

# Visualization of Advection-Diffusion in Unsteady Fluid Flow

Grzegorz Karol Karch<sup>1</sup>, Filip Sadlo<sup>1</sup>, Daniel Weiskopf<sup>1</sup>, Claus-Dieter Munz<sup>2</sup>, and Thomas Ertl<sup>1</sup>

<sup>1</sup>Visualization Research Center, University of Stuttgart, Germany

<sup>2</sup>Institute of Aerodynamics and Gas Dynamics, University of Stuttgart, Germany

---

## Abstract

*Advection has been the standard transport mechanism in flow visualization. Diffusion, in contrast, has not been considered important in visual flow field analysis so far, although it is inherent to many physical processes. We present a novel technique that allows for interactive 3D visualization of both advection and diffusion in unsteady fluid flow. We extend texture-based flow visualization, which is advection-oriented, by diffusion. Our finite volume approach based on WENO (weighted essentially non-oscillatory) reconstruction is well parallelizable and features low numerical diffusion at interactive rates. Our scheme contributes to three different applications: (a) high-quality dye advection at low numerical diffusion, (b) physically-based dye advection accounting for diffusivity of virtual media, and (c) visualization of advection-diffusion fluxes in physical media where the velocity field is accompanied by a concentration field. Interactive rendering of the virtual dye is accomplished by ray casting. We apply our GPU implementation to CFD examples of thermal convection and evaporation phenomena.*

Categories and Subject Descriptors (according to ACM CCS): I.6.6 [Simulation and Modeling]: Simulation Output Analysis—; J.2 [Physical Sciences and Engineering]: Physics—

---

## 1. Introduction

Flow visualization plays an important role in science and engineering. In particular, the field of computational fluid dynamics (CFD) produces data of ever growing size and complexity, and visualization serves as one of the primary means of obtaining qualitative insight into these data. Many visualization techniques rely on characteristic curves of vector fields, such as stream lines and path lines—from direct visualization of individual curves or sets of curves [TB96], over surfaces [Hu192, GKT\*08] and volumes constructed from these field lines [XZC04], to implicit representations by scalar fields derived from them [vW93, WJE00]. Whereas these techniques proved to be successful in a wide area of applications where the purely advective transport property of the field is of interest, only few techniques have been devised so far for the visualization of diffusion effects. It is the aim of this paper to fill this gap—by providing a technique for appropriate analysis of flow fields where diffusion is taken into account. Advection-diffusion processes are widely present, besides diffusion of velocity in standard Navier-Stokes solutions due to viscosity, they gain importance when additional diffusive quantities such as temperature or concentration of soluble substances are involved.

Dye advection is a prominent technique for visualizing

advective processes in flows. It visualizes the transport of virtually inserted dye and hence represents the computational counterpart to injection of ink or smoke in physical experiments. In fluid dynamics, this simulation of transport of quantities due to prescribed velocity is called *passive advection*. Both in simulation and visualization major effort is spent to avoid the involved *numerical diffusion*, i.e., artificial diffusion of the quantity due to numerics—primarily caused by repeated interpolation of the quantity during advection in discrete grids. Although numerical diffusion tends to introduce blur, dye advection has the advantage that, in contrast to, e.g., streak surface generation, no geometric representation has to be maintained, allowing for visualization of arbitrarily complex flows. Whereas traditional passive advection solves the *advection* equation for prescribed flows, we solve for the *advection-diffusion* equation, accounting for diffusion processes. To obtain appropriate visualization of advection-diffusion data, it is important to keep numerical diffusion low during visualization. We achieve this by adopting a finite volume approach based on weighted essentially non-oscillatory (WENO) reconstruction. Our technique contributes to three different scenarios:

- Interactive dye advection in 3D flow fields at high quality and low numerical diffusion, requiring velocity data only.

- Physically plausible dye advection including diffusion of dye. By adjusting diffusivity to that of ink in water or smoke in air, one obtains the computational counterpart to these common experiments. This also requires only velocity data. If the flow field resulted from an advection-diffusion simulation with known diffusivity, the diffusivity of the virtual dye can be set accordingly for inspection.
- Visualization of transport due to diffusion. We introduce the novel concept of *passive diffusion*, the advection of dye according to the diffusion flux derived from concentration data. It can be either used alone for visualization of diffusion, or in combination with the velocity data for visualizing advection-diffusion. Besides these fields, it requires the constant of diffusivity of the concentration data.

## 2. Related Work

In scientific visualization, one line of research adopts a Lagrangian view on texture-based flow visualization—with line integral convolution (LIC) being their role model [CL93]. The other area of research, which is relevant to this paper, advects a texture from frame to frame. The majority of previous work in this research direction is based on semi-Lagrangian advection or similar schemes. For example, texture advection techniques address 2D visualization [JEH00, vW02, WHE01], 3D visualization [TvW03, WSE07], and visualization on surfaces [LJH03, vW03, LTWH08]. It is also possible to combine dye advection techniques with LIC, as described by Shen et al. [SJM96]. An overview of texture-based flow visualization techniques in general, including further references to prior work, is provided by Laramee et al. [LHD\*04]. A serious problem of semi-Lagrangian advection is the high level of numerical diffusion introduced by repeated resampling of the transported texture. In particular, resampling with bilinear or trilinear interpolation leads to strong blurring, which can be understood from a signal-processing perspective [Wei09]. Therefore, higher-order reconstruction filters can reduce blurring [AB06]. An alternative approach adopts the concept of level-set advection to avoid blurring for advected dye. In this approach, the boundary between dye and background is modeled as an interface transported without blur [CKSW08, Wei04].

However, none of the above methods from scientific visualization modeled diffusion explicitly. In fact, there is previous work that uses diffusion for flow visualization. For example, Sanderson et al. [SJK04] employ reaction-diffusion for flow visualization, adapting reaction-diffusion methods for generic texture synthesis in computer graphics [Tur91, WK91]. Similarly, Markov random field texture synthesis can be adopted for flow visualization [TA03]. Also, anisotropic (non-physical) diffusion, known from image filtering, can be applied to flow visualization [DPR00]. An extension to unsteady flow is provided in [BPR01], by adapting the diffusion tensor and blending the transport diffusion

evolution results started at successively incremented times. Finally, there is a large body of research on diffusion tensor visualization, mostly in medical imaging and visualization [VZKL06]. It is important to note that none of these diffusion-related papers use physical diffusion in combination with physical advection.

Most closely related work is the physically-based dye advection by Li et al. [LTH08]. They applied the piecewise parabolic method [CW84] (PPM) for visualization by dye advection in time-dependent 2D flow fields, providing a technique exhibiting low numerical diffusion. The technique allows for comparably large time steps by advecting each cell backward in time and sampling a parabolic reconstruction of the dye inside the resulting polygon. Several reasons motivated our choice for the weighted essentially non-oscillatory scheme [LOC94] (WENO) instead: it represents a generalization to arbitrary order of accuracy, is based on a clear mathematical foundation, lends itself better to parallelization, and, last but not least, allows for the incorporation of diffusion using a finite volume formulation.

