

Velocity-Based Monte Carlo Fluids

Ryusuke Sugimoto
University of Waterloo
Waterloo, Ontario, Canada
rsugimot@uwaterloo.ca

Christopher Batty
University of Waterloo
Waterloo, Ontario, Canada
christopher.batty@uwaterloo.ca

Toshiya Hachisuka
University of Waterloo
Waterloo, Ontario, Canada
toshiya.hachisuka@uwaterloo.ca




Figure 1: Our velocity-based Monte Carlo fluid solver can naturally handle scenes for which the existing vorticity-based Monte Carlo method [Rioux-Lavoie and Sugimoto et al. 2022] yields incorrect results. Our solver allows the red and blue smoke densities associated with a pair of vortices to flow between two obstacles (left, top), similar to the conventional non-Monte Carlo velocity-based method [Batty et al. 2007] (left, middle), while the vorticity-based method produces an incorrect result, in which smoke deviates around the outside of the obstacles (left, bottom). Our solver also readily supports commonly available features in traditional velocity-based Eulerian grid fluid solvers, such as buoyancy effects: we simulate a smoke plume rising from a complex-shaped source (right, top) and past a sphere-shaped or bunny-shaped obstacle (right, bottom), where the motion is driven solely by buoyancy forces.

ABSTRACT

We present a velocity-based Monte Carlo fluid solver that overcomes the limitations of its existing vorticity-based counterpart. Because the velocity-based formulation is more commonly used in graphics, our Monte Carlo solver can be readily extended with various techniques from the fluid simulation literature. We derive our method by solving the Navier-Stokes equations via operator splitting and designing a pointwise Monte Carlo estimator for each sub-step. We reformulate the projection and diffusion steps as integration problems based on the recently introduced walk-on-boundary technique [Sugimoto et al. 2023]. We transform the volume integral arising from the source term of the pressure Poisson equation into a form more amenable to practical numerical evaluation. Our resulting velocity-based formulation allows for the proper simulation of scenes that the prior vorticity-based Monte Carlo method [Rioux-Lavoie and Sugimoto et al. 2022] either simulates incorrectly or cannot support. We demonstrate that our method can easily incorporate advancements drawn from conventional non-Monte Carlo methods by showing how one can straightforwardly add buoyancy effects, divergence control capabilities, and numerical dissipation reduction methods, such as advection-reflection and PIC/FLIP methods.

CCS CONCEPTS

• **Computing methodologies** → **Physical simulation**; Ray tracing; • **Mathematics of computing** → *Partial differential equations*; *Integral equations*.

KEYWORDS

Monte Carlo methods, fluid simulation, walk-on-boundary

1 INTRODUCTION

Researchers in computer graphics have recently revisited Monte Carlo PDE solvers [Muller 1956; Sabelfeld 1982] due to their attractive properties. Sawhney and Crane [2020] showed that geometry processing tasks can reap benefits from Monte Carlo solvers, such as flexibility with respect to geometric boundary representations, robustness to noise in the input geometry, support for pointwise estimation of the solution, and trivial parallelization. Monte Carlo methods have already been used in light transport simulation for decades since the work of Kajiya [1986] due to these same properties.

Looking beyond light transport and geometry processing, Rioux-Lavoie and Sugimoto et al. [2022] developed the first Monte Carlo

method to perform smoke simulation based on the Navier-Stokes equations. They discussed theoretical and practical aspects of Monte Carlo-based fluid simulation; most advantages in the geometry processing context carry over directly to fluid simulation applications. However, their method relies on a vorticity-based formulation, whereas velocity-based formulations are more common in fluid animation.

Since the introduction of velocity-based Eulerian grid fluid simulation techniques to graphics [Foster and Metaxas 1996; Stam 1999], researchers have improved on many aspects of the velocity-based formulation, adding features such as buoyancy modeling for smoke [Fedkiw et al. 2001; Foster and Metaxas 1997] and velocity divergence control for expansion/contraction and artistic effects [Feldman et al. 2003], as well as PIC/FLIP solvers [Brackbill and Ruppel 1986; Harlow 1962; Zhu and Bridson 2005] and advection-reflection solvers [Narain et al. 2019; Zehnder et al. 2018] for reducing numerical dissipation. These and many other improvements cannot be straightforwardly incorporated into the Monte Carlo fluid solver of Rioux-Lavoie and Sugimoto et al. [2022] because of its vorticity-based formulation. Furthermore, Yin et al. [2023] recently showed that existing vorticity-based methods can fail to simulate harmonic velocity fields, leading to incorrect results when the fluid domain is not simply connected (e.g., is disjoint or has holes). Conveniently, however, velocity-based methods can capture these physics without any change.

We therefore introduce a Monte Carlo fluid solver that relies on a velocity-based formulation. Similarly to the vorticity-based method [Rioux-Lavoie and Sugimoto et al. 2022], our method does not require the boundaries of solid objects to be explicitly discretized; for example, there is no need for a cut-cell method [Batty et al. 2007] or a conforming mesh [Feldman et al. 2005] as in traditional solvers, and the only requirement on the geometry is that it support ray intersection queries. We use the recently introduced walk-on-boundary method [Sugimoto et al. 2023], which is a Monte Carlo ray tracing solver for PDEs, to perform pressure projection in our solver, and thus, our solver can be trivially accelerated by GPU ray tracing. Our method can additionally incorporate all of the aforementioned techniques developed for velocity-based simulations: buoyancy, divergence control, PIC/FLIP, and advection-reflection.

To solve the Navier-Stokes equations, we apply operator splitting [Stam 1999], and the core of our solver depends on reformulating the projection and diffusion steps as integration problems. In the absence of obstacles, we derive a non-recursive integral for each step and apply Monte Carlo integration. The projection step requires a careful evaluation of the volume integral term arising from the Poisson equation’s source term, and we describe in detail how to transform it into a form that is more amenable to numerical evaluation. Domains with boundaries require solving a boundary integral equation, which we treat using the walk-on-boundary method [Sugimoto et al. 2023]. The diffusion equation for the diffusion step is a time-dependent equation, and we apply a *time-dependent* walk-on-boundary method [Sabelfeld and Simonov 1994] for problems with boundaries; we are the first to apply it to a computer graphics application. After constructing Monte Carlo solvers for each substep, we combine them to design a Navier-Stokes solver with flexible options for how we store the intermediate velocity field in a discretized cache. To summarize, our contributions include:

- a velocity-based Monte Carlo fluid simulator that uses operator splitting,
- an integral formulation for the projection step with careful Poisson source term handling,
- extension of the projection step to problems with boundaries using the walk-on-boundary method,
- application of the diffusion walk-on-boundary method for the diffusion step,
- adaptations of more advanced velocity-based solver techniques to our Monte Carlo formulation.

2 METHOD

Our method aims to numerically solve for a velocity field that satisfies the incompressible Navier-Stokes equations with constant density and viscosity:

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} &= -(\mathbf{u} \cdot \nabla) \mathbf{u} - \frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f}, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned} \quad (1)$$

where $\mathbf{u}$ is velocity, p is pressure, t is time, $\mathbf{f}$ is acceleration due to external forces, ν is kinematic viscosity, and ρ is density.

Vorticity-Based Monte Carlo. Let us recap the method by Rioux-Lavoie and Sugimoto et al. [2022] to highlight the challenges in developing a Monte Carlo solver for the velocity-based formulation in Eq. 1. Their method uses the vorticity representation of the fluid. For instance, for inviscid fluid in 2D, the vorticity-based formulation simplifies the Navier-Stokes equations to the vorticity transport equation, $\frac{D\omega}{Dt} = 0$, where $\frac{D}{Dt}$ is the material derivative. We can write the velocity field used to advect the vorticity field as an integral, and combined with a semi-Lagrangian advection gives the equation they used to advance the vorticity field in each time step in the absence of obstacles:

$$\omega(\mathbf{x}, t) \approx \omega \left(\mathbf{x} - \Delta t \int_{\mathbb{R}^2} B(\mathbf{x}, \mathbf{y}) \times \omega(\mathbf{y}, t - \Delta t) dV(\mathbf{y}, t - \Delta t) \right),$$

where $B(\mathbf{x}, \mathbf{y})$ is the Biot-Savart kernel. They applied Monte Carlo integration to evaluate the integral.

The vorticity-based formulation offered a straightforward application of the Monte Carlo method. It is nontrivial to design a similar scheme for the velocity-based formulation in Eq. 1 because the Navier-Stokes equations, in their velocity form, couple the velocity and the pressure variables in an intricate way: the pressure serves as a Lagrange multiplier that makes the velocity field incompressible. We address this challenge by adopting operator splitting.

