \mathbf{N}_b . Coordinates are initialized by

$$\begin{cases} \mathbf{C}_x(i, j) = i + \text{rand}(1) \\ \mathbf{C}_y(i, j) = j + \text{rand}(1), \end{cases} \quad (5)$$

where $\text{rand}(1)$ is a real number in $[0, 1)$. The random offset distributes coordinates on a jitter grid to avoid regular patterns that might otherwise appear during the first several steps of the advection. Note that the integer part of each coordinate serves as an index into the cell.

\mathbf{N} is initialized with a two-valued noise function (0 or 1) to ensure maximum contrast and its values are copied into \mathbf{N}_b . Coordinates and noise values are stored in floating-point format to ensure sufficient accuracy in the calculations.

4.3 Coordinate Integration

A first order discretization of (3) with a constant time step h gives

$$\begin{cases} \mathbf{C}'_x = \mathbf{C}_x - (l_{\max}/V_{\max})\mathbf{V}_x(r_{W_V}\mathbf{C}_x, r_{H_V}\mathbf{C}_y) \\ \mathbf{C}'_y = \mathbf{C}_y - (l_{\max}/V_{\max})\mathbf{V}_y(r_{H_W}\mathbf{C}_x, r_{H_V}\mathbf{C}_y), \end{cases} \quad (6)$$

where $r_{W_V} = (W_V - 1)/W$ and $r_{H_V} = (H_V - 1)/H$ for a vector field resolution of W_V by H_V . The two scaling factors r_{W_V} and r_{H_V} ensure that the coordinates of the velocity arrays stay within proper bounds. We replaced the constant integration time step h with the quotient l_{\max}/V_{\max} , where V_{\max} is the maximum velocity magnitude over the entire space-time domain. Therefore, the new quantity l_{\max} represents the maximal possible displacement of a particle over all iterations, measured in units of cell widths. The actual displacement of a particle is proportional to the local velocity, also measured in cell widths. The velocity arrays \mathbf{V}_x and \mathbf{V}_y at the current time are linearly interpolated between the two closest available vector fields in the data set. Section 6 discusses the relationship between parameters necessary to generate animations consistent with the physics of the problem.

A useful property of a first order formulation is that the velocity arrays are never accessed at a point outside the physical domain. We have also implemented a second order discretization, but found no noticeable improvement due to the small extent of the spatio-temporal correlations in the final display. However, high accuracy is important for dye advection (Section 5).

4.4 Noise Advection

The advection of noise described by (4) is applied twice to \mathbf{N} to produce two noise arrays \mathbf{N}' and \mathbf{N}_a , one for advection, one for display. \mathbf{N}' is an internal noise array whose purpose is to carry on the advection process and to reinitialize \mathbf{N} for the next iteration. To maintain a sufficiently high contrast in the advected noise, \mathbf{N}' is computed with a constant interpolation. Before \mathbf{N}' can be used in the next iteration, it must undergo a series of corrections to account for edge effects, the presence of arbitrary domains, and the deleterious consequences of flow divergence.

\mathbf{N}_a serves to create the current animation frame and no longer participates in the noise advection. It is computed using linear interpolation of \mathbf{N} to reduce spatial aliasing effects. \mathbf{N}_a is then blended into \mathbf{N}_b (Section 4.9).

A straightforward implementation of (4) leads to conditional expressions to handle the cases when

$$\mathbf{x}'(i, j) = (\mathbf{C}'_x(i, j), \mathbf{C}'_y(i, j))$$

is exterior to the physical domain. A more efficient implementation eliminates the need to test for boundary conditions by surrounding \mathbf{N} and \mathbf{N}' with a buffer zone of constant width. From (6), $\mathbf{x}'(i, j)$ refers to cells located at most a distance of $b = \lceil h \rceil$ cell widths away from the array borders. An expanded noise array of size $(W + 2b) \times (H + 2b)$ is therefore sufficient to prevent out of bound array accesses (see Fig. 3). The advected arrays \mathbf{N}' and \mathbf{N}_a are computed according to




Fig. 3. Noise arrays \mathbf{N}' and \mathbf{N}_a are expanded with a surrounding region $b = \lceil h \rceil$ cells wide.

$$\begin{cases} \mathbf{N}'(i + b, j + b) = \mathbf{N}(\lfloor \mathbf{C}'_x(i, j) \rfloor + b, \lfloor \mathbf{C}'_y(i, j) \rfloor + b) \\ \mathbf{N}_a(i, j) = \mathbf{N}(r_W \mathbf{C}'_x(i, j) + b, r_H \mathbf{C}'_y(i, j) + b), \end{cases} \quad (7)$$

for all $(i, j) \in \{0, \dots, W - 1\} \times \{0, \dots, H - 1\}$, where $r_W = (W - 1)/W$ and $r_H = (H - 1)/H$. The two scaling factors r_W and r_H ensure a properly constructed linear interpolation.

4.5 Edge Treatment

A recurring issue with texture advection that must be addressed is the proper treatment of information flowing *into* the physical domain. Within the context of this paper, we must determine the user-specified value in (4). To address this, we recall that the advected image contains a two-valued random noise with little or no spatial correlation. We take advantage of this property to replace the user-specified value by a random value (0 or 1). At each iteration, we simply store new random noise in the buffer zone, at negligible cost.