We base our approach on the WENO scheme as described by Dumbser and Käser [DK07]. We perform dimensional splitting, i.e., we apply the 1D scheme subsequently in the three spatial dimensions. For incorporating active diffusion, we follow Chou and Shu [CS07] and use linear instead of WENO weights for computing the concentration gradient.

## 3. Advection and Diffusion

Dye advection in visualization, i.e., the simulation of the behavior of virtual dye inside prescribed (time-dependent) vector fields, is called *passive advection* in computational fluid dynamics. It leads to a linear problem that contrasts *active advection*, the advection of velocity itself during flow simulation resulting in a nonlinear problem that is harder to solve. This is one of the reasons why visualization by dye advection tends to be fast, compared to flow simulation itself. Nevertheless, numerical diffusion is a major problem also with passive advection and substantial effort is taken to reduce it, as, for example, in the dye advection method by Li et al. [LTH08], typically at the cost of reduced performance.

If physical diffusion is included in the advection problem, it typically leads to the advection-diffusion equation. Whereas we address passive advection only, *diffusion* traditionally plays an active role in science and engineering, i.e., it is the concentration variation of the quantity itself that governs its diffusion. If the virtual dye takes the role of the quantity, we obtain physical dye advection accounting for diffusion of solubles in typical flow media. We call this approach *active diffusion* to contrast it from our second variant that we call *passive diffusion*. We introduce passive diffusion as a means of visualizing the mechanisms “behind” concentration changes due to diffusion. In this case, the flow field needs to be accompanied by a (time-dependent) concentration field for analysis by our technique.

### 3.1. Advection-Diffusion Equation

The traditional advection-diffusion equation [BSL60] reads

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} = D \Delta \phi \quad (1)$$

with the concentration  $\phi = \phi(\mathbf{x}, t)$ , velocity  $\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ , constant of diffusivity  $D$ , and Laplacian  $\Delta$ . To allow for passive diffusion, introduced in Section 3.3, we generalize Eq. 1:

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} + (\nabla \phi) \mathbf{j} = D_\phi \Delta \phi \quad (2)$$

by introducing flux density  $\mathbf{j}$ , and account for diffusion flux:

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} - (\nabla \phi) D_\psi \nabla \psi = D_\phi \Delta \phi \quad (3)$$

with  $\phi$  being the concentration of the quantity of interest (here the virtual dye), its constant of diffusivity  $D_\phi$ ,  $\psi = \psi(\mathbf{x}, t)$  being the concentration field (e.g., of solubles or quantities such as heat) governing the passive diffusion process, and its constant of diffusivity  $D_\psi$  (please see Section 3.3 for a thorough explanation). Here,  $\mathbf{j} = -D_\psi \nabla \psi$  represents the diffusion flux due to the concentration variation of the quantity  $\psi$ , according to Fick's law [Fic55]. Please note that our implementation solves Eq. 3 for dye concentration  $\phi$  only;  $\mathbf{u}$ ,  $\psi$ , and  $D_\psi$  are obtained from simulation data. In our formulation of passive diffusion, we aim at visualizing the transport of  $\psi$ , by making the virtual dye move according to it. This is achieved by using  $\mathbf{j}$  as another advective mechanism for the transport of  $\phi$  as described by Eq. 2. That is, the concentration gradient of  $\psi$  induces (additional) transport of the virtual dye.


### 3.2. Active Diffusion

In the traditional case of diffusion that we call *active*, the diffusion is governed solely by the gradient of the virtual dye itself, hence  $D_\psi = 0$  and Eq. 3 simplifies to the traditional advection-diffusion equation:

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} = D_\phi \Delta \phi. \quad (4)$$

This variant builds only on the velocity field—it does not require accompanying time-dependent simulation data of the concentration field. In general flow visualization, diffusivity  $D_\phi$  can be chosen, e.g., from physics textbooks, to achieve physically correct interaction of the fluid and tracers used in experiments, such as smoke in the analysis of automotive design (Figure 1). If the velocity field resulted from a flow simulation that included diffusion, i.e., an advection-diffusion problem, and if the used diffusivity is known, it can be used as  $D_\phi$  to obtain corresponding dye behavior. Although we restrict ourselves here to isotropic diffusion, it is worth noticing that our finite volume approach also allows for, e.g., data-driven anisotropic diffusion. Setting  $D_\phi = 0$  turns Eq. 4 into traditional (passive) dye advection (Figure 2):

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} = 0. \quad (5)$$



**Figure 1:** Flow over a “blunt body” (front to back). Advection of dye without diffusion model (top, with a streak line for comparison). Active diffusion (bottom) mimics diffusion of smoke in experimental analysis. Resolution  $600 \times 125 \times 121$ .

Due to the low numerical diffusion of the WENO scheme (Section 4.1) and its ability to provide interactive dye advection in 3D, this reduced mode, to the best of our knowledge, is already a new contribution to the field of dye advection.

### 3.3. Passive Diffusion

Our central contribution is the visualization of transport of quantities due to diffusion. It requires an accompanying concentration field  $\psi$ , which often represents temperature or solubles such as ink or vapor, together with its constant of diffusivity,  $D_\psi$ . If passive diffusion is not combined with active diffusion ( $D_\phi = 0$ ), as in our experiments, Eq. 3 simplifies to

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) \mathbf{u} - (\nabla \phi) D_\psi \nabla \psi = 0. \quad (6)$$


Interestingly, our passive advection-diffusion model can be incorporated into the traditional passive advection scheme (Eq. 5) by combining advection flux and diffusion flux:

$$\frac{\partial \phi}{\partial t} + (\nabla \phi) (\mathbf{u} - D_\psi \nabla \psi) = 0. \quad (7)$$

In Section 4.2, we describe in detail how our approach is formulated in terms of the finite volume scheme. Interestingly, the passive diffusion term is implemented similarly to the active diffusion term there, not the advection term.

### 3.4. Numerical Diffusion

So far, we neglected inaccuracies resulting from diffusion, either on the simulation or visualization side. Grid-based CFD faces the serious issue that accuracy of advection of quantities such as heat or solubles highly depends on grid resolution, i.e., undersampling and numerical diffusion are omnipresent problems, typically leading to results that deviate substantially from the true physical behavior. Imagine



**Figure 2:** Unsteady buoyant air flow (CFD) in a closed container heated at bottom (red plate,  $75^{\circ}\text{C}$ ) and cooled at the top (blue plate,  $5^{\circ}\text{C}$ ). Two isosurfaces, one at  $38^{\circ}\text{C}$  (blue) and one at  $42^{\circ}\text{C}$  (red). Advected dye (green, without diffusion model) seeded at center (black box). Finite volume method using WENO reconstruction at both dye resolutions  $244 \times 124 \times 244$  (a) and  $122 \times 62 \times 122$  (b) exhibits much lower numerical diffusion than trilinear reconstruction at  $244 \times 124 \times 244$  (c).

a flow exhibiting foliation, i.e., repeated thinning and folding of the fluid, transforming an initially compact region that contains a quantity of arbitrarily many and arbitrarily finely folded sheets. Even if there would be no numerical diffusion during advection, already the sampling of these sheets would require unmanageably fine resolution.