Operator splitting with pointwise estimators. The basic formulation of our method follows the standard operator splitting framework [Stam 1999]. We solve Eq. 1 by taking discrete time steps with step size Δt . For each time step, we decompose the Navier-Stokes equations into four substeps:

- (1) advection: $\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla) \mathbf{u}$,
- (2) external force integration: $\frac{\partial \mathbf{u}}{\partial t} = \mathbf{f}$,
- (3) diffusion: $\frac{\partial \mathbf{u}}{\partial t} = \nu \nabla^2 \mathbf{u}$, and
- (4) projection: $\frac{\partial \mathbf{u}}{\partial t} = -\frac{1}{\rho} \nabla p$ such that $\nabla \cdot \mathbf{u} = 0$.

In contrast to traditional discretization-based solvers, we develop a pointwise (Monte Carlo) estimator for each substep. We design a

solver for each substep that can estimate the velocity at *any* spatial point we are interested in, assuming we can likewise query the velocity estimates from the previous steps at any spatial point.

Basing our approach on this pointwise formulation offers unique advantages. Our formulation is agnostic to the underlying discretization of the velocity field, which allows for flexible integration of velocity-based techniques used by traditional discretization-based solvers, including the common Stable Fluids-style [Stam 1999] grid-based methods and grid-particle hybrid PIC/FLIP methods, into our Monte Carlo formulation. While the vorticity-based Monte Carlo method [Rioux-Lavoie and Sugimoto et al. 2022] has similar advantages of having flexible discretization options, their method cannot benefit from common velocity-based techniques. Moreover, we will show in Section 3 that our formulation allows us to remove some of the excessive grid interpolation errors that are otherwise introduced in traditional solvers.

Below, we formulate a pointwise estimate of the solution to each substep. Later, in Section 3, we will describe how we employ these estimators in a few different implementations, some using uniform grid velocity storage and others using hybrid grid-particle velocity storage. To simplify notation, we use $\mathbf{u}_0$ to $\mathbf{u}_4$ to indicate the velocity field at each substep within a single time step. The initial velocity field for the current time step $\mathbf{u}_0$ is defined as the output velocity field $\mathbf{u}_4$ from the last time step. The advection step takes $\mathbf{u}_0$ as its input and outputs $\mathbf{u}_1$, and so on.

2.1 Advection

We employ semi-Lagrangian advection, which estimates the new velocity at a point by tracing trajectories backward in time through the velocity field. For example, a forward Euler time discretization for the velocity at point $\mathbf{x}$ yields the update rule

$$\mathbf{u}_1(\mathbf{x}) \leftarrow \mathbf{u}_0(\mathbf{x} - \Delta t \mathbf{u}_0(\mathbf{x})). \quad (2)$$

This semi-Lagrangian update satisfies our pointwise evaluation requirement: we can estimate the velocity after the step by evaluating the input velocities at a few points. While the time-discretized equation above is identical to what appears in grid-based fluid solvers, we emphasize that the point where we evaluate the advected velocity, $\mathbf{x}$, does not necessarily need to be aligned with discretized locations (e.g., grid points), and when we evaluate the pre-advection velocity at points $\mathbf{x}$ and $\mathbf{x} - \Delta t \mathbf{u}_0(\mathbf{x})$, we are not restricted to evaluation by interpolation of some grid data. Thus, the pointwise form offers flexibility that is unavailable with classical solvers. In practice, rather than forward Euler, we use a higher-order time integration scheme (third-order Runge-Kutta) to improve the accuracy of the traced trajectories, and the adaptations of error-correcting schemes such as MacCormack [Selle et al. 2008] and BFECC [Dupont and Liu 2007] are also trivial; all of these schemes likewise support pointwise evaluation.

2.2 External Force Integration

In the presence of external forces leading to acceleration, such as a buoyancy force, we update the post-advection velocity using a forward Euler discretization:

$$\mathbf{u}_2(\mathbf{x}) \leftarrow \mathbf{u}_1(\mathbf{x}) + \Delta t \mathbf{f}(\mathbf{x}). \quad (3)$$

This expression is again a pointwise update of the velocity field, where $\mathbf{x}$ is not necessarily at a discretized location.

2.3 Projection

As the diffusion step is necessary only for viscous fluids and we can simply let $\mathbf{u}_3 \leftarrow \mathbf{u}_2$ for inviscid fluids, we first discuss the projection step. The objective of the projection step is to find the pressure field that projects out the divergent velocity mode of the input field. The post-projection velocity at $\mathbf{x}$ is given by

$$\mathbf{u}_4(\mathbf{x}) \leftarrow \mathbf{u}_3(\mathbf{x}) - \nabla \bar{p}(\mathbf{x}), \quad (4)$$

where $\bar{p}$ satisfies the Poisson equation

$$\nabla^2 \bar{p}(\mathbf{x}) = \nabla \cdot \mathbf{u}_3(\mathbf{x}), \quad (5)$$

and $\bar{p} = \frac{\Delta t}{\rho} p$. Since only $\nabla \bar{p}$ is required for the update in Eq. 4, the task is reduced to estimating the gradient of the solution of the Poisson equation; the value of $\bar{p}$ itself is not needed. In each subsection below, we discuss the integral formulation followed by its evaluation with Monte Carlo integration.

2.3.1 Without solid boundaries. To simplify the problem, we will first consider a case without any solid boundaries (i.e., the domain is unbounded with no obstacles inside it). Under this assumption, it is well-known that we can write the solution to the Poisson equation using the fundamental solution G of the Laplace operator. The fundamental solution G satisfies $\nabla^2 G(\mathbf{x}, \mathbf{y}) + \delta(\mathbf{x} - \mathbf{y}) = 0$ in the unbounded domain, where δ is the Dirac delta function. In 2D, $G(\mathbf{x}, \mathbf{y}) = -\frac{1}{2\pi} \log r$, and in 3D, $G(\mathbf{x}, \mathbf{y}) = \frac{1}{4\pi r}$, where $r = \|\mathbf{y} - \mathbf{x}\|_2$. We can then write the solution to Eq. 5 as a convolution of the source term of the Poisson equation and the fundamental solution:

$$\bar{p}(\mathbf{x}) = - \int_{\mathbb{R}^d} G(\mathbf{x}, \mathbf{y}) \nabla_{\mathbf{y}} \cdot \mathbf{u}_3(\mathbf{y}) dV(\mathbf{y}), \quad (6)$$

where dimension $d = 2, 3$. We used the subscript $\mathbf{y}$ to indicate that we perform differentiation with respect to $\mathbf{y}$, and we will use similar notation throughout this section. By applying the Laplacian to Eq. 6, one can confirm that $\bar{p}$ satisfies Eq. 5. To evaluate the gradient of $\bar{p}$, we take the gradient of Eq. 6:

$$\nabla_{\mathbf{x}} \bar{p}(\mathbf{x}) = - \int_{\mathbb{R}^d} \nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{y}) \nabla_{\mathbf{y}} \cdot \mathbf{u}_3(\mathbf{y}) dV(\mathbf{y}). \quad (7)$$

This integral is still not suitable for numerical computation for our purposes, as we describe below. We therefore extend beyond the prior work by proposing a further transformation.

Since we typically consider storing information and performing volume integrals over a bounded simulation domain, we replace the (infinite) integral domain $\mathbb{R}^d$ in Eq. 7 with a bounded simulation domain Ω^s by assuming that the velocity divergence $\nabla \cdot \mathbf{u}_3$ is zero outside the simulation domain.

Attempting to use Eq. 7 in this form for Monte Carlo integration still requires explicit evaluation of velocity divergence inside the domain, whereas we desire a solver that takes only a pointwise-evaluated velocity field as input. There are a few possible approaches to obtain the required divergence. First, we could apply finite differences to the velocity field by accepting some errors. Second, we could differentiate the substep that precedes projection so that it outputs the necessary velocity divergence, in addition to

the velocity itself. We propose instead a velocity-only design that we believe fits better in our Monte Carlo framework.