At the next iteration, \mathbf{N} will contain these values and some of them will be advected to the interior of the physical domain by (7). Since random noise has no spatial correlation, the advection of the surrounding buffer values into the interior region of \mathbf{N}' produces no visible artifacts.

To treat periodic flows in one or more directions, the noise values are copied from an inner strip of width b along the interior edge of \mathbf{N}' onto the buffer zone at the opposite boundary. As a result, particles leaving one side of the domain seamlessly reappear at its opposite side.

4.6 Incoming Flow in Arbitrary Shaped Domains

It often happens that the physical domain is nonrectangular or contains interior regions where the flow is not defined. Denote by B the boundaries interior to \mathbf{N} that delineates these regions. LEA handles this case with no modification by simply setting the velocity to zero where it is not defined. The stationary noise in these regions is hidden from the animation frame by superimposing a semitransparent map that is opaque where the flow is undefined. For example, the opaque regions of this map might represent shorelines or islands.

When the flow velocity normal to B is nonzero and points into the physical domain, the advection of stationary noise values will create noticeable artifacts in the form of streaks. This might happen if an underground flow, not visible in the display, emerges into view at B . If necessary, we suppress these streaks with the help of a precomputed Boolean mask (or, alternatively, a Boolean function) $\mathbf{M}(i, j)$ that determines whether or not the velocity field is defined




Fig. 4. Noise duplication. A single noise value is duplicated into four cells in a uniform 45 deg flow.

at (i, j) . $N'(i, j)$ is then updated with random noise where $M(i, j)$ is false.

4.7 Noise Injection

We propose a simple procedure to counteract a *duplication effect* that occurs during the computation of N' in (7). Effectively, if particles in neighboring cells of N' retrieve their property value from within the same cell of N , this value will be duplicated in the corresponding cells of N . Single property values in N may be duplicated onto neighboring cells in N' , where coordinates (C'_x, C'_y) have identical integer values.

To illustrate the source of noise duplication, we consider an example. Fig. 4 shows the evolution of property values and particle positions for four neighboring pixels during one integration step and one advection step. The vector field is uniform, is oriented at 45 degrees to the x axis, and points toward the upper right corner. At the start of the iteration, each particle has a random position within its pixel (Fig. 4a). To determine the new property value of each pixel, the particle positions are integrated backward in time (Fig. 4b). The property value of the lower left corner pixel is duplicated onto the four pixels (worst-case scenario) (Fig. 4c). The fractional position of each particle is then reinitialized for the next iteration.

Over time, the average size of contiguous constant color regions in the noise increases. This effect is undesirable since lower noise frequency reduces the spatial resolution of the features that can be represented. This duplication effect is further reinforced in regions where the flow has a strong positive divergence.

To break the formation of uniform blocks and to maintain a high frequency random noise, we inject a user-specified percentage of noise into N' . Random cells are chosen in N' and their value is inverted (a zero value becomes one and vice versa). The number of cells randomly inverted must be sufficiently high to eliminate the appearance of pixel duplication, but low enough to maintain the temporal correlation introduced by the advection step.

To quantify the effect of noise injection, we compute the energy content of the advected noise in N at different scales as a function of time. Although the Fourier transform would appear to be the natural tool for this analysis, the two-valued nature of the noise image suggests instead the use of the Haar wavelet (linear combination of Heaviside functions). We perform a two-dimensional Haar wavelet transform and compute the level of energy in different scale bands (the spatial scale of consecutive bands varies by a

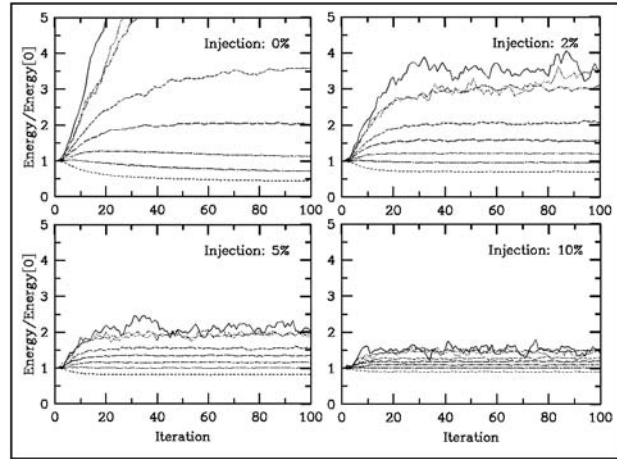


Fig. 5. Energy content of the flow at different scales based on a 2D Haar wavelet decomposition of the two-valued noise function (assumed to be isotropic). The energy in each band is scaled with respect to its initial value. Results are shown for injection rates of 0 percent, 2 percent, 5 percent, and 10 percent.

factor of two). The two-dimensional energy spectrum is reduced to a one-dimensional spectrum by assuming that the noise texture is isotropic at any point in time. (The smooth anisotropic flow results from *blending* multiple noise textures.) The energy in each band is scaled by its value after the initial noise injection. Ideally, we would like to preserve the initial spectrum at all times.