Unfortunately, numerical diffusion (including undersampling) is hard to model. It depends, among other things, on the varying speed and direction at which the quantity is propagated through each cell of the grid during advection. It is therefore not practical to quantify numerical diffusion, neither during simulation, nor during dye advection.

All our three scenarios are subject to numerical diffusion. Numerical diffusion of the dye  $\phi$  is, however, kept comparably low due to WENO reconstruction. In the scenario of dye advection without diffusion modeling (Eq. 5), the issue is, like in other texture advection techniques, usually not addressed. In the scenario of dye advection with active diffusion (Eq. 4), the dye  $\phi$  is subject to both active diffusion due to diffusivity  $D_{\phi} > 0$  and numerical diffusion. Hence, the effective diffusion of  $\phi$  tends to exceed that prescribed by  $D_{\phi}$ . Since numerical diffusion is not quantified, one approach to judge its influence is to compare the result with dye advection using  $D_{\phi} = 0$ . We compare the results visually: small difference indicates that numerical diffusion affects (is in the order of) the one modeled by  $D_{\phi}$ . Finally, in the scenario of passive diffusion (Eq. 6), one can make use of the concentration field  $\psi$  from simulation data if there are identifiable sources therein. For the example of the Buoyant Flow data set, one can obtain the region where the room is heated by applying a threshold filter to the  $\psi$  field. By continuously seeding dye in this region and using our active diffusion model (Eq. 4), one can adjust  $D_{\phi}$  until  $\phi$  matches the (time-dependent)  $\psi$  field, see Figure 3. This can compensate for inappropriate  $D_{\psi}$  due to both numerical diffusion in the simulation and dye advection. The obtained  $D_{\phi}$  can then also be used as  $D_{\psi}$  in our passive diffusion model.

#### 4. WENO-Based Finite Volume Method

As we have seen in Sections 3.2 and 3.3, the three different modes of our scheme (traditional dye advection, active diffusion, and passive diffusion) can be accomplished by solving for Eq. 4 (passive advection with active diffusion). Since diffusion requires comparably small time steps, Eq. 4 lends itself well to solution by the finite volume method, where the temporal change of the quantity within a cell is modeled by its fluxes over the boundaries of the cell. The quantity itself is represented in a per-cell manner. Hence, the evaluation of the necessary fluxes at the cell boundaries requires reconstruction. As demonstrated in Figure 2, tensor-product linear interpolation leads to excessive blur. The PPM addresses this problem by parabolic reconstruction within the cell. As demonstrated in Figures 7 and 8, WENO reconstruction achieves even better results. Furthermore, it allows for dimensional splitting (see Section 4.2) which simplifies implementation, and more important, substantially accelerates the dye advection, allowing for interactive rates.


We first describe WENO reconstruction in 1D (Section 4.1), then we detail how Eq. 4 is formulated using finite volume in 1D using WENO reconstruction (Section 4.2). Finally we describe our GPU implementation (Section 4.3).

##### 4.1. WENO Reconstruction

The finite volume formulation requires the evaluation of fluxes at the boundaries of the cells, necessitating the reconstruction of  $\phi$  within each cell. The advection flux is simply its value at the boundary multiplied by the velocity component normal to the boundary, whereas the diffusion flux requires the evaluation of its gradient at the boundary, i.e., the derivative of  $\phi$ , in direction normal to the boundary.

We describe WENO reconstruction with respect to our application based on a 1D finite volume. The reader is referred to [DK07] for a thorough introduction to the topic and further details. We describe the reconstruction for the example





**Figure 3:** Diffusivity adjustment in buoyant flow example (vertical cross section, low temperature - blue, high temperature - red). Virtual dye (green, mixed with color from temperature) seeded at hot plate at lower image border. Using  $D_\phi = 1.11 \cdot 10^{-5}$  (a),  $D_\phi = 1.65 \cdot 10^{-5}$  (b), and (theoretical) diffusivity from simulation  $D_\phi = 2.19 \cdot 10^{-5}$  (c). In (a) the dye region is too small, in (c) too large, and in (b) the reduced diffusivity compensates numerical diffusion well.

of third-order accurate reconstruction (using second-order polynomials). The extension to higher degrees is straightforward [DK07]. In our experiments, we obtained a good trade-off between efficiency and accuracy using fourth-order accurate WENO, i.e., cubic polynomials.

Let us consider second-order reconstruction polynomials

$$\phi_k(x) = \sum_{j=0}^2 \hat{w}_j^k x^j$$

with  $k$  being the index of their central cell, and coefficients  $\hat{w}_j^k$ , as shown in Figure 4. The coefficients have to be chosen conservatively, i.e., the integrals of  $\phi_k$  over the cells have to be identical to the cell-centered values  $\bar{\phi}_i$ :

$$\int_{x_{i-\frac{3}{2}}}^{x_{i-\frac{1}{2}}} \phi_k dx = \bar{\phi}_{i-1}, \quad \int_{x_{i-\frac{1}{2}}}^{x_{i+\frac{1}{2}}} \phi_k dx = \bar{\phi}_i, \quad \int_{x_{i+\frac{1}{2}}}^{x_{i+\frac{3}{2}}} \phi_k dx = \bar{\phi}_{i+1}.$$