To eliminate the dependency on velocity divergence in Eq. 7, we use the identity $\nabla_{\mathbf{x}}G = -\nabla_{\mathbf{y}}G$ and apply integration by parts:

$$\nabla_{\mathbf{x}}\bar{p}(\mathbf{x}) = \int_{\Omega^s} \{\nabla_{\mathbf{y}}G(\mathbf{x}, \mathbf{y})\} \nabla_{\mathbf{y}} \cdot \mathbf{u}_3(\mathbf{y}) dV(\mathbf{y}) \quad (8)$$

$$= - \int_{\Omega^s} \mathbf{H}(\mathbf{x}, \mathbf{y}) \mathbf{u}_3(\mathbf{y}) dV(\mathbf{y}) \quad (9)$$

$$- \int_{\partial\Omega^s} \{\nabla_{\mathbf{x}}G(\mathbf{x}, \mathbf{y})\} \mathbf{n}(\mathbf{y}) \cdot \mathbf{u}_3(\mathbf{y}) dA(\mathbf{y}),$$

where $\mathbf{n}$ is the outward unit normal. The Hessian of the fundamental solution denoted $\mathbf{H}$, is given by

$$\mathbf{H}(\mathbf{x}, \mathbf{y}) = \mathbf{S}(\mathbf{x}, \mathbf{y}) - \frac{\delta(\mathbf{r})}{d} \mathbf{I}, \quad \mathbf{S}(\mathbf{x}, \mathbf{y}) = \frac{1}{|\partial B| r^{d+2}} (d\mathbf{r}\mathbf{r}^\top - r^2 \mathbf{I}), \quad (10)$$

where $|\partial B|$ is the surface area of a unit $(d-1)$ -sphere, $\mathbf{r} = \mathbf{y} - \mathbf{x}$, $r = \|\mathbf{r}\|_2$ and $\mathbf{I}$ is the identity matrix. While we use the notation above for ease of understanding, technically, it needs to be understood in the sense of the generalized (distributional) derivative, and interested readers can refer to the discussion by Frahm [1983] and Hnizdo [2011] for further details. Substituting Eq. 10 into Eq. 9, we get

$$\begin{aligned} \nabla_{\mathbf{x}}\bar{p}(\mathbf{x}) &= - \int_{\Omega^s} \mathbf{S}(\mathbf{x}, \mathbf{y}) \mathbf{u}_3(\mathbf{y}) dV(\mathbf{y}) + \frac{1}{d} \mathbf{u}_3(\mathbf{x}) \\ &\quad - \int_{\partial\Omega^s} \{\nabla_{\mathbf{x}}G(\mathbf{x}, \mathbf{y})\} \mathbf{n}(\mathbf{y}) \cdot \mathbf{u}_3(\mathbf{y}) dA(\mathbf{y}) \\ &= - \int_{\Omega^s} \mathbf{S}(\mathbf{x}, \mathbf{y}) \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dV(\mathbf{y}) \\ &\quad - \int_{\partial\Omega^s} \{\nabla_{\mathbf{x}}G(\mathbf{x}, \mathbf{y})\} \mathbf{n}(\mathbf{y}) \cdot \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dA(\mathbf{y}). \end{aligned} \quad (11)$$

Again, technically, the domain of the first integral should exclude an infinitesimal ball around $\mathbf{x}$ [Hnizdo 2011], but we use this simple notation for readability. In Eq. 11, to get the final expression, we replaced the original velocity field with a velocity field that is globally shifted by the constant velocity at point $\mathbf{x}$, $\mathbf{u}_3(\mathbf{x})$. The computed pressure gradient remains unchanged because the divergence of the shifted velocity field is the same as the original one. This global shift cancels the zeroth order term of the Taylor expansion of $\mathbf{u}_3(\mathbf{y})$ around $\mathbf{x}$, and the remaining terms of the Taylor expansion multiplied by the function $\mathbf{S}$ will have a singularity of $O(1/r^{d-1})$ instead of the original $O(1/r^d)$ as $r \rightarrow 0$. This lower-order singularity can be handled by an appropriate importance-sampling strategy, as we discuss next.

Monte Carlo Estimation. Now that we have an integral representation of the pressure gradient, we can define a Monte Carlo estimator for the pressure gradient. We sample N_V points $\mathbf{y}^i$ inside the simulation domain Ω^s and N_A points $\mathbf{y}^j$ on the simulation domain boundary $\partial\Omega^s$ using sampling strategies with probability density functions (PDFs) P_V and P_A , respectively, to get

$$\nabla_{\mathbf{x}}\bar{p}(\mathbf{x}) \approx \langle E_V(\mathbf{x}) \rangle + \langle E_A(\mathbf{x}) \rangle \quad (12)$$

where

$$\langle E_V(\mathbf{x}) \rangle = -\frac{1}{N_V} \sum_{i=1}^{N_V} \frac{\mathbf{S}(\mathbf{x}, \mathbf{y}^i)}{P_V(\mathbf{y}^i|\mathbf{x})} \{\mathbf{u}_3(\mathbf{y}^i) - \mathbf{u}_3(\mathbf{x})\} \quad (13)$$

$$\langle E_A(\mathbf{x}) \rangle = -\frac{1}{N_A} \sum_{j=1}^{N_A} \frac{\nabla_{\mathbf{x}}G(\mathbf{x}, \mathbf{y}^j)}{P_A(\mathbf{y}^j|\mathbf{x})} \mathbf{n}(\mathbf{y}^j) \cdot \{\mathbf{u}_3(\mathbf{y}^j) - \mathbf{u}_3(\mathbf{x})\}. \quad (14)$$

This formulation is an unbiased estimator for the pressure gradient if the probability that we sample any point with a nonzero contribution is nonzero. If the integrand is non-negative everywhere, importance sampling according to the PDF that is roughly proportional to the integrand can decrease the variance of the estimator. While our integrands may contain negative values, we follow this idea to design our sampling strategy. In our implementation, for Eq. 13, to handle the singular integral, we draw samples such that PDF P_V is proportional to $1/r^{(d-1)}$ inside of the smallest ball around $\mathbf{x}$ that fully contains the entire simulation domain, and set zero contributions from the samples outside the simulation domain. This approach is equivalent to evaluating the integral we get by extending the original integral domain to the ball and putting zero integrands in the extended part. We also use antithetic sampling here: in addition to a point $\mathbf{y}^i$ sampled this way, we always sample the symmetric point with respect to $\mathbf{x}$, $2\mathbf{x} - \mathbf{y}^i$ to further reduce variance for smooth velocity fields. For Eq. 14, we uniformly sample points on the simulation boundary. Properly sampling the positive and negative contributions separately may further reduce the variance of the estimators [Chang et al. 2023; Owen 2013, Section 9.12], but we leave it as future work. Equation 12 only requires the ability to evaluate the velocity at the position of sample points to perform projection, in contrast to traditional solvers that typically require a globally coupled linear system solve.

2.3.2 With solid boundaries. Next, we consider the case when solid object boundaries are involved (Fig. 2). We still solve Eq. 5 for $\nabla\bar{p}$, but with free slip boundary conditions,

$$\frac{\partial\bar{p}}{\partial\mathbf{n}} = \mathbf{n} \cdot (\mathbf{u}_3 - \mathbf{u}_s) \quad (15)$$

on solid boundaries, where $\mathbf{n}$ is the unit outward normal from the fluid domain Ω , $\frac{\partial}{\partial\mathbf{n}} = \mathbf{n} \cdot \nabla$ is a normal derivative, and $\mathbf{u}_s$ is the solid velocity. If the simulation domain is bounded, this problem is an interior Neumann problem for the Poisson equation, and otherwise, it is an exterior Neumann problem; we need a Monte Carlo solver that is capable of handling these problems.

The basic walk-on-spheres method [Sawhney and Crane 2020] can only handle Dirichlet problems, and its applicability to unbounded domain problems leaves some concerns about the termination of paths without additional bias. Rioux-Lavoie and Sugimoto et al. [2022] arbitrarily terminated their walk-on-spheres paths after a few steps in their simulation, leaving some additional bias. Nabizadeh et al. [2021] suggested inverting the unbounded domain into a bounded one using the Kelvin transform so that paths always terminate, but using this approach would require very different treatment for bounded and unbounded domains.

To solve Neumann problems, the walk-on-stars method [Sawhney and Miller et al. 2023] extended the walk-on-spheres method

with carefully-designed recursion relationships and tailored sampling techniques based on the extensions by Simonov [2008] and Ermakov and Sipin [2009]. For a Neumann problem in an unbounded domain, however, the Kelvin transform [Nabizadeh et al. 2021] is still required; it transforms a Neumann problem into a Robin problem, to which the application of walk-on-stars has not yet been attempted. Among others, the walk-on-boundary method [Sabelfeld 1982; Sabelfeld and Simonov 1994], introduced recently to graphics by Sugimoto et al. [2023], can be applied to Neumann boundary problems in both bounded and unbounded domains under a unified framework. The method builds upon a boundary integral equation analogous to the rendering equation for light transport simulation [Kajiya 1986], and uses ray-tracing to solve the boundary integral equation, and we use this method in our solver.