Fig. 5 illustrates the influence of the noise injection on the time evolution of the energy spectrum. Without injection, the energy in the larger scales (regions of pixel duplication) increases rapidly without stabilizing. This comes at the expense of some energy loss in the smaller scales (which decreases in the figure). As the percentage of noise injection increases, the spread of the scaled spectrum decreases continuously toward zero (the ideal state). However, excessive injection deteriorates the quality of the temporal correlation.

The necessary percentage of injected noise is clearly a function of the particular flow and depends on both space and time. It should be modeled as the contribution of two terms: a constant term that accounts for the duplication effects at zero divergence and a term that is a function of the velocity divergence. In the interest of simplicity and efficiency, we use a fixed percentage of two to three percent, which provides adequate results over a wide range of flows.

4.8 Coordinate Reinitialization

The coordinate arrays are reinitialized to prepare a new collection of particles to be integrated backward in time for the next iteration. However, coordinates are not reinitialized to their initial values. The advection equations presented in Section 3 assume that the particle property is computed at the previous time step via a linear interpolation. Unfortunately, the lack of spatial correlation in the noise image would lead to a rapid loss of contrast, which justifies our use of a constant interpolation scheme. However, the choice of constant interpolation implies that a property value can only change if it originates from a

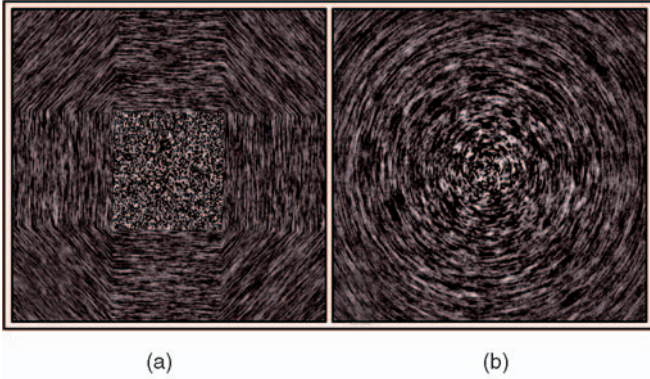


Fig. 6. Circular flow without and with accumulation of fractional displacement ($h = 2$).

different cell. If the coordinate arrays were reinitialized to their original values at each iteration, subcell displacements would be ignored and the flow would be frozen where the velocity magnitude is too low. This is illustrated in Fig. 6, which shows the advection of a steady circular vector field. Constant interpolation without fractional coordinate tracking clearly shows that the flow is partitioned into distinct regions within which the integer displacement vector is constant (Fig. 6a).

To prevent this from happening, we track the fractional part of the displacement within each cell. Instead of reinitializing the coordinates to their initial values, the fractional part of the displacement is added to cell indices (i, j) :

$$\begin{cases} C_x(i, j) = i + C'_x(i, j) - [C'_x(i, j)] \\ C_y(i, j) = j + C'_y(i, j) - [C'_y(i, j)]. \end{cases} \quad (8)$$

The effect of this correction is shown in Fig. 6b.

The coordinate arrays have now returned to the state they were in after their initialization phase ((5)); they verify the relations $[C_x(i, j)] = i$ and $[C_y(i, j)] = j$.

4.9 Noise Blending

Although successive advected noise arrays are correlated in time, each individual frame remains devoid of spatial correlation. By applying a temporal filter to successive frames, spatial correlation is introduced. We store the result of the filtering process in an array N_b . We have found the exponential filter to be convenient since its discrete version only requires the current advected noise and the previous filtered frame. It is implemented as an alpha blending operation

$$N_b = (1 - \alpha)N_b + \alpha N_a, \quad (9)$$

where α represents the opacity of the current advected noise array. A typical range for α is $[0.05, 0.2]$. Fig. 7 shows the effect of α on images based on the same set of noise arrays.

The blending stage is crucial because it introduces spatial correlation along *pathline* segments in every frame. To show clearly that the spatial correlation occurs along pathlines passing through each cell, we conceptualize the algorithm in 3D space; the x and y axes represent the spatial coordinates, whereas the third axis is time. To understand




Fig. 7. Frames obtained with different values of α .

the effect of the blending operation, let's consider an array N with black cells and change a single cell to white. During advection, a sequence of noise arrays (stacked along the time axis) is generated in which the white cell is displaced along the flow. By construction, the curve followed by the white cell is a pathline. The temporal filter blends successive noise arrays N_a with the most recent data weighted more strongly. The temporal blend of these noise arrays produces the projection of the pathline onto the $x - y$ plane, with an exponentially decreasing intensity as one travels back in time along the pathline. When the noise array with a single white cell is replaced by a two-color noise distribution, the blending operation introduces spatial correlation along a dense set of short pathlines.

Streamlines and pathlines passing through the same cell at the same time are tangent to each other, so a streamline of short extent is well-approximated by a short pathline. Therefore, the collection of short pathlines serves to approximate the instantaneous direction of the flow. With our LEA technique, a single frame represents the instantaneous structure of the flow (streamlines), whereas an animated sequence of frames reveals the motion of a dense collection of particles released into the flow.

The filtering phase completes one pass of the advection algorithm. The image N_b can be displayed to the screen or stored as an animation frame. N' is used as the initial noise texture N for the next iteration. It is worthwhile to mention that each iteration ends with data having the exact same property as when it started. In particular, the coordinate arrays satisfy

$$[C_x(i, j)] = i, \quad [C_y(i, j)] = j$$

and N contains a two-color noise without degradation of contrast.