This represents


$$\begin{aligned} \hat{w}_0^k(x_{i-\frac{1}{2}} - x_{i-\frac{3}{2}}) + \frac{\hat{w}_1^k}{2}(x_{i-\frac{1}{2}}^2 - x_{i-\frac{3}{2}}^2) + \frac{\hat{w}_2^k}{3}(x_{i-\frac{1}{2}}^3 - x_{i-\frac{3}{2}}^3) &= \bar{\phi}_{i-1} \\ \hat{w}_0^k(x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}) + \frac{\hat{w}_1^k}{2}(x_{i+\frac{1}{2}}^2 - x_{i-\frac{1}{2}}^2) + \frac{\hat{w}_2^k}{3}(x_{i+\frac{1}{2}}^3 - x_{i-\frac{1}{2}}^3) &= \bar{\phi}_i \\ \hat{w}_0^k(x_{i+\frac{3}{2}} - x_{i+\frac{1}{2}}) + \frac{\hat{w}_1^k}{2}(x_{i+\frac{3}{2}}^2 - x_{i+\frac{1}{2}}^2) + \frac{\hat{w}_2^k}{3}(x_{i+\frac{3}{2}}^3 - x_{i+\frac{1}{2}}^3) &= \bar{\phi}_{i+1} \end{aligned}$$

and can be formulated as a matrix-vector product:

$$\mathbf{L}_k \begin{pmatrix} \hat{w}_0^k \\ \hat{w}_1^k \\ \hat{w}_2^k \end{pmatrix} = \begin{pmatrix} \bar{\phi}_{i-1} \\ \bar{\phi}_i \\ \bar{\phi}_{i+1} \end{pmatrix}$$

with

$$\mathbf{L}_k = \begin{pmatrix} (x_{i-\frac{1}{2}} - x_{i-\frac{3}{2}}) & \frac{1}{2}(x_{i-\frac{1}{2}}^2 - x_{i-\frac{3}{2}}^2) & \frac{1}{3}(x_{i-\frac{1}{2}}^3 - x_{i-\frac{3}{2}}^3) \\ (x_{i+\frac{1}{2}} - x_{i-\frac{1}{2}}) & \frac{1}{2}(x_{i+\frac{1}{2}}^2 - x_{i-\frac{1}{2}}^2) & \frac{1}{3}(x_{i+\frac{1}{2}}^3 - x_{i-\frac{1}{2}}^3) \\ (x_{i+\frac{3}{2}} - x_{i+\frac{1}{2}}) & \frac{1}{2}(x_{i+\frac{3}{2}}^2 - x_{i+\frac{1}{2}}^2) & \frac{1}{3}(x_{i+\frac{3}{2}}^3 - x_{i+\frac{1}{2}}^3) \end{pmatrix}.$$



**Figure 4:** Example of 1D WENO reconstruction. Cell centers  $x$  with cell-averaged values  $\bar{\phi}$  and reconstruction polynomials  $\phi(x)$ . Resulting WENO reconstruction  $\phi^{WENO}(x)$  for cells  $i-1$  and  $i$  (bold). Note that neither  $\phi(x)$  nor  $\phi^{WENO}(x)$  pass through the values at the centers.

During dye advection, the coefficients  $\hat{w}_j^k$  are determined by

$$\begin{pmatrix} \hat{w}_0^k \\ \hat{w}_1^k \\ \hat{w}_2^k \end{pmatrix} = \mathbf{L}_k^{-1} \begin{pmatrix} \bar{\phi}_{i-1} \\ \bar{\phi}_i \\ \bar{\phi}_{i+1} \end{pmatrix}.$$

The reconstruction  $\phi_i^{WENO}(x)$  of the concentration  $\phi(x)$  inside cell  $i$  is a linear combination of the polynomials  $\phi_k(x)$ :

$$\phi_i^{WENO}(x) = \sum_{k=i-1}^{i+1} \omega_k \phi_k(x)$$

with weights  $\omega_k = \tilde{\omega}_k / \sum_{j=i-1}^{i+1} \tilde{\omega}_j$  using  $\tilde{\omega}_l = \lambda_l / (\varepsilon + \sigma_l)^r$ . As reported in [DK07] we use  $\varepsilon = 10^{-5}$ ,  $r = 4$ , and

$$\lambda_l = \begin{cases} 10^3 & \text{if } l = i, \\ 1 & \text{otherwise.} \end{cases}$$


The oscillation indicators  $\sigma_l$  can be obtained from a matrix  $\Sigma$  by  $\sigma_l = \hat{\mathbf{w}}^l \Sigma \hat{\mathbf{w}}^l$ . In our third-order WENO example,

$$\Sigma_{mn} = \sum_{r=1}^2 \int_V \frac{\partial^r x^m}{\partial x^r} \frac{\partial^r x^n}{\partial x^r} dx.$$

Both  $\Sigma$  and  $\mathbf{L}_k^{-1}$  are precomputed, e.g., using a computer algebra system, since, after transformation to a reference space, they neither depend on the mesh nor on the problem.

## 4.2. Finite Volume Method

On uniform (or rectilinear) grids the WENO reconstruction and the whole finite volume approach can be formulated in an efficient manner using dimensional splitting. Since cell face normals coincide with coordinate axes in these grids, one can apply the 1D-procedure sequentially in  $x$ -,  $y$ -, and  $z$ -direction to accomplish a time step of the 3D dye advection (see [JS96]). Hence, the 1D reconstruction described in Section 4.1 together with the 1D finite volume approach described below are applied for each direction, denoted here as  $x$  with respective velocity component  $u$ .



**Figure 5:** Riemann solution for fluxes ((a) and (b)). WENO reconstruction (bold) exhibits discontinuities at cell faces. Flux (transparent rectangle) is determined from advection direction  $u$ , choosing “donor” side. For diffusion (c), gradient at faces is computed from central reconstruction polynomial (green, blue) and averaged (black) (cf. Figure 4).

Let  $\bar{\phi}_i(t)$  be the amount of dye inside cell  $i$  at time  $t$ , and  $\bar{\phi}_i(t + \delta t)$  the amount after time step  $\delta t$ . The concentration change  $\delta\bar{\phi}_i = \bar{\phi}_i(t + \delta t) - \bar{\phi}_i(t)$  according to Eq. 5 computes in the finite volume scheme as  $\delta\bar{\phi}_i = (\bar{f}_{i+} - \bar{f}_{i-})\delta t$ , i.e., the balance over time  $\delta t$  between influx  $\bar{f}_{i+}$  of dye into cell  $i$  and outflux  $\bar{f}_{i-}$  of dye from cell  $i$ . For convergence,  $\delta t$  has to be chosen sufficiently small, usually by prescribing a Courant-Friedrichs-Lewy condition ( $CFL$ ) value.  $CFL$  relates to maximum velocity  $u_{\max}$  in the field and cell size  $h$  by  $CFL = u_{\max}\delta t/h$ .