Following Sabelfeld and Simonov [1994], we can write the gradient of the solution to the Poisson equation of Eq. 5 with the boundary condition of Eq. 15 using the single layer potential formulation and define an integral equation for the unknown density function. For our application, these equations contain the volume integral terms arising from the source term in the Poisson equation of Eq. 5 specific to our problem. Similar to Section 2.3.1, we transform these volume integral terms as follows. If the domain Ω is unbounded, we suggest replacing it with a bounded simulation domain Ω^s , assuming zero velocity divergence outside of the simulation domain; otherwise, we let $\Omega^s = \Omega$. With this definition, we have $\partial\Omega \subseteq \partial\Omega^s$ for the boundaries, and we distinguish these two sets of boundaries in the following equations. We also remove the explicit dependencies on the velocity divergence by integration by

parts similarly to Eq. 9 to get

$$\begin{aligned} \nabla_{\mathbf{x}} \bar{p}(\mathbf{x}) = & \int_{\partial\Omega} \{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{y})\} [\mu(\mathbf{y}) - \mathbf{n}(\mathbf{y}) \cdot \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\}] dA(\mathbf{y}) \\ & - \int_{\Omega^s} \mathbf{S}(\mathbf{x}, \mathbf{y}) \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dV(\mathbf{y}) \\ & - \int_{\partial\Omega^s \setminus \partial\Omega} \{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{y})\} \mathbf{n}(\mathbf{y}) \cdot \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dA(\mathbf{y}) \end{aligned} \quad (16)$$

for $\mathbf{x} \in \Omega$ and

$$\begin{aligned} \mu(\mathbf{x}) = & - \int_{\partial\Omega} 2 \frac{\partial G}{\partial \mathbf{n}_{\mathbf{x}}}(\mathbf{x}, \mathbf{y}) [\mu(\mathbf{y}) - \mathbf{n}(\mathbf{y}) \cdot \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\}] dA(\mathbf{y}) \\ & + \int_{\Omega^s} 2 \mathbf{n}(\mathbf{x})^\top \mathbf{S}(\mathbf{x}, \mathbf{y}) \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dV(\mathbf{y}) \\ & + \int_{\partial\Omega^s \setminus \partial\Omega} 2 \frac{\partial G}{\partial \mathbf{n}_{\mathbf{x}}}(\mathbf{x}, \mathbf{y}) \mathbf{n}(\mathbf{y}) \cdot \{\mathbf{u}_3(\mathbf{y}) - \mathbf{u}_3(\mathbf{x})\} dA(\mathbf{y}) \\ & + 2 \mathbf{n}(\mathbf{x}) \cdot \{\mathbf{u}_3(\mathbf{x}) - \mathbf{u}_s(\mathbf{x})\} \end{aligned} \quad (17)$$

for $\mathbf{x} \in \partial\Omega$. In both Eq. 16 and Eq. 17, we have the boundary integral terms that involve the unknown density function μ on the solid boundaries $\partial\Omega$, and all the other integrals involve only known quantities. Note Eq. 16 is a generalization of the strategy described in Section 2.3.1 as we can recover Eq. 11 by dropping the solid boundary integral term.

Monte Carlo Estimation. Based on Eq. 16 and Eq. 17, we get a biased walk-on-boundary Monte Carlo estimator [Sabelfeld and Simonov 1994; Sugimoto et al. 2023]. Compared to the simplest formulation for the Laplace equation [Sugimoto et al. 2023], we have some additional non-recursive terms in Eq. 16 and Eq. 17 as a result of the transformations discussed above; we must consider how to sample these added terms. We choose to sample the non-recursive contributions on $\partial\Omega$ using ray intersection sampling for importance sampling and estimate the other non-recursive terms similarly to Section 2.3.1. To improve efficiency, we additionally employ boundary value caching, which lets us share the walk-on-boundary subpaths among evaluation points akin to the virtual point light (VPL) method [Keller 1997] in rendering. We describe the details of our particular sampling strategy in the supplemental note. Our Monte Carlo method is compatible with a variety of sampling techniques, so further variance reduction is likely possible.

2.4 Diffusion

Simulating viscous fluids requires an additional diffusion step before the projection step. The set of equations we solve here is the constant-coefficient diffusion equation,

$$\begin{aligned} \frac{\partial \bar{\mathbf{u}}(\mathbf{x}, s)}{\partial s} &= \nu \nabla^2 \bar{\mathbf{u}}(\mathbf{x}, s) \quad \text{for } \mathbf{x} \in \Omega, s \in (0, \Delta t), \\ \bar{\mathbf{u}}(\mathbf{x}, s) &= \mathbf{u}_s(\mathbf{x}) \quad \text{for } \mathbf{x} \in \partial\Omega, s \in (0, \Delta t), \quad \text{and} \\ \bar{\mathbf{u}}(\mathbf{x}, 0) &= \mathbf{u}_2(\mathbf{x}) \quad \text{for } \mathbf{x} \in \Omega, \end{aligned} \quad (18)$$

and the output of this diffusion step is $\mathbf{u}_3(\mathbf{x}) = \bar{\mathbf{u}}(\mathbf{x}, \Delta t)$. The time s here represents the time within each diffusion time step. We use the overbar to indicate that the velocity variable $\bar{\mathbf{u}}$ takes time s as defined here, and differs from the other sections. The solid velocity gives the no-slip boundary condition, and the output from the preceding advection step gives the initial conditions.




Figure 2: Our method can correctly simulate scenes with different boundary types (resolution 256^3). From the left, we simulate the colored smoke advected with the velocity field induced by vortex pairs moving with no obstacles (left), in an unbounded domain with a moving square obstacle (middle), and in a domain bounded by an obstacle (right). For each simulation, we show the time evolution from the left to the right, but the time stamps differ for each setup.  1m17s

When there are no solid boundaries or the time step size is small enough to ignore boundary effects, we can omit the second equation describing the boundary condition. In such a case, evaluating a convolution of the input velocity field with a Gaussian function would give the exact solution to the diffusion equation. Though they apply it to the vorticity field rather than the velocity field, Rioux-Lavoie and Sugimoto et al. [2022] use such a convolution-based approach to model diffusion. This approach fails when either the time step size is relatively large or the boundary effect is important. Traditional grid-based solvers similarly cannot rely on a simple Gaussian filter in the presence of large time steps or obstacles, and must instead solve a globally coupled linear system for viscosity [Bridson 2015].

To address this problem, we utilize the walk-on-boundary method for the diffusion equation [Sabelfeld and Simonov 1994, Chapter 4] to solve Eq. 18. Similarly to Section 2.3.2, the diffusion walk-on-boundary method is a pointwise estimator and is a natural generalization of the case when there are no boundaries. One difference is that the boundary integral equation and the walk-on-boundary steps are now defined in the space-time domain, and we must sample paths from the point where we want to evaluate the solution with time Δt towards the initial time within the time step, in the negative time direction.

Notably, for the diffusion walk-on-boundary, the recursion depth does not need to be predefined. By sampling the space-time paths using a PDF proportional to the integral kernel of the integral equation for the diffusion equation, we can terminate recursions when the sampled time is negative. In our experiments, we observed that the diffusion step is cheaper than the projection step, and we did not attempt any algorithmic improvements, such as the reuse of subpaths. However, there remain many choices for sampling and various efficiency improvement strategies should apply here, too. The supplemental material summarizes the diffusion walk-on-boundary method for readers’ convenience.

3 RESULTS

We explain the results of our practical implementation of the method with discretization structures and other possible extensions. We implemented our method using CUDA for GPU parallelism and the NVIDIA OptiX ray tracing engine [Parker et al. 2010] via the OWL wrapper library [Wald 2023] to accelerate the ray intersection queries. For all 2D results, we use consistent parameters as listed in the supplemental note except for Fig. 7. The figures have time stamps for the corresponding animations in the supplemental video.

Caching of velocity fields. While our fundamental formulation is agnostic to a choice of intermediate storage of velocity fields, or absence of it by the expensive recursive application of estimators to the initial time step (see Section 4), our basic implementation uses a simple uniform grid structure to store the velocities at grid nodes at each time step, and we refer to such a strategy a caching strategy in contrast to the fully time-recursive (cacheless) alternative. Since we do not evaluate any finite differences on the stored data, we do not need to use a staggered grid [Harlow and Welch 1965]. When we query a velocity from the cache, we bilinearly or trilinearly interpolate the values, and for points outside the cached domain, we take the velocity at the nearest cache point. As with classical grid-based solvers, interpolation- and advection-induced errors

sometimes cause us to query values from points inside of solid obstacles; for this issue, we assign velocities inside solid obstacles to take velocities from the solid itself, giving no-slip behavior on the boundary. Alternatively, one could fill the velocities inside solid obstacles with some form of extrapolation to approximate free slip behavior, for example, but we use the first approach throughout this paper.