4.10 Postprocessing

A series of optional postprocessing steps is applied to N_b to enhance the image quality and to remove features of the flow that are uninteresting to the user. We present two




Fig. 8. Frame without (bottom) and with (top) LIC filter. A velocity mask is applied to both images.

filters. A fast version of LIC can be applied to remove high frequency content in the image, while a *velocity mask* serves to draw attention to regions of the flow with strong currents.

4.10.1 Directional Low-Pass Filtering (LIC)

Although the temporal filter (noise blending phase) converts high frequency noise images into smooth spatially correlated images, aliasing artifacts remain visible in regions where the noise is advected over several cells in a single iteration. As a rule, aliasing artifacts become noticeable where noise advect more than one or two cells in a single time step (see Fig. 8, bottom). Experimentation with different low-pass filters led us to conclude that a Line Integral Convolution filter applied to N_b is the best filter to remove the effect of artifacts while preserving and enhancing the directional correlation resulting from the blending phase. This follows from the fact that temporal blending and LIC bring out the structure of pathlines and streamlines, respectively, and these curves are tangent to one another at each point. Although the image quality is often enhanced with longer kernel lengths, it is detrimental here since the resulting streamlines will have significant deviations from the actual pathlines. The partial destruction of the temporal correlation between frames would then lead to flashing effects in the animation. A secondary effect of longer kernels is decreased contrast.

While any LIC implementation can be used, our algorithm can advect an entire texture at interactive rates. Therefore, we are interested in the fastest possible LIC implementation. To the best of our knowledge, FastLIC [16] and Hardware-Accelerated LIC [4] are the fastest algorithms to date and both are well-suited to the task. However, we propose a simple, but very efficient, software version of Heidrich's hardware implementation to post-process the data when the highest quality is desired.

Besides the input noise array N_b , the algorithm requires two additional coordinate arrays, C_{xx} and C_{yy} , and an array N_{LIC} to store the result of the line integral convolution. The

```

float* ALIC(const float* Vx, const float* Vy,
            int Wv, int Hv,
            const float* Nb, int W, int H
            int L, float* N_LIC)
rWv = Wv/W ; rHv = Hv/H
rW = W / (W-1); rH = H / (H-1)
L2 = L div 2
r = 1 / (2*L2+1)
Loop over pixels i,j { N_LIC(i,j) = r*Nb(i,j) }
sgn = 1/Vmax
for n = 1 to 2 // Forward and backward advection
  Loop over pixels i,j
    Cxx(i,j)=i; Cyy(i,j)=j
    for k = 0 to L2
      Loop over pixels i,j //Coordinate integration
        Cxx(i,j)= ( Cxx(i,j)
          + sgn*Vx(rWv*Cxx(i,j), rHv*Cyy(i,j))+W)%W
        Cyy(i,j)= ( Cyy(i,j)
          + sgn*Vy(rWv*Cxx(i,j), rHv*Cyy(i,j))+H)%H

      Loop over pixels i,j // Noise advection
        // and Accumulation
        N_LIC(i,j) += r*Nb(rW*Cxx(i,j), rH*Cyy(i,j))
    end for
  sgn = -1/Vmax
end for
return N_LIC

```

Fig. 9. Pseudocode for ALIC (Array LIC).

length of the convolution kernel is denoted by L . For reference, we include the pseudocode for *Array-LIC* in Fig. 9.

In general, $L \approx h$ produces a smooth image with no aliasing. However, large values of h speed up the flow, with a resulting increase in aliasing effects. If the quality of the animation is important, L must be increased, with a resulting slowdown in the frame rate. The execution time of the LIC filter is commensurate with the timings of FastLIC for $L < 10$. Beyond 10, a serial FastLIC [16] should be used instead. An OpenMP implementation of our ALIC algorithm on shared memory architectures is straightforward. Results are presented in Section 7. As shown in Table 1, smoothing the velocity field with LIC reduces the frame rate by a factor of three across architectures. We recommend exploring the data at higher resolution without the filter or at low resolution with the filter.

4.10.2 Velocity Mask

A straightforward implementation of the texture advection algorithm described so far produces static images that show the flow streamlines and interactive animations that show the motion for the flow along pathlines. The length of the streaks is statistically proportional to the flow velocity magnitude. Additional information can be encoded into the images by modulating the color intensity according to one or more secondary variables.

It is often advantageous to superimpose the representation of flow advection over a background image that provides additional context. Two examples are shown in Fig. 10. In order to implement this capability, the image must become partially transparent.

Two approaches have been implemented. First, we couple the opacity of a pixel to its color intensity. Second, we modulate the pixel transparency with the magnitude of the velocity.