In case of passive advection without diffusion (Eq. 5), we compute  $\bar{f}_{i+}$  and  $\bar{f}_{i-}$  from velocity  $u$  and concentration  $\phi_i^{\text{WENO}}$ . Please note that, according to Figure 5, the Riemann solution of  $\bar{f}_{i+}$  and  $\bar{f}_{i-}$  depends on the sign of  $u$  ( $u = u_{\text{sim}}$  is interpolated in space and time from the simulation data). Let  $x_{i+\frac{1}{2}} = x_{j-\frac{1}{2}}$  be the face between cell  $i$  and  $j$ . In case (a)

$$\bar{f}_{i-} = u \cdot \phi_i^{\text{WENO}}(x_{i+\frac{1}{2}}), \quad \bar{f}_{j+} = u \cdot \phi_i^{\text{WENO}}(x_{i+\frac{1}{2}})$$

and  $\bar{f}_{i+} = \bar{f}_{j-} = 0$ , whereas in case (b)  $\bar{f}_{i-} = \bar{f}_{j+} = 0$  and

$$\bar{f}_{i+} = u \cdot \phi_j^{\text{WENO}}(x_{j-\frac{1}{2}}), \quad \bar{f}_{j-} = u \cdot \phi_j^{\text{WENO}}(x_{j-\frac{1}{2}}).$$

In a straightforward approach  $\phi$  could be reconstructed at time  $t$ :  $\phi_{i,j}^{\text{WENO}}(x_{i+\frac{1}{2}}, t)$  and  $u$  could be time-averaged:  $u = [u_{\text{sim}}(x_{i+\frac{1}{2}}, t) + u_{\text{sim}}(x_{i+\frac{1}{2}}, t + \delta t)]/2$ . For a higher order of accuracy, and hence better convergence of the scheme allowing for larger time steps, we apply *prediction steps* instead (see below).

Next, we include active diffusion to obtain the finite volume formulation of Eq. 4. Applying the vector identity  $\Delta = \nabla \cdot \nabla$  to its right hand side we obtain  $D_\phi \Delta \phi = D_\phi \nabla \cdot \nabla \phi$ . Assuming uniform  $D_\phi$  and applying the Gauss theorem:

$$D_\phi \int_V \nabla \cdot \nabla \phi dV = D_\phi \int_{S=\partial V} \nabla \phi \cdot \mathbf{n} dS.$$

$D_\phi \nabla \phi \cdot \mathbf{n}$  is the diffusion flux  $D_\phi \nabla \phi$  of the virtual dye through the cell face with normal  $\mathbf{n}$ . In our 1D scheme and the example in Figure 5(c) this gives rise to diffusion fluxes  $\bar{d}_{i-} = \bar{d}_{j+} = D_\phi \partial \phi_{ij} / \partial x$ , with, for mass conservation,

$\partial \phi_{ij} / \partial x = (\partial \phi_i / \partial x + \partial \phi_j / \partial x) / 2$  (exchange  $i$  and  $j$  if opposite orientation of  $\partial \phi_{ij} / \partial x$ ). Hence, including active diffusion,

$$\delta\bar{\phi}_i = (\bar{f}_{i+} - \bar{f}_{i-} + \bar{d}_{i+} - \bar{d}_{i-})\delta t. \quad (8)$$

Finally, passive diffusion is included (Eq. 7), i.e., for passive diffusion we derive  $u = -D_\psi \partial \psi / \partial x$ , and for passive advection together with passive diffusion  $u = u_{\text{sim}} - D_\psi \partial \psi / \partial x$ . Again, in a straightforward approach  $D_\psi \partial \psi / \partial x$  could be time-averaged:  $D_\psi [\partial \psi(x_{i+\frac{1}{2}}, t) / \partial x + \partial \psi(x_{i+\frac{1}{2}}, t + \delta t) / \partial x] / 2$  from simulated  $\psi$  by interpolation, and  $u_{\text{sim}}$  also time-averaged, as above. However, we will use prediction steps here, too. Please note that whereas the concentration is evaluated from  $\phi^{\text{WENO}}$ , its gradient  $\partial \phi / \partial x$  is computed from the central polynomial  $\phi_k$  [CS07] (Figure 5).

To improve the accuracy of the finite volume scheme, we apply *prediction steps* instead of the straightforward approaches based on averaging, sketched above. The underlying idea is to obtain a better accuracy for the fluxes  $\bar{f}$  and  $\bar{d}$  within  $]t, t + \delta t]$  by generating predictions using integration. Whereas  $u_{\text{sim}}$  and  $\psi$  are simulated and hence available at any “future” time  $\tau \in ]t, t + \delta t]$ ,  $\bar{\phi}_i$  and hence  $\phi_i^{\text{WENO}}$  and  $\phi_i$  are only available at time  $t$ . However, Eq. 4 can be used to provide integration-based per-cell approximation values  $\tilde{\phi}_i(\tau)$  at time  $\tau$ . By rearranging Eq. 4 to

$$\tilde{\phi}_t := \partial \phi / \partial t = D_\phi \Delta \phi - (\nabla \phi) \mathbf{u}$$

one obtains an approximation  $\tilde{\phi}_t$  for  $\partial \phi / \partial t$  [DK07]. Each of the  $n$  prediction steps of size  $\delta t/n$  consists of evaluating  $\tilde{\phi}_t$  at the center of each cell  $i$  and estimating  $\tilde{\phi}_i(\tau + \delta t/n) = \tilde{\phi}_i(\tau) + \tilde{\phi}_t \delta t/n$ , in parallel for all cells. In our 1D case  $\tilde{\phi}_t = D_\phi \partial^2 \phi / \partial x^2 - \partial \phi / \partial x u$ . The procedure is repeated  $n$  times ( $n = 1$  in our experiments) and each time WENO reconstruction is applied to  $\tilde{\phi}_t(\tau)$  to obtain  $\tilde{\phi}_i^{\text{WENO}}(\tau)$  and  $\tilde{\phi}_i(\tau)$ . In each step,  $\bar{f}_{i+}(\tau)$  and  $\bar{f}_{i-}(\tau)$  as well as  $\bar{d}_{i+}(\tau)$  and  $\bar{d}_{i-}(\tau)$  are computed from  $\tilde{\phi}_i(\tau)$  and  $\tilde{\phi}_i^{\text{WENO}}(\tau)$ . During prediction steps, these fluxes are accumulated and finally used in Eq. 8 to obtain  $\bar{\phi}_i(t + \delta t) = \bar{\phi}_i(t) + \delta\bar{\phi}_i$ .

### 4.3. Implementation

The WENO-based finite volume method substantially benefits from GPU parallelization in that there is spatial and temporal locality in the algorithm. We employ the CUDA API (single-precision) to compute advection-diffusion and use data streaming for memory efficiency. Figure 6(a) describes the overall procedure. The input to our implementation includes a set of parameters (including the polynomial WENO order), can contain geometry objects to be rendered together with the ray-casted volume, and the vector (and scalar concentration) data. If the data are given on an unstructured grid, the interpolation weights for each node of our (arbitrary) uniform dye advection grid are precomputed and used for interpolating the data ( $\mathbf{u}$ ,  $\psi$ ) before transfer to the GPU, avoiding expensive point location at each time step. Device (GPU)

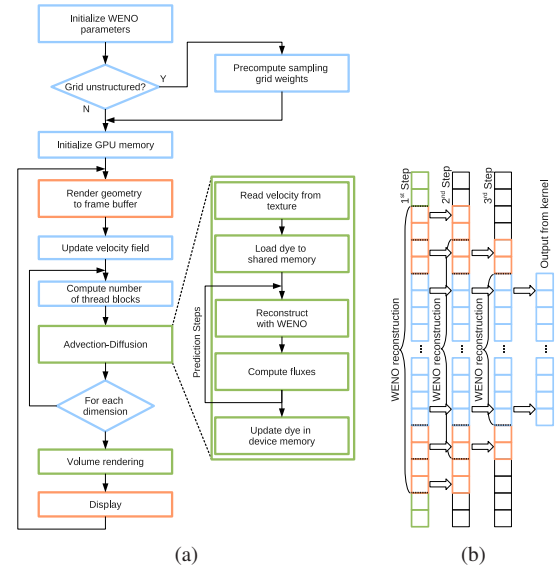
memory is allocated for the dye, the two consecutive time steps of simulation data, and the OpenGL buffers.