While we could cache the velocity field after each substep, the pointwise estimation capability of our approach lets us avoid some of the excessive velocity caching between substeps and reduce the associated errors. For example, we need no caching between the advection and the following projection step for inviscid simulation if we query the advected velocity values directly on the fly whenever the projection estimator needs such a pointwise input velocity. Alternatively, we can also design a method without velocity caching between the projection and the next advection so the velocity field is always advected according to the pointwise divergence-free velocity field. We implemented these two options in addition to the option where we cache the resulting velocity field both after advection and after projection in Fig. 3 (b) to (d). Our experiments showed consistent results among the three options; further investigations are needed to understand when having fewer caching errors can make critical differences.

Boundary conditions. Fig. 2 shows that our method can correctly simulate the velocity field due to vortex pairs under different boundary conditions. In Fig. 1, we further show our method’s application to a domain with two disjoint obstacles. In such cases, vorticity-based methods typically use an incorrect boundary condition, and the vorticity-based Monte Carlo method [Rioux-Lavoie and Sugimoto et al. 2022] fails to simulate the physics correctly. Applying the modification proposed by Yin et al. [2023] for vorticity-based simulation to a Monte Carlo method is not straightforward and has not been demonstrated. Hence, ours is currently the only Monte Carlo method that can produce correct simulation results in these scenarios. While the relative performance depends heavily on parameter choices, for the specific simulations in Fig. 2, ours took 101.0s per time step while the vorticity-based method took 87.6s, with a difference of about 15%.

Buoyancy. In addition to the velocity itself, we can additionally simulate the advection and diffusion of the temperature field T and the smoke concentration field s in a manner similar to the velocity field to add buoyancy effects (Figures 1, 4, and 5). We use the Boussinesq approximation, which assumes the fluid density variation to be negligible, so that the buoyancy force can be computed using $\mathbf{f} = [\alpha s - \beta(T - T_{\text{ambient}})]\mathbf{g}$, where $\mathbf{g}$ is the gravitational constant vector, α and β are positive parameters, and T_{ambient} is the ambient temperature constant [Fedkiw et al. 2001; Foster and Metaxas 1997]. Naively incorporating buoyancy into the vorticity-based Monte Carlo formulation by Rioux-Lavoie and Sugimoto et al. [2022] would require the curl of $\mathbf{f}$ [Park and Kim 2005], which would, in turn, require finite difference approximations with some additional errors and ruin the cache-structure-agnostic nature of pointwise estimators.

Divergence control. As our formulation is based on the standard pressure solve framework, we can easily add velocity sinks and




Figure 3: Variants of our method with different configurations tested on the Kármán vortex street scene with resolution 128×256 . Except for (a), all results are generated with our Monte Carlo method. (b) to (d) compare the caching alternatives, (e) to (g) demonstrate simulations with various Reynolds numbers Re , (h) shows the application of the advection-reflection method to yield reduced numerical dissipation, and (i) to (k) show the application of PIC/FLIP methods. Note that (i) to (k) have no density diffusion because they employ density advection by particles. Given their agreement with the result of the traditional grid-based method (a), we consider all variants of our method to produce qualitatively reasonable results. 1m56s

sources in our simulation for better artistic control (Fig. 4 (b)) or other expansion/contraction effects, by modifying the pressure solve step similar to the grid-based formulation [Feldman et al. 2003]. We add another term to the right-hand side of Eq. 5, which leads to an additional integral term in our formulation. Our implementation supports Dirac delta-type sinks and sources, and we sample the location of sinks and sources directly. It should also generalize straightforwardly to more general volumetric sinks and sources with an appropriate sampling technique. This divergence control technique is yet another advantage of adopting a velocity-based formulation, as it is impossible under vorticity formulations.

Viscous fluids. We test our solver in the case of viscous fluids using the von Kármán vortex street scenario (Fig. 3 (e) to (g)). For this simulation, as we increase the viscosity coefficient, or equivalently, as we decrease the Reynolds number, the flow is expected to become less turbulent. We observe that our simulation qualitatively

produces the expected type of flow for each Reynolds number configuration we tested. For all the results other than those in Fig. 3 (e) to (g), we disabled the diffusion step.

Advection-reflection solver. In contrast to the case of high viscosity, we may want to introduce less artificial viscosity to the solver. It is known that even without the diffusion step, grid-based solvers suffer from artificial viscosity in the solver because of the dissipation of energy in the advection and projection steps. To address this problem, Zehnder et al. [2018] recently introduced an advection-reflection solver, which still uses the same building blocks of advection and projection steps as the standard advection-projection solvers use, but in a more careful way, to reduce artificial viscosity with the same number of projection steps. We adapted its second-order variant [Narain et al. 2019] to our Monte Carlo solver (Fig. 3 (h)), and we observe our method similarly benefits from the use of the advection-reflection method, producing more turbulent animation.

Monte Carlo PIC/FLIP. The numerical diffusion effect we have discussed above is partially due to the semi-Lagrangian advection scheme with some grid data interpolation that our method and many traditional methods employ. To address this problem in the context of classical grid-based solvers, the PIC/FLIP methods [Brackbill and Ruppel 1986; Harlow 1962; Zhu and Bridson 2005] have been popularized. The idea of PIC/FLIP is that we handle the advection of velocity using particles by (forward) Lagrangian advection, transfer the velocity from particles to grid, perform projection on the grid, and update particle velocities by transferring grid data back to the particles. PIC and FLIP differ in their last grid-to-particle transfer step: PIC transfers the velocity itself whereas FLIP transfers the velocity update. FLIP is expected to produce more turbulent and noisy results. One can also blend the two transfer methods. Our velocity-based Monte Carlo framework can utilize this idea straightforwardly (Fig. 3 (i) to (k)). In our version, at each step, we first perform the standard particle-to-grid transfer. Then, instead of advecting the particles using the projected grid velocities at the particles' locations via interpolation, we evaluate the velocity (or



Figure 4: Buoyancy and divergence control (resolution 256^2). We simulate a hot smoke plume rising towards a blue circular obstacle in 2D using the Boussinesq buoyancy model (left). We can also add a velocity sink to the scene, indicated with a red dot, causing the smoke to be sucked into the sink (right). 3m49s

velocity update) from the grid data exactly at the positions required for particle advection. Exploiting the pointwise estimation capability in this way lets us avoid extra interpolation errors for this step.

Performance. We measured the timing of our simulations in Fig. 3 using their GPU implementations against a basic single-thread CPU implementation of traditional grid-based fluids on a workstation with an Intel Xeon Silver 4316 CPU and an NVIDIA RTX A5000 GPU. Stable fluids with cut-cell boundary handling [Batty et al. 2007] (a) runs at 0.064 seconds per time step, while the fastest implementation of our solver that caches the velocity field after each substep (b) runs at 7.8 seconds per time step, two orders of magnitude slower. If we disable the VPL with this setup, the runtime increases to 109 seconds per time step. For the other caching choices, the runtime for the one without caching after projection (c) increases to 32.1 seconds per time step because we perform the projection at all points used in the RK3 advection, and the one without caching after advection (d) runs at 10.1 seconds per time step due to the increased number of evaluation points for the advection. In general, the walk-on-boundary method for handling obstacles requires a large number of paths, and further efficiency improvements would be desirable to improve the method’s practicality. The addition of diffusion steps in (e) to (g) resulted in the increase of runtime by 48% (e) to 84% (g) compared to their baseline counterpart implementation (c).

Error analysis. We performed a convergence test for the projection step in the absence of obstacles. The estimator is unbounded in this case, and we can observe in Fig. 6 that the error decreases at the inverse square root rate, as we increase the number of samples to estimate the volume term and the area term, respectively. We also observe that using a simpler formulation or a simpler sampling technique makes the estimator converge slower or completely fail to converge. In Fig. 7, we compared our simulations with different numbers of samples. We observe that using low sample counts can blow up the simulation, and the noise of the velocity estimate is typically larger around the boundary. As Sugimoto et al. [2023] discuss, the variance of the walk-on-boundary method is particularly large for interior Neumann problems in non-convex domains, and Fig. 7 shows such a scene.

4 CONCLUSION AND DISCUSSION

We have presented a velocity-based Monte Carlo method for fluid simulation. We developed a pointwise estimator for each substep and demonstrated flexible integration of existing velocity-based techniques into our method.