TABLE 1
Timings in Frames/Second as a Function
of Options and Resolutions

Options Resolutions	Advection		Advection + Velocity Mask ($m = n = 3$)		Advection + Velocity Mask + ALIC filter ($L = 6$)	
300 × 300	9.7	14.0	8.7	8.8	2.6	3.0
	16.3	39.0	10.4	27.0	3.6	11.6
500 × 500	3.5	4.7	3.2	3.1	0.9	1.0
	6.3	18.0	3.7	10.5	1.3	4.5
1000 × 1000	NA	1.2	NA	0.7	NA	0.2
	1.4	4.1	0.9	2.7	0.3	1.1

Each configuration has been tested on four different configurations: PC workstation (Xeon Dell Precision Workstation 530, 1.7 GHz, 1GB, 256KB L2 cache) (upper left), Octane (SGI Octane, R12000, 300MHz, 2GB, 2MB L2 cache, 64KB L1 cache) (upper right), Onyx2 (SGI Onyx2, R12000, 300MHz, 2GB, 8MB L2 cache, 64KB L1 cache) (lower left), and Onyx2 with four processors (lower right).

The blended image pixel color ranges from black to white. Neither color has a predominant role in representing the velocity streaks. Therefore, one of these colors can be eliminated and made partially transparent. We consider a black pixel to be transparent and a white pixel to be fully opaque. The transfer function that links these two states is a power law. Regions of the flow that are nearly stationary add little useful information to the view. For example, regions of high velocity are often of most interest in wind and ocean current data. Accordingly, we also modulate the transparency of each pixel according to the velocity magnitude. This produces a strong correlation between the length of the velocity streaks and their opacity.

The ideas described in the two previous paragraphs are implemented through an opacity map, also referred to as a velocity mask. Once computed, the velocity mask is combined with N_b into an intensity-alpha texture that is blended with the background image. We define the opacity map

$$A = (1 - (1 - V)^m)(1 - (1 - N_b)^n) \quad (10)$$

as a product of a function of local velocity magnitude and a function of the noise intensity. Higher values of the exponents m and n increase the contrast between regions of low and high velocity magnitude and low and high intensity, respectively. When $m = n = 1$, the opacity map reduces to

$$A = VN_b.$$

As the exponents are increased, the regions of high velocity magnitude and of high noise intensity increase their importance relative to other regions in the flow.

Higher quality pictures that emphasize the velocity magnitude can also be obtained by replacing the noise texture with a scalar map of the velocity magnitude (with color ranging from black to white as the magnitude ranges from zero to one) with an opacity defined by (8). As a result, the texture advection is seen *through* the opacity map.




Fig. 10. Noise-based advection with velocity mask and LIC filtering. Top: Sea currents in the Gulf of Mexico (data courtesy of COAPS, Florida State University). Bottom: Cyclone formation over Europe (data courtesy of MeteoSwiss, Switzerland).

5 DYE-BASED ADVECTION

By replacing the advected noise texture by a texture with a smooth pattern, one can emulate the process of dye advection, a standard technique used in experimental flow dynamics [17]. In this approach, a dye is released into the flow and advected with the fluid. Dye advection has been considered outside the context of texture advection techniques [11], [14]. Standard approaches include tracking clouds of particles or defining the dye with a polygonal enclosure. More recently, dye has been incorporated within the framework of texture advection schemes simply by including the dye into the advected texture [7], [18]. If the noise texture is replaced by a texture of uniform background color upon which the dye is deposited, the high frequency nature of the texture is removed. As a result, many of the treatments detailed in Section 4 are no longer required and the implementation is greatly simplified. The principal assumption in dye advection is that the color of the texture varies smoothly in space.

Effects that can safely be neglected include particle injection into the edge buffer, random particle injection,




Fig. 11. Dye advection in a circular flow. (a) First order time integration. (b) Second order time integration.

constant interpolation of the noise texture, coordinate reinitialization, and noise blending. We introduce here two new arrays \mathbf{D} and \mathbf{D}' that contain the dye at times t and $t + h$.

5.1 Dye Injection

Starting from a background texture of uniform color, we inject, at every iteration, dye into the fluid. The injection is often localized in space and assumes the shape of spots or curves distributed across the fluid. The dye can be released once and tracked (which approximates the path of a particle), released continuously at a single point (generates a streakline), or pulsed (the dye is turned on intermittently). To better discern the resulting patterns, colored dye is very useful. In this case, three-color components must be independently advected.

5.2 Dye Advection

The algorithm is similar to that used to advect the noise texture. However, when advecting dye, it is important to accurately track its long-time evolution. High order time integration schemes are therefore desirable. We have found that a second order Runge-Kutta midpoint algorithm provides sufficient accuracy. The scheme introduces an intermediate time step and an associated intermediate position \mathbf{x}^* . Although always within the extended physical, the velocity field at \mathbf{x}^* might not be defined. To avoid this situation, we set $\mathbf{x}' = \mathbf{x}^*$ whenever \mathbf{x}^* is outside the physical domain.

A second order algorithm offers no particular advantages for noise-based texture advection schemes since the accumulated error over the length of a typical streak is imperceptible. Noise-based representations cannot give clear information about the origin of a particle and its destination. When dye is released into the flow at a point, its subsequent trajectory lies along a particle path and the accumulation of error due to an inaccurate time integration scheme is very noticeable. Fig. 11 compares the result of first and second order integration schemes on the trajectory of dye continuously released into a circular flow. Unlike the first order scheme, the second order scheme produces closed circles.