To increase performance, each GPU thread block loads an array of data from device memory into its shared memory along the current axis  $\alpha$  of dimensional splitting. The number of thread blocks equals  $b_d = r((d+1) \bmod 3) \times r((d+2) \bmod 3) \times \lceil r(d)/s \rceil$ , where  $r(d)$  is the resolution in dimension  $d$ , and  $s$  is the size of a thread block limited by GPU shared memory size. This constraint means that if the resolution along  $\alpha$  is too large, the arrays must be divided into overlapping subsets. The overlaps (ghost cells) are necessary to perform WENO reconstruction, and their size depends on the WENO order and the number of prediction steps (Figure 6(b)). The WENO reconstruction is carried out, and the advection and diffusion fluxes are computed. Finally, the dye concentration is updated in device memory. For optimization (early rejection), a data fragment processed by a thread block is skipped if no dye is present. In order to determine the presence of dye, parallel reduction (i.e., summation of array elements) is performed. Coalesced access, which significantly reduces transfer time between device memory and kernels, occurs only in computation in  $x$ -dimension. However, regardless of the dimension, transfer time is well hidden by computation of advection and diffusion. To avoid fluctuations in physical time during interaction, we apply respective corrections to the dye advection time step  $\delta t$ . For further details please refer to the supplementary material.

For volume rendering, ray casting is utilized, following the example in the NVIDIA's CUDA SDK. Our implementation additionally supports volume lighting, ray-casted iso-surfaces [Lev88] of the  $\psi$  field, and geometric objects.

## 5. Results

We demonstrate and evaluate our technique by visualizing advection-diffusion in three CFD examples. Additionally, we investigate the performance of our technique with respect to numerical diffusion, mass conservation, and speed. All examples were run on GeForce GTX 580 (3GB). The first CFD example, a buoyant flow inside a closed container, exhibits strong time-dependency and was simulated on a uniform grid with dimension  $61 \times 31 \times 61$  and 2000 time steps (spanning 50 seconds). Heat takes in this case the role of the diffusing quantity. The second example, a flow around a heating coil, is quasi-stationary and was simulated on an unstructured tetrahedral grid comprising 93227 cells. Also in this case we visualize the diffusion of heat. The last CFD example is the quasi-stationary simulation of an evaporating drop, conducted on a uniform grid with resolution  $192 \times 128 \times 128$ . In this case vapor takes the role of the diffusing quantity. All results were obtained with fourth-order WENO, i.e., using third-order polynomials. Please see also the accompanying video for more details on the results including performance information.




**Figure 6:** (a) Advection-diffusion procedure. Operations in blue blocks are performed on CPU, OpenGL is used in orange blocks, green blocks are done in CUDA on GPU. (b) Blocking if dye grid dimension exceeds GPU shared memory. Number of prediction steps and order of WENO reconstruction define the number of ghost cells. For efficiency, blocks should be large, e.g., contain 128 output cells.

### 5.1. Comparison to Physically-Based Dye Advection

We compare our dye advection scheme to the most closely related work, the physically-based dye advection (PBDA) due to Li et al. [LTH08]. For comparison, we reduced our approach to 2D and implemented PBDA on the GPU. In Zalesak's disk test for numerical diffusion, one revolution of a disk-shaped dye distribution is performed about an axis situated at its bottom (Figure 7). We used  $CFL = 0.4$  for our approach and  $CFL = 3.0$  for PBDA. Whereas in our technique  $CFL$  is limited only by an upper bound for stability reasons, we experienced artifacts when choosing small  $CFL$  with PBDA. As in [LTH08], we used  $20 \times 20$  samples per cell for this comparison. Both approaches perform comparably well for high resolutions. However, it is apparent that our WENO-based approach performs considerably better at low resolutions. Low numerical diffusion at low resolution is of major importance for our technique because today's graphics hardware does not yet allow for interactive rates at high resolution. Our approach also performs better in terms of mass conservation (Table 1). It is apparent that resolution plays a major role with respect to mass conservation only for finite volume based on bilinear interpolation (Lin-FV).

We also used a 2D CFD simulation of a von Kármán vortex street consisting of 151 time steps for comparing our approach to [LTH08] (Figure 8). In this case, we used  $60 \times 60$



**Figure 7:** Zalesak's disk test. Our WENO-based approach exhibits lower numerical diffusion than PBDA.

**Table 1:** Relative change in mass (%) for bilinear finite volume (Lin-FV), PBDA, and WENO-based finite volume in Zalesak's disk test (Figure 7).

Resolution	Lin-FV	PBDA	WENO
100x100	-5.5057	-4.9869	0.0016
200x200	-1.4950	-4.9018	0.0003
400x400	-0.2392	-5.1294	-0.0004


instead of  $20 \times 20$  samples per cell with a  $CFL$  of 6.0 in PBDA, and a  $CFL$  of 0.85 in our approach. Still, it is apparent that our approach exhibits lower numerical diffusion. Table 2 gives performance details for the von Kármán example. In favor of PBDA, we applied for the measurements  $20 \times 20$  samples per cell. Still, each dye advection step is one order of magnitude faster in our approach, but since  $CFL$  is larger in PBDA, the resulting times for advecting dye over a given physical time period are comparable. On the other hand, this causes our approach to produce one order of magnitude higher frame rates, allowing for *interactive* dye advection.

## 5.2. CFD Examples

Buoyant flow driven by a heat gradient is a prominent candidate for visualization by our technique because it exhibits both advection and diffusion of heat, see Figure 2 for an introduction. The data set exhibits a time-dependent circular

**Table 2:** Performance comparison between WENO and PBDA for the von Kármán vortex street. Time for one step (Comp. time) vs. the whole simulation (Total time).