One could also consider using our velocity-based pointwise estimators to design a cacheless scheme without any spatial discretization whatsoever, similar to the vorticity-based strategy described by Rioux-Lavoie and Sugimoto et al. [2022]. As Rioux-Lavoie and Sugimoto et al. [2022] pointed out, this kind of fully recursive strategy offers truly pointwise estimation at the cost of large variance and thus has a high computational cost. While of theoretical value, this approach is impractical today because the estimation cost grows exponentially as the simulation proceeds in time, depending on the specific path sampling strategy.

Our method’s computational speed is not yet comparable to traditional methods as a relatively large number of samples is required

to produce results with low enough error. Moreover, having boundaries in the scene necessitates the use of an advanced Monte Carlo solver for (Neumann) boundary value problems, with corresponding increased variance, leading to higher costs.

In addition to efficiency improvements, our method will benefit from any future improvements to the walk-on-boundary method. Support for spatially varying diffusion equations under the walk-on-boundary method, analogous to what was proposed for the walk-on-spheres method [Sawhney and Seyb et al. 2022], would allow us to support variable fluid densities in our framework. Applying differentiable rendering techniques similar to Yilmazer et al. [2022] would enable the use of our solver for inverse problems. Supporting double-sided boundary conditions by extending the boundary integral formulation would also be a useful extension.

We designed our solver with a velocity-only formulation, but it may be beneficial to revisit this choice as (grid-based) methods utilizing velocity gradient information [Feng et al. 2023; Jiang et al. 2015; Nabizadeh et al. 2022] have recently shown promising improvements. Other ideas in the velocity-based fluid literature, such as the method of characteristic mapping [Qu et al. 2019] to lessen advection errors or energy-preserving integrators [Mullen et al. 2009] to reduce splitting errors, are also worth investigating in our Monte Carlo context.

Lastly, we believe our method paves the way for a full Monte Carlo simulation of liquids with free surfaces, and we look forward to the continued advancement of Monte Carlo techniques for physical simulations more broadly.

ACKNOWLEDGMENTS

This research was partially funded by NSERC Discovery Grants (RGPIN-2021-02524 & RGPIN-2020-03918), CFI-JELF (Grant 40132), and a grant from Autodesk. Ryusuke Sugimoto was partially funded by David R. Cheriton Graduate Scholarship. This research was enabled in part by support provided by SHARCNET and the Digital Research Alliance of Canada. We thank Joel Wretborn, Brooke Dolny, Rikin Gurditta, and Clara Kim for proofreading. We would like to thank the anonymous reviewers for their constructive evaluations and feedback.

REFERENCES

- Christopher Batty, Florence Bertails, and Robert Bridson. 2007. A fast variational framework for accurate solid-fluid coupling. *ACM Trans. Graph.* 26, 3 (jul 2007), 100–es. <https://doi.org/10.1145/1276377.1276502>
- Blender Foundation. 2023. *Blender 4.0*. Blender Institute, Amsterdam. <http://www.blender.org>
- J.U. Brackbill and H.M. Ruppel. 1986. FLIP: A method for adaptively zoned, particle-in-cell calculations of fluid flows in two dimensions. *J. Comput. Phys.* 65, 2 (1986), 314–343. [https://doi.org/10.1016/0021-9991\(86\)90211-1](https://doi.org/10.1016/0021-9991(86)90211-1)
- Robert Bridson. 2015. *Fluid simulation for computer graphics*. A K Peters/CRC Press, New York. <https://doi.org/10.1201/9781315266008>
- Wesley Chang, Venkataram Sivaram, Derek Nowrouzezahrai, Toshiya Hachisuka, Ravi Ramamoorthi, and Tzu-Mao Li. 2023. Parameter-Space ReSTIR for Differentiable and Inverse Rendering. In *ACM SIGGRAPH 2023 Conference Proceedings* (Los Angeles, CA, USA) (SIGGRAPH '23). Association for Computing Machinery, New York, NY, USA, Article 18, 10 pages. <https://doi.org/10.1145/3588432.3591512>
- Todd F. Dupont and Yingjie Liu. 2007. Back and Forth Error Compensation and Correction Methods for Semi-Lagrangian Schemes with Application to Level Set Interface Computations. *Math. Comp.* 76, 258 (2007), 647–668. <http://www.jstor.org/stable/40234399>
- S. Ermakov and A. Sipin. 2009. The “walk in hemispheres” process and its applications to solving boundary value problems. *Vestnik St. Petersburg University: Mathematics* 42 (09 2009), 155–163. <https://doi.org/10.3103/S1063454109030029>
- Ronald Fedkiw, Jos Stam, and Henrik Wann Jensen. 2001. Visual Simulation of Smoke. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. Association for Computing Machinery, New York, NY, USA, 15–22. <https://doi.org/10.1145/383259.383260>
- Bryan E. Feldman, James F. O'Brien, and Okan Arikan. 2003. Animating suspended particle explosions. *ACM Trans. Graph.* 22, 3 (jul 2003), 708–715. <https://doi.org/10.1145/882262.882336>
- Bryan E. Feldman, James F. O'Brien, and Bryan M. Klingner. 2005. Animating gases with hybrid meshes. *ACM Trans. Graph.* 24, 3 (jul 2005), 904–909. <https://doi.org/10.1145/1073204.1073281>
- Fan Feng, Jinyuan Liu, Shiyong Xiong, Shuqi Yang, Yaorui Zhang, and Bo Zhu. 2023. Impulse Fluid Simulation. *IEEE Transactions on Visualization and Computer Graphics* 29, 6 (jun 2023), 3081–3092. <https://doi.org/10.1109/TVCG.2022.3149466>
- Nick Foster and Dimitris Metaxas. 1996. Realistic Animation of Liquids. *Graphical Models and Image Processing* 58, 5 (1996), 471–483. <https://doi.org/10.1006/gnip.1996.0039>
- Nick Foster and Dimitris Metaxas. 1997. Modeling the Motion of a Hot, Turbulent Gas. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 181–188. <https://doi.org/10.1145/258734.258838>
- Charles P. Frahm. 1983. Some novel delta-function identities. *American Journal of Physics* 51, 9 (09 1983), 826–829. <https://doi.org/10.1119/1.13127>
- Francis H Harlow. 1962. *The particle-in-cell method for numerical solution of problems in fluid dynamics*. Technical Report. Los Alamos National Lab. <https://doi.org/10.2172/4769185>
- Francis H. Harlow and J. Eddie Welch. 1965. Numerical Calculation of Time-Dependent Viscous Incompressible Flow of Fluid with Free Surface. *The Physics of Fluids* 8, 12 (12 1965), 2182–2189. <https://doi.org/10.1063/1.1761178>
- V Hnizdo. 2011. Generalized second-order partial derivatives of $1/r$. *European Journal of Physics* 32, 2 (jan 2011), 287. <https://doi.org/10.1088/0143-0807/32/2/003>
- Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. 2013. Robust Inside-Outside Segmentation Using Generalized Winding Numbers. *ACM Trans. Graph.* 32, 4, Article 33 (jul 2013), 12 pages. <https://doi.org/10.1145/2461912.2461916>
- Chenfanfu Jiang, Craig Schroeder, Andrew Selle, Joseph Teran, and Alexey Stomakhin. 2015. The Affine Particle-in-Cell Method. *ACM Trans. Graph.* 34, 4, Article 51 (jul 2015), 10 pages. <https://doi.org/10.1145/2766996>
- James T. Kajiya. 1986. The Rendering Equation. *SIGGRAPH Comput. Graph.* 20, 4 (aug 1986), 143–150. <https://doi.org/10.1145/15886.15902>
- Alexander Keller. 1997. Instant Radiosity. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 49–56. <https://doi.org/10.1145/258734.258769>
- Patrick Mullen, Keenan Crane, Dmitry Pavlov, Yiyang Tong, and Mathieu Desbrun. 2009. Energy-preserving integrators for fluid animation. *ACM Trans. Graph.* 28, 3, Article 38 (jul 2009), 8 pages. <https://doi.org/10.1145/1531326.1531344>
- Mervin E. Muller. 1956. Some Continuous Monte Carlo Methods for the Dirichlet Problem. *The Annals of Mathematical Statistics* 27, 3 (1956), 569 – 589. <https://doi.org/10.1214/aoms/1177728169>
- Mohammad Sina Nabizadeh, Ravi Ramamoorthi, and Albert Chern. 2021. Kelvin Transformations for Simulations on Infinite Domains. *ACM Trans. Graph.* 40, 4, Article 97 (jul 2021), 15 pages. <https://doi.org/10.1145/3450626.3459809>
- Mohammad Sina Nabizadeh, Stephanie Wang, Ravi Ramamoorthi, and Albert Chern. 2022. Covector Fluids. *ACM Trans. Graph.* 41, 4, Article 113 (jul 2022), 16 pages. <https://doi.org/10.1145/3528223.3530120>
- Rahul Narain, Jonas Zehnder, and Bernhard Thomaszewski. 2019. A Second-Order Advection-Reflection Solver. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 2, Article 16 (jul 2019), 14 pages. <https://doi.org/10.1145/3340257>
- Art B. Owen. 2013. *Monte Carlo theory, methods and examples*. <https://artowen.su.domains/mc/>
- Sang Il Park and Myoung Jun Kim. 2005. Vortex Fluid for Gaseous Phenomena. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Los Angeles, California) (SCA '05). Association for Computing Machinery, New York, NY, USA, 261–270. <https://doi.org/10.1145/1073368.1073406>
- Steven G. Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, and Martin Stich. 2010. OptiX: A General Purpose Ray Tracing Engine. *ACM Trans. Graph.* 29, 4, Article 66 (jul 2010), 13 pages. <https://doi.org/10.1145/1778765.1778803>
- Ziyin Qu, Xinxin Zhang, Ming Gao, Chenfanfu Jiang, and Baoquan Chen. 2019. Efficient and conservative fluids using bidirectional mapping. *ACM Trans. Graph.* 38, 4, Article 128 (jul 2019), 12 pages. <https://doi.org/10.1145/3306346.3322945>
- Damien Rioux-Lavoie, Ryusuke Sugimoto, Tümay Özdemir, Naoharu H. Shimada, Christopher Batty, Derek Nowrouzezahrai, and Toshiya Hachisuka. 2022. A Monte Carlo Method for Fluid Simulation. *ACM Trans. Graph.* 41, 6, Article 240 (nov 2022), 16 pages. <https://doi.org/10.1145/3550454.3555450>
- Karl K. Sabelfeld. 1982. Vector algorithms in the Monte-Carlo method for solving systems of second-order elliptic equations and Lamé's equation. *Doklady Akademii Nauk SSSR* 262, 5 (1982), 1076–1080. <https://www.mathnet.ru/eng/dan45069> (In Russian).
- Karl K. Sabelfeld and Nikolai A. Simonov. 1994. *Random Walks on Boundary for Solving PDEs*. De Gruyter, Berlin. <https://doi.org/10.1515/9783110942026>
- Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-Based Methods on Volumetric Domains. *ACM Trans. Graph.* 39, 4, Article 123 (aug 2020), 18 pages. <https://doi.org/10.1145/3386569.3392374>
- Rohan Sawhney, Bailey Miller, Ioannis Gkioulekas, and Keenan Crane. 2023. Walk on Stars: A Grid-Free Monte Carlo Method for PDEs with Neumann Boundary Conditions. *ACM Trans. Graph.* 42, 4 (aug 2023), 22 pages. <https://doi.org/10.1145/3592398>
- Rohan Sawhney, Dario Seyb, Wojciech Jarosz, and Keenan Crane. 2022. Grid-Free Monte Carlo for PDEs with Spatially Varying Coefficients. *ACM Trans. Graph.* 41, 4, Article 53 (jul 2022), 17 pages. <https://doi.org/10.1145/3528223.3530134>
- Andrew Selle, Ronald Fedkiw, ByungMoon Kim, Yingjie Liu, and Jarek Rossignac. 2008. An Unconditionally Stable MacCormack Method. *Journal of Scientific Computing* 35, 2 (2008), 350–371. <https://doi.org/10.1007/s10915-007-9166-4>
- Nikolai A Simonov. 2008. Walk-on-Spheres Algorithm for Solving Boundary-Value Problems with Continuity Flux Conditions. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, Berlin, Heidelberg, 633–643. https://doi.org/10.1007/978-3-540-74496-2_38
- Jos Stam. 1999. Stable Fluids. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 121–128. <https://doi.org/10.1145/311535.311548>
- Ryusuke Sugimoto, Terry Chen, Yiti Jiang, Christopher Batty, and Toshiya Hachisuka. 2023. A Practical Walk-on-Boundary Method for Boundary Value Problems. *ACM Trans. Graph.* 42, 4, Article 81 (jul 2023), 16 pages. <https://doi.org/10.1145/3592109>
- Ingo Wald. 2023. *OWL - The Optix 7 Wrapper Library*. <http://owl-project.github.io>
- Hang Yin, Mohammad Sina Nabizadeh, Baichuan Wu, Stephanie Wang, and Albert Chern. 2023. Fluid Cohomology. *ACM Trans. Graph.* 42, 4, Article 126 (jul 2023), 25 pages. <https://doi.org/10.1145/3592402>
- Ekrem Fatih Yilmazer, Delio Vicini, and Wenzel Jakob. 2022. Solving Inverse PDE Problems using Grid-Free Monte Carlo Estimators. arXiv:2208.02114 [cs.GR]
- Jonas Zehnder, Rahul Narain, and Bernhard Thomaszewski. 2018. An Advection-Reflection Solver for Detail-Preserving Fluid Simulation. *ACM Trans. Graph.* 37, 4, Article 85 (jul 2018), 8 pages. <https://doi.org/10.1145/3197517.3201324>
- Yongning Zhu and Robert Bridson. 2005. Animating Sand as a Fluid. *ACM Trans. Graph.* 24, 3 (jul 2005), 965–972. <https://doi.org/10.1145/1073204.1073298>