In noise advection schemes, constant interpolation of the coordinate position was necessary to maintain maximum contrast in the advected texture, which led to the notion of subcell particle tracking (Section 4.8). Because of the smoothness of the dye images, we can directly apply the

advection scheme defined by (4). New property values of \mathbf{D}' are linearly interpolated at position \mathbf{x}' in \mathbf{D} . With linear interpolation, subcell coordinates are no longer required; at each iteration, the integration proceeds from the lower left corner of each cell. Coordinate arrays are not necessary for dye advection.

The new property of \mathbf{D}' cells is then computed from \mathbf{D} according to

$$\mathbf{D}'(i + b, j + b) = \mathbf{D}(r_W \mathbf{x}'_x(i, j) + b, r_H \mathbf{x}'_y(i, j) + b)$$

for all $(i, j) \in \{0, \dots, W - 1\} \times \{0, \dots, H - 1\}$. r_H and r_W coefficients are defined in Section 4.4. $\mathbf{x}'_k(i, j)$ is the k component of the backward integration from coordinates (i, j) .

The dye advection proceeds independently from the noise advection and the two resulting textures \mathbf{N}_b and \mathbf{D}' can be combined to form composite images.

During the advection phase, new property values in \mathbf{D}' are linearly interpolated from \mathbf{D} . Since the dye color now includes contributions from neighboring pixels with the background color, the dye is numerically diffused, producing a smoke-like effect. Although sometimes useful, there are times when we wish the dye to retain its sharp interface with the background flow. One solution to counter this problem is presented in the next section.

5.3 Diffusion Correction

The operation of linear interpolation acts as a smooth filter. A cross section of a white spot over a black background would show an intensity profile that gets progressively smoother as the dot is advected. To maintain a sharp profile, each color component could be corrected according to the simple prescription

$$\text{if } c < 0.5 \text{ then } c = 0 \text{ else } c = 1$$

for each color component $c \in [0, 1]$. However, this simple function produces excessive aliasing effects. Instead, we designed a filter function to steepen the smoothed profile toward the ideal discontinuous profile that would result in the absence of diffusion errors. Imposed constraints are that $c = 0, 0.5, 1$ be invariant to the filter. The rational function

$$c' = \frac{c - 1/2}{s(2|c - 1/2| - 1) + 1} + 1/2,$$

with the sharpness parameter $s \in [0, 1]$, satisfies all the requirements. The identity function is recovered when $s = 0$, while the simple thresholding function corresponds to $s = 1$.

Fig. 12 shows the effect of diffusion on the advection of a 16 pixel-wide square spot of dye along a circular vector field rotating clockwise and centered at $\mathbf{x} = 0$. Each time step ($h = 25$), the dye is released at four evenly spaced points along $x = 0$. The sharpness parameter is $s = 0, 0.2, 0.5, 1$. The sharpness is clearly improved when the filter is turned on. However, undesirable artifacts are clearly visible when $s = 1$. We have found that $s = 0.2$ is the minimum acceptable value.

A downside to the correction filter is that some dye is removed that is not within the diffused region. The effect is more noticeable if there are an insufficient number of pixels in the regions with dye. In Fig. 13, dye is released into the

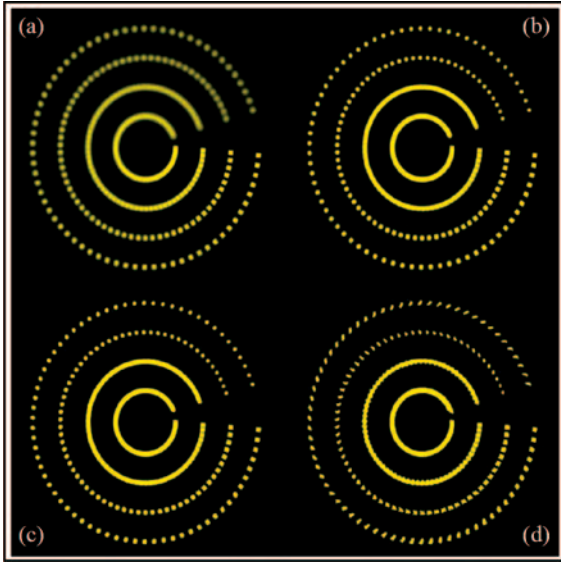


Fig. 12. Effect of sharpness parameter. (a) $s = 0$, (b) $s = 0.2$, (c) $s = 0.5$, (d) $s = 1$.

flow with a thickness of 16 (top) and eight (bottom) pixels. No filter is applied in the left column, while $s = 0.2$ in the right column. Although sharper with the filter on, the thinner dye has disappeared from the image before completing one revolution (bottom-right figure).

6 PARAMETER SETTINGS

In Section 4.3, we found it convenient to control the speed of particles along a pathline with the parameter l_{\max} which represents the maximum distance a property can be advected, measured in cell widths. In a stationary flow, the velocity field is fixed and the user can freely control the apparent rate of property motion in the flow through l_{\max} . For time-dependent vector fields, however, the spatial integration step size and time increment between consecutive vector fields must respect a fixed ratio so that resulting animation conform to the physical flow.