Resolution	Comp. time [ms]		Total time [s]	
	WENO	PBDA	WENO	PBDA
100x300	2.84	19.93	4.74	4.68
500x1500	10.76	85.76	44.85	50.51
1000x3000	34.94	283.78	291.12	334.58
2000x6000	114.44	963.27	1907.15	2272.36



**Figure 8:** Dye advection in 2D von Kármán vortex street with our WENO-based finite volume method (top) and PPM (bottom). WENO exhibits lower numerical diffusion.

flow behavior. No diffusion model was used for the dye in this figure, it visualizes the mixing behavior, i.e., the stretching of the fluid into sheets, leading to foliation. The dye advection in Figure 9(a) shows advection alone (Eq. 5): the dye follows the hot plume upward and thereafter it is advected downward by the cold plume. Figure 9(b) shows only passive diffusion (Eq. 6 with  $\mathbf{u} = \mathbf{0}$ ): the dye, i.e., the heat, flows from the hot plate through the cold plume into the cold plate at the top of the container. Figure 9(c) shows the superposition of the two, the true advection-diffusion of heat: like in Figure 9(a), the dye is advected upward but then part of it is caught by the cold plate instead of being advected downward again by the cold plume. Please see also the performance details in Table 3.

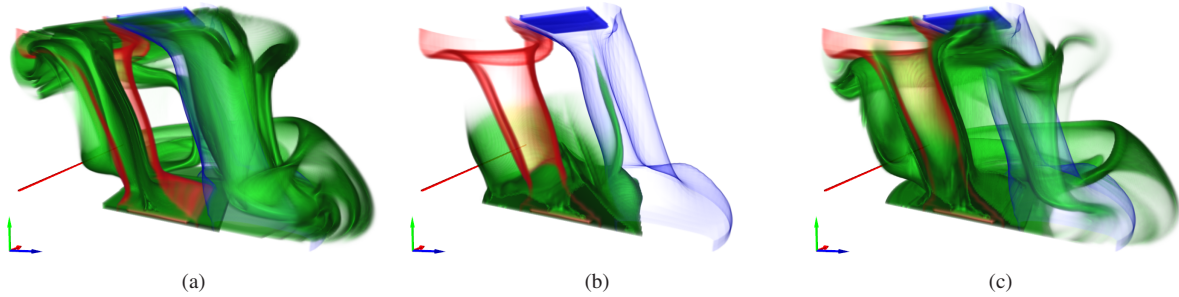
As a second example, we visualized advection-diffusion of heat in the heating coil data set. The coil is located between two pipes: an inner and an outer one, both cooled. An airflow passes the coil from bottom to top and is heated. Our technique reveals different transport behavior with respect to passive advection (Eq. 5) and passive advection-diffusion (Eq. 6), see Figure 10. One can easily see how heat is transported by advection-diffusion to the cooled walls.

As a third example, we applied our technique for the vi-


**Table 3:** Buoyant flow example at different configurations (no lighting and no isosurfaces). 1) Single dye at resolution  $122 \times 62 \times 122$  with early rejection (worst case in brackets), and 2) without. 3) Two independent dyes, and 4) one dye at  $244 \times 124 \times 244$ . See also the accompanying video.

Config.	Avg. FPS	Render [ms]	Dye comp. [ms]
Conf.1)	47.1 (22.6)	7.7	14.3 (35.0)
Conf.2)	22.7	9.0	35.1
Conf.3)	25.6 (12.3)	9.6	36.3 (73.4)
Conf.4)	10.3 (3.2)	7.8	97.6 (302.0)





**Figure 9:** Buoyant unsteady flow example (with clipping) by green dye seeded at the hot (red) plate. Pure advection (a), only diffusion flux (b) (passive diffusion), and both advective and diffusion flux (c) (passive advection-diffusion). In (b), diffusion flux transports the dye outward hot air (red) and toward cold air (blue). In (c), the dye reveals the true transport of heat: it is diffusing toward the cold plume and partially leaving at the cold (blue) plate before it is advected downward by the cold plume.




**Figure 10:** Heating coil example. Air flow from bottom to top, dye (green) seeded at two points at lower side of coil, and isosurfaces of temperature (red). Dye advection by advection only (a) vs. dye advection by advection-diffusion of heat (b). It is apparent that heat is repelling the dye from the coil and transporting it to the cooled walls in (b).

visualization of advection-diffusion in the context of an evaporating drop, i.e., we visualize advection-diffusion of vapor. The simulation was conducted using a finite volume-based direct numerical simulation employing the volume-of-fluid method for tracking of different phases [SW08]. The main air flow direction is from the left to the right. Figure 11 shows advection-diffusion visualized by our technique. In this case advection dominates diffusion—the passive advection-diffusion visualization (Eq. 6) is indistinguishable from that of advection only. Nevertheless, passive diffusion only (Eq. 6 with  $\mathbf{u} = \mathbf{0}$ ) reveals that diffusion is strongest at the upstream front of the drop.

## 6. Conclusion

We have presented an interactive visualization technique for advection-diffusion processes based on the finite vol-



**Figure 11:** Evaporating drop example with isosurfaces at vapor concentration 0.0001 and 0.005 (blue), drop located to the left. Dye (green) is seeded once (a) and continuously (b) at the drop. Transport outward the drop is revealed by passive diffusion only (a). However, due to the high drop speed, advection dominates the transport behavior (b).

ume method with WENO reconstruction. We further have introduced the concept of *passive diffusion* to visualize diffusion fluxes individually or in combination with advection. Our technique contributes to three scenarios: (a) high-quality interactive dye advection at low numerical diffusion, (b) physically-based dye advection accounting for diffusivity of the virtual media, and (c) visualization of advection-diffusion in simulation data. As demonstrated, our approach is able to reveal the transport mechanisms due to advection-diffusion that are responsible for concentration changes. A possible direction for future work is the extension of our approach to physical processes beyond advection-diffusion.

## Acknowledgments

The authors wish to thank Jan Schlottke from the ITLR Stuttgart for the drop data set and many fruitful discussions, and Mark Haas from the IAG Stuttgart for support with WENO code development and many useful comments. This work was supported in part by the Cluster of Excellence in Simulation Technology (EXC 310/1) and the Collaborative Research Centre SFB-TRR 75 at the University of Stuttgart.