Figure 5: Buoyancy simulation in 3D. We simulate a smoke plume rising toward a bunny-shaped obstacle (top, resolution 128^3) and a sphere obstacle (middle, resolution 128^3) in 3D using the Boussinesq buoyancy model. We can also incorporate a more complicated smoke source shape as well (bottom, resolution $288 \times 128 \times 128$). We rendered the images with the principled volume shader in Blender [Blender Foundation 2023]. 0m15s




Figure 6: Projection error without solid boundaries. Our projection step projects out the divergent velocity mode from the velocity field, converting the top left field into the one on the bottom left (resolution 256^2). We visualize the velocity field using arrows and we show its magnitude using the darkness of the red color as well. We measure the root mean squared errors measured against the analytical solution and plot them by altering the number of volume samples N_V with a fixed number of area samples $N_A = 10^7$ in the middle and by altering N_A with $N_V = 10^7$ on the right. We observe the inverse square root convergence rate in both cases with our formulation with the default sampling strategy (a) as described in Section 2.3.1, while the bias due to the volume term eventually dominates the error on the plot for N_A as we increase N_A . We also tested the estimator by (separately) replacing the volume term importance sampling according to the PDF proportional to $1/r^{(d-1)}$ with a uniform sampling (b), and disabling the global velocity shift described in Eq. 11 (c). Both of these alternatives either converge slower than the proposed method or fail to converge.




Figure 7: Flow in a 2D bounded non-convex domain (resolution 256^2). We simulate colored smoke advected by a velocity field induced by a vortex pair inside a bunny shape. On the left, the velocity field at the 40th frame is visualized. In these simulations, the number of samples used in each simulation is 10 times (top), 1 times (middle), and 0.1 times (bottom) the number of samples used in other 2D scenes in this paper. We observe that using a small sample count can lead to outright failure of the simulation (bottom). Using a modest amount of samples (middle) in this scenario still exhibits large errors and jitter in the animation (see the supplemental video), likely due to larger velocity errors near the boundary and large globally correlated error from the boundary value caching strategy. With the highest sample count (top) the simulation yields a smooth velocity field. **4m18s**

A WALK-ON-BOUNDARY FOR PROJECTION

For Eq. 16 and Eq. 17, we defined the normal directions as pointing outward from the fluid domain to describe both interior and exterior problems with a single pair of equations, whereas the previous work [Sabelfeld and Simonov 1994; Sugimoto et al. 2023] had two separate pairs of equations for interior and exterior problems. Hence, our equations for exterior problems differ from Sabelfeld and Simonov [1994] and Sugimoto et al. [2023] in their signs.

Monte Carlo Estimation. Based on Eq. 16 and Eq. 17, we use a walk-on-boundary Monte Carlo estimator [Sabelfeld and Simonov 1994; Sugimoto et al. 2023]. Note that Eq. 17 contains the unknown density function μ itself on the left-hand side of the equation and in one of the integrands on the right-hand side, making it a recursive integral equation similar to the rendering equation [Kajiya 1986] used in path tracing for light transport simulation rather than a simple non-recursive one as in Section 2.3.1. Following Sugimoto et al. [2023], we truncate the recursion after M steps and multiply the contribution of the longest path by 0.5 to design a biased estimator.