Spatially, l_{\max} is related to the physical dimensions by

$$\frac{L_{\max} N_{\text{images}}}{W} = \frac{V_{\phi \max} t_{\phi}}{W_{\phi}},$$

where a subscript ϕ refers to a physical variable. $V_{\phi \max}$, t_{ϕ} , and W_{ϕ} represent, respectively, the maximum velocity, the total duration, and the width of the domain for the physical phenomenon. If we fix l_{\max} and the image resolution W , we deduce the number of frames N_{images} necessary to visualize the phenomenon with such constraints.

Temporally, we often have a set of N_{vf} instantaneous vector fields. From them, a total of N_{images} vector fields (twice that for second order integration) are interpolated at constant intervals h_{vf} :

$$h_{\text{vf}} = (N_{\text{vf}} - 1) / (N_{\text{images}} - 1).$$

While parameters conformity with physical dimensions is not essential in the case of noise-based animation, it becomes necessary for accurate animation of dye spreading.




Fig. 13. Effect of dye width. (a), (b) Sixteen pixels wide and (c), (d) eight pixels wide. (a), (c) without and (b), (d) with diffusion correction.

For instance, for the case of the cylinder (Fig. 14), the physical dimensions are $V_{\phi \max} = 9.9m/s$, $t_{\phi} = 4s$, and $W_{\phi} = 8m$. The data set is composed of 16 vector fields. If we choose $l_{\max} = 6$ pixels, the resulting animation will show the entire phenomenon in 248, 495, or 743 frames for image resolutions of 300, 600, and 900 pixels, respectively. Respectively, about 15, 30, and 45 interpolated vector fields would be required between the vector fields available in the data set.

7 RESULTS

The next section presents timings of our algorithm. We conducted experiments to evaluate the efficiency of the algorithm at four resolutions (300^2 through 1000^2 pixels). We present in Table 1 timings in frames/second, using several of the available options. Three different computers were used.

The organization of the algorithm as a series of array operations makes it particularly straightforward to parallelize on shared memory architectures. Furthermore, operations on the array elements only make accesses within h rows or columns. For small h , the locality of these accesses is sufficient not to produce cache misses on a CPU with a cache of moderate size (e.g., 512 kbytes). Table 1 also includes timings from an OpenMP implementation running on four processors of an Onyx2.

7.1 Examples

Fig. 14 shows two animations of flow around a circular cylinder. Dye is released at 10 evenly spaced points along the line $x = 0.05$ (the horizontal scale is normalized to 1). The dye thickness is 32 pixels. We compare the animations with $s = 0$ (Fig. 14 left) and $s = 0.2$ (Fig. 14 right). The dye is released once every 20 frames with alternating colors. The effect of the sharpening filter is evident. Depending on how the dye is released, additional information can be extracted from the flow. In this case, as the fluid elements are affected by the velocity gradient tensor, the elements distort and change direction. This can be quantified by measuring the




Fig. 14. Animation frames of flow around a circular cylinder. Dye is released periodically and with alternating colors from 10 locations evenly placed along a vertical strip. (left row) Sharpness $s = 0$ and (right row) $s = 0.2$.

change in area and shape in a postprocessing step. Only dye advection is capable of providing this type of data. Note that continuous dye release alone cannot provide this information.

Another example of dye injection is to release it periodically into the flow along a curve (Fig. 15). Ocean currents in the Gulf of Mexico are emphasized by releasing dye along two vertical lines every 60 frames. The antidiffusion filter is turned off to produce a smoke-like effect. The structure of the flow is nicely put into evidence by the resulting timelines.




Fig. 15. Three frames from an animation of flow in the Gulf of Mexico. Colored dye 32 pixels wide is released every 60 frames along two vertical strips.

We believe that novel combinations in the use of dye, including release frequency, location, and dye color, combined with new quantitative analysis techniques of the resulting images will help better understand the dynamics of the flow and should prove useful to the flow modeling community.

8 CONCLUSION

We presented a software algorithm for the visualization of unsteady vector fields, based on a combination of Euler and Lagrangian advection of noise textures. We paid particular attention to the treatment of edge effects, uniformity of the noise textures, and spatial and temporal coherence. Various postprocessing options were considered, including a fast line integral convolution technique to remove aliasing artifacts and masking to enhance user-selectable flow features.

Dye advection strategies come at no extra cost through the consideration of smooth textures. They are cheaper to implement than noise advection due to the inherent spatial coherence built into the textures. A new filter was proposed to counter the inherent diffusion of dye inherent in our integration scheme.

ACKNOWLEDGMENTS

The authors would like to thank David Banks for lively discussions in all areas of visualization, including several valuable suggestions to improve the quality of this paper. Some of the data sets used to illustrate the techniques presented in this paper were provided courtesy of Z. Ding (FSU), J. O'Brien (COAPS, FSU), J. Quiby (MeteoSwiss, Switzerland), and R. Arina (Ecole Polytechnique of Turin, Italy). We acknowledge the support of the US National Science Foundation under grant NSF-9872140.