## References

- [AB06] ACAR R., BOULANGER P.: Digital marbling: A multi-scale fluid model. *IEEE Trans. on Vis. and Comp. Graph.* 12, 4 (2006), 600–614. 2
- [BPR01] BÜRKLE D., PREUSSER T., RUMPF M.: Transport and anisotropic diffusion in time-dependent flow visualization. In *IEEE Visualization* (2001), pp. 61–67. 2
- [BSL60] BIRD B. R., STEWART W. E., LIGHTFOOT E. N.: *Transport Phenomena*. John Wiley & Sons, 1960. 3
- [CKSW08] CUNTZ N., KOLB A., STRZODKA R., WEISKOPF D.: Particle level set advection for the interactive visualization of unsteady 3D flow. *Comp. Graph. Forum* 27, 3 (2008), 719–726. 2
- [CL93] CABRAL B., LEEDOM L. C.: Imaging vector fields using line integral convolution. In *SIGGRAPH Conference Proceedings* (1993), pp. 263–270. 2
- [CS07] CHOU C.-S., SHU C.-W.: High order residual distribution conservative finite difference weno schemes for convection-diffusion steady state problems on non-smooth meshes. *Journal of Computational Physics* 224 (2007), 992–1020. 2, 6
- [CW84] COLELLA P., WOODWARD P. R.: The piecewise parabolic method (ppm) for gas-dynamical simulations. *Journal of Computational Physics* 54, 1 (1984), 174–201. 2
- [DK07] DUMBSER M., KÄSER M.: Arbitrary high order non-oscillatory finite volume schemes on unstructured meshes for linear hyperbolic systems. *Journal of Computational Physics* 221, 2 (2007), 693–723. 2, 4, 5, 6
- [DPR00] DIEWALD U., PREUSSER T., RUMPF M.: Anisotropic diffusion in vector field visualization on euclidean domains and surfaces. *IEEE Trans. on Vis. and Comp. Graph.* 6, 2 (2000), 139–149. 2
- [Fic55] FICK A.: *Phil. Mag.* 10 (1855), 30. 3
- [GKT\*08] GARTH C., KRISHNAN H., TRICOCHÉ X., BOBACH T., JOY K. I.: Generation of accurate integral surfaces in time-dependent vector fields. *IEEE Trans. on Vis. and Comp. Graph.* 14, 6 (2008), 1404–1411. 1
- [Hul92] HULTQUIST J. P.: Constructing stream surfaces in steady 3-d vector fields. In *IEEE Visualization* (1992), pp. 171–178. 1
- [JEH00] JOBARD B., ERLEBACHER G., HUSSAINI M. Y.: Hardware-accelerated texture advection for unsteady flow visualization. In *IEEE Visualization* (2000), pp. 155–162. 2
- [JS96] JIANG G.-S., SHU C.-W.: Efficient implementation of weighted eno schemes. *Journal of Computational Physics* 126, 1 (1996), 202–228. 5
- [Lev88] LEVOY M.: Display of surfaces from volume data. *IEEE Computer Graphics and Applications* 8, 3 (1988), 29–37. 7
- [LHD\*04] LARAMEE R. S., HAUSER H., DOLEISCH H., VROLIJK B., POST F. H., WEISKOPF D.: The state of the art in flow visualization: Dense and texture-based techniques. *Comp. Graph. Forum* 23, 2 (2004), 203–221. 2
- [LJH03] LARAMEE R. S., JOBARD B., HAUSER H.: Image space based visualization of unsteady flow on surfaces. In *IEEE Visualization* (2003), pp. 131–138. 2
- [LOC94] LIU X.-D., OSHER S., CHAN T.: Weighted essentially non-oscillatory schemes. *Journal of Computational Physics* 115, 1 (1994), 200–212. 2
- [LTH08] LI G.-S., TRICOCHÉ X., HANSEN C. D.: Physically-based dye advection for flow visualization. *Comp. Graph. Forum* (2008), 727–734. 2, 7
- [LTWH08] LI G.-S., TRICOCHÉ X., WEISKOPF D., HANSEN C.: Flow charts: Visualization of vector fields on arbitrary surfaces. *IEEE Trans. on Vis. and Comp. Graph.* 14, 5 (2008), 1067–1080. 2
- [SJK04] SANDERSON A. R., JOHNSON C. R., KIRBY R. M.: Display of vector fields using a reaction-diffusion model. In *IEEE Visualization* (2004), pp. 115–122. 2
- [SJM96] SHEN H.-W., JOHNSON C. R., MA K.-L.: Visualizing vector fields using line integral convolution and dye advection. In *1996 Volume Visualization Symposium* (1996), pp. 63–70. 2
- [SW08] SCHLOTTKE J., WEIGAND B.: Direct numerical simulation of evaporating droplets. *Journal of Computational Physics* 227, 10 (May 2008), 5215–5237. 9
- [TA03] TAPONETTO F., ALEXA M.: Vector field visualization using markov random field texture synthesis. In *Eurographics / IEEE TCVG Symposium on Visualization* (2003), pp. 195–202. 2
- [TB96] TURK G., BANKS D.: Image-guided streamline placement. In *SIGGRAPH* (1996), pp. 453–460. 1
- [Tur91] TURK G.: Generating textures on arbitrary surfaces using reaction-diffusion. *Computer Graphics (SIGGRAPH 91)* 25, 4 (1991), 289–298. 2
- [TvW03] TELEA A., VAN WIJK J. J.: 3D IBFV: Hardware-accelerated 3D flow visualization. In *IEEE Visualization* (2003), pp. 233–240. 2
- [vW93] VAN WIJK J. J.: Implicit stream surfaces. In *IEEE Visualization* (1993), pp. 245–252. 1
- [vW02] VAN WIJK J. J.: Image based flow visualization. *ACM Transactions on Graphics* 21, 3 (2002), 745–754. 2
- [vW03] VAN WIJK J. J.: Image based flow visualization for curved surfaces. In *IEEE Visualization* (2003), pp. 123–130. 2
- [VZKL06] VILANOVA A., ZHANG S., KINDLMANN G., LAIDLAW D.: An introduction to visualization of diffusion tensor imaging and its applications. In *Visualization and Processing of Tensor Fields* (2006), Springer, pp. 121–153. 2
- [Wei04] WEISKOPF D.: Dye advection without the blur: A level-set approach for texture-based visualization of unsteady flow. *Comp. Graph. Forum* 23, 3 (2004), 479–488. 2
- [Wei09] WEISKOPF D.: Iterative twofold line integral convolution for texture-based vector field visualization. In *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*, Möller T., Hamann B., Russell R., (Eds.). Springer, 2009, pp. 191–211. 2
- [WHE01] WEISKOPF D., HOPF M., ERTL T.: Hardware-accelerated visualization of time-varying 2D and 3D vector fields by texture advection via programmable per-pixel operations. In *Vision, Modeling, and Visualization (VMV)* (2001), pp. 439–446. 2
- [WJE00] WESTERMANN R., JOHNSON C., ERTL T.: A level-set method for flow visualization. In *IEEE Visualization* (2000), pp. 147–154. 1
- [WK91] WITKIN A., KASS M.: Reaction-diffusion textures. *Computer Graphics (SIGGRAPH 91)* 25, 4 (1991), 299–308. 2
- [WSE07] WEISKOPF D., SCHAFHITZEL T., ERTL T.: Texture-based visualization of unsteady 3D flow by real-time advection and volumetric illumination. *IEEE Trans. on Vis. and Comp. Graph.* 13, 3 (2007), 569–582. 2
- [XZC04] XUE D., ZHANG C., CRAWFIS R.: Rendering implicit flow volumes. In *IEEE Visualization* (2004), pp. 99–106. 1