REFERENCES

- [1] B.G. Becker, D.A. Lane, and N.L. Max, "Unsteady Flow Volumes," *Proc. IEEE Visualization '95*, G.M. Nielson and D. Silver, eds., Oct. 1995.
- [2] B. Cabral and L.C. Leedom, "Imaging Vector Fields Using Line Integral Convolution," *Computer Graphics Proc.*, J.T. Kajiya, ed., pp. 263-269, Aug. 1993.
- [3] W.C. de Leeuw and R. van Liere, "Spotting Structure in Complex Time Dependent Flow," technical report, CWI—Centrum voor Wiskunde en Informatica, Sept. 1998.
- [4] W. Heidrich, R. Westermann, H.-P. Seidel, and T. Ertl, "Applications of Pixel Textures in Visualization and Realistic Image Synthesis," *Proc. ACM Symp. Interactive 3D Graphics*, pp. 127-134, Apr. 1999.
- [5] M.Y. Hussaini, G. Erlebacher, and B. Jobard, "Real-Time Visualization of Unsteady Vector Fields," *Proc. 40th AIAA Aerospace Sciences Meeting*, 2002. AIAA J., submitted.
- [6] B. Jobard, G. Erlebacher, and M.Y. Hussaini, "Tiled Hardware-Accelerated Texture Advection for Unsteady Flow Visualization," *Proc. Graphicon 2000*, pp. 189-196, Aug. 2000.
- [7] B. Jobard, G. Erlebacher, and M.Y. Hussaini, "Hardware-Accelerated Texture Advection for Unsteady Flow Visualization," *Proc. Visualization 2000*, T.E. Ertl, B. Hamann, and A. Varshney, eds., pp. 155-162, Oct. 2000.
- [8] B. Jobard, G. Erlebacher, and M.Y. Hussaini, "Lagrangian-Eulerian Advection for Unsteady Flow Visualization," *Proc. IEEE Visualization 2001*, Oct. 2001.
- [9] D.A. Lane, "UFAT—A Particle Tracer for Time-Dependent Flow Fields," *Proc. IEEE Visualization '94*, R.D. Bergeron and A.E. Kaufman, eds., pp. 257-264, 1994.
- [10] D.A. Lane, "Visualizing Time-Varying Phenomena in Numerical Simulations of Unsteady Flows," NASA Ames Research Center, Feb. 1996.
- [11] N. Max, B. Becker, and R. Crawfis, "Flow Volumes for Interactive Vector Field Visualization," *Proc. IEEE Visualization '94*, pp. 19-24, 1994.
- [12] N. Max and B. Becker, "Flow Visualization Using Moving Textures," *Data Visualization Techniques*, Chandrajit Bajaj, ed., pp. 99-105, John Wiley & Sons, 1999.
- [13] M. Rumpf and J. Becker, "Visualization of Time-Dependent Velocity Fields by Texture Transport," *Proc. Eurographics Workshop Scientific Visualization '98*, pp. 91-101, 1998.
- [14] H.-W. Shen, C.R. Johnson, and K.L. Ma, "Visualizing Vector Fields Using Line Integral Convolution and Dye Advection," *Proc. Symp. Volume Visualization*, pp. 63-70, 1996.
- [15] H.-W. Shen and D.L. Kao, "A New Line Integral Convolution Algorithm for Visualizing Time-Varying Flow Fields," *IEEE Trans. Visualization and Computer Graphics*, vol. 4, no. 2, pp. 98-108, Apr.-June 1998.
- [16] D. Stalling and H.-C. Hege, "Fast and Resolution Independent Line Integral Convolution," *Proc. SIGGRAPH '95*, pp. 249-256, 1995.
- [17] M. van Dyke, *An Album of Fluid Motion*. Stanford, Calif.: The Parabolic Press, 1982.
- [18] D. Weiskopf, M. Hopf, and T. Ertl, "Hardware-Accelerated Visualization of Time-Varying 2D and 3D Vector Fields by Texture Advection via Programmable Per-Pixel Operations. Vision, Modeling, and Visualization," *Vision, Modeling, and Visualization (VMV '01) Conf. Proc.*, T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, eds., pp. 439-446, 2001.



Bruno Jobard received the PhD degree in computer science from the University of Littoral Cote d'Opale, France, in 2000. He was a postdoctoral fellow at Florida State University from August 1999 to April 2001. Currently, he holds the position of researcher at the Swiss Center for Scientific Computing (Switzerland). His research interests concern flow visualization algorithms.



Gordon Erlebacher received the undergraduate degree from the Polytechnique Institute at the Free University of Brussels in 1979 and the PhD degree from Columbia University in 1983 in applied physics. He spent 13 years at the NASA Langley Research Center at ICASE working on the numerical simulation of flow transition, flow turbulence, and flow visualization. He is a professor at Florida State University in the Department of Mathematics and in the School of Computational Science & Information Technology. His current interests are in large-scale numerical simulations, real-time/interactive scientific visualization, and face recognition.



M. Yousuff Hussaini received the PhD degree in engineering from the University of California at Berkeley in 1970 and the BS and MS degrees in mathematics and physics from the University of Madras. He is a professor of mathematics, holder of the TMC Eminent Scholar Chair in High Performance Computing, Affiliate Professor of Mechanical Engineering, Industrial Engineering and Computer Science at Florida State University. He helped set up the School of Computational Science and Information Technology, for which he was awarded the honorary title of Founding Director. Before coming to Florida State University in 1996, he was the director of the Institute for Computer Applications in Science and Engineering (ICASE) at the NASA Langley Research Center, Hampton, Virginia.

► For more information on this or any computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.