We can consider various sampling strategies to estimate the solution to the truncated recursive integral equation with Monte Carlo methods. Following Sugimoto et al. [2023], we use a forward estimator to solve this in our implementation: to sample a path consisting of multiple boundary points, we randomly sample a boundary point first, take a random walk on randomly sampled boundary points,

and finally connect them to each individual evaluation point. In the below, for $P_U(\mathbf{x}_0)$, we use uniform sampling to sample a point $\mathbf{x}_0$ on the solid boundary $\partial\Omega$. Then, for $P_R(\mathbf{x}_{m+1}|\mathbf{x}_m)$, we uniformly randomly sample a direction from a unit hemisphere for a line that goes through the point $\mathbf{x}_m$ and uniformly randomly sample one of the intersection points between the line and the solid boundary $\partial\Omega$ to sample the next point $\mathbf{x}_{m+1}$. Recursively applying this sampling strategy lets us sample a series of points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{M+1}$ to form a path, given previous points in the path. We use this strategy because the PDF associated with this specific forward sampling strategy is known to be proportional to the integral kernel $\frac{\partial G}{\partial \mathbf{n}_x}$ we have in Eq. 17.

Compared to the simplest formulation for the Laplace equation [Sugimoto et al. 2023], we have some additional non-recursive terms in Eq. 16 and Eq. 17 as a result of the transformations we discussed in the main paper, and we need to consider how to sample such terms, too. We choose to sample the non-recursive contributions on $\partial\Omega$ using the forward estimator as well for the purpose of importance sampling while estimating the other terms similarly to Section 2.3.1. To do so, we first define two estimators for length m contributions, $\langle \mu_m^1(\mathbf{x}) \rangle$ and $\langle \mu_m^2(\mathbf{x}) \rangle$, recursively, based on Eq. 17: we define the base case for length 1 contributions as

$$\langle \mu_1^1(\mathbf{x}_1) \rangle = \frac{2 \frac{\partial G}{\partial \mathbf{n}_x}(\mathbf{x}_1, \mathbf{x}_0)}{P_R(\mathbf{x}_1|\mathbf{x}_0)} \mathbf{n}(\mathbf{x}_0) \cdot \{\mathbf{u}_3(\mathbf{x}_0) - \mathbf{u}_3(\mathbf{x}_1)\}, \quad (19)$$

$$\langle \mu_1^2(\mathbf{x}_0) \rangle = 2\mathbf{n}(\mathbf{x}_0) \cdot \{-\langle E_V(\mathbf{x}_0) \rangle - \langle E_A(\mathbf{x}_0) \rangle + \mathbf{u}_3(\mathbf{x}_0) - \mathbf{u}_s(\mathbf{x}_0)\}, \quad (20)$$

where $\langle \mu_1^1(\mathbf{x}_1) \rangle$ corresponds to the non-recursive contribution from the first integral in Eq. 17 and $\langle \mu_2^1(\mathbf{x}_0) \rangle$ corresponds to the rest of the non-recursive contributions in Eq. 17. Then, we define the contributions from longer paths as

$$\langle \mu_{m+1}^1(\mathbf{x}_{m+1}) \rangle = \frac{-2 \frac{\partial G}{\partial \mathbf{n}_x}(\mathbf{x}_{m+1}, \mathbf{x}_m)}{P_R(\mathbf{x}_{m+1} | \mathbf{x}_m)} \langle \mu_m^1(\mathbf{x}_m) \rangle, \quad (21)$$

$$\langle \mu_{m+1}^2(\mathbf{x}_m) \rangle = \frac{-2 \frac{\partial G}{\partial \mathbf{n}_x}(\mathbf{x}_m, \mathbf{x}_{m-1})}{P_R(\mathbf{x}_m | \mathbf{x}_{m-1})} \langle \mu_m^2(\mathbf{x}_{m-1}) \rangle. \quad (22)$$

Note in particular, when we use the above-mentioned uniform line intersection sampling for $P_R(\mathbf{y} | \mathbf{x})$,

$$\frac{2 \frac{\partial G}{\partial \mathbf{n}_x}(\mathbf{x}, \mathbf{y})}{P_R(\mathbf{x} | \mathbf{y})} = \kappa(\mathbf{y}, \mathbf{x}) \cdot \text{sgn}(\mathbf{n}(\mathbf{x}) \cdot (\mathbf{x} - \mathbf{y})), \quad (23)$$

where $\kappa(\mathbf{y}, \mathbf{x})$ is the number of intersection points that the line that goes through the two points has, excluding point $\mathbf{y}$, and sgn is the sign function. Then, based on Eq. 16, we can estimate the pressure gradient by taking N_P sample paths in addition to the terms in Eq. 12:

$$\begin{aligned} & \langle \nabla_{\mathbf{x}} \bar{p}(\mathbf{x}) \rangle \\ &= \langle E_V(\mathbf{x}) \rangle + \langle E_A(\mathbf{x}) \rangle \\ &+ \frac{1}{N_P} \sum_{k=1}^{N_P} \left[-\frac{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}_0^k)}{2P_U(\mathbf{x}_0^k)} \mathbf{n}(\mathbf{x}_0^k) \cdot \{\mathbf{u}_3(\mathbf{x}_0^k) - \mathbf{u}_3(\mathbf{x})\} \right. \\ &+ \frac{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}_{M+1}^k)}{2P_U(\mathbf{x}_0^k)} \langle \mu_M^1(\mathbf{x}_{M+1}^k) \rangle + \frac{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}_M^k)}{2P_U(\mathbf{x}_0^k)} \langle \mu_M^2(\mathbf{x}_M^k) \rangle \\ &\left. + \sum_{m=1}^{M-1} \frac{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}_{m+1}^k)}{P_U(\mathbf{x}_0^k)} \langle \mu_m^1(\mathbf{x}_{m+1}^k) \rangle + \frac{\nabla_{\mathbf{x}} G(\mathbf{x}, \mathbf{x}_m^k)}{P_U(\mathbf{x}_0^k)} \langle \mu_m^2(\mathbf{x}_m^k) \rangle \right]. \quad (24) \end{aligned}$$

On the right-hand side, the first line estimates the last two integrals in Eq. 16, the second line estimates the non-recursive terms in the first integral, the third line estimates the longest path contributions, and the last line estimates the shorter path contributions. Note that while we do not explicitly indicate so, the contributions from paths of any length $\langle \mu_m^1(\mathbf{x}_{m+1}^k) \rangle$ and $\langle \mu_m^2(\mathbf{x}_m^k) \rangle$ implicitly depend on the sampled boundary point $\mathbf{x}_0^k$ due to the sampling strategy we employ.

To use the estimator Eq. 24, the most naive approach would be to generate sample paths for each individual evaluation point, which has a high computational cost. Instead, we use a boundary value caching strategy similar to the virtual point light method in rendering [Keller 1997]. We first compute the contributions of N_P subpaths up to the point before we connect them to the evaluation points and cache them at $M+1$ boundary points per path. We then connect all of them to all evaluation points. This significantly increases the effective number of sample paths per evaluation point without increasing the computational cost too much while introducing a correlation of estimates between evaluation points. Even this correlation can be preferable for our application because it guarantees that the contributions from these paths changes smoothly across evaluation points.

Sugimoto et al. [2023, Figure 9] apply a similar caching technique to their Neumann problem walk-on-boundary solver based on the same single-layer boundary integral formulation as well, but they use a backward estimator, which works only with a less efficient resampled importance sampling strategy. In contrast, our method can utilize a forward estimator with more efficient line intersection importance sampling. This is at the cost of increasing storage per cache point by a small constant multiplication factor and increasing computation per evaluation point when we sum up the contributions from all cache points.

Similar to Section 2.3.1, for the estimation of volume integrals $\langle E_V(\mathbf{x}) \rangle$ in Eq. 20 and Eq. 24, we use the importance sampling strategy so PDF P_V is proportional to $1/r^{(d-1)}$, and let $\mathbf{S}(\mathbf{x}, \mathbf{y}) = \mathbf{0}$ for all $\mathbf{y}$ outside the simulation domain. Additionally, we need to set $\mathbf{S}(\mathbf{x}, \mathbf{y}) = \mathbf{0}$ for all $\mathbf{y}$ inside of solid obstacles, too. We perform this insidedness test by casting a ray from point $\mathbf{y}$ in a random direction and taking a sum of its intersection signs. This strategy can be considered a Monte Carlo estimator for the generalized winding number [Jacobson et al. 2013]. We use uniform sampling for $\langle E_A(\mathbf{x}) \rangle$.

For all 2D examples in the paper except Fig. 7, we use $N_V = N_A = 5 \cdot 10^5$ for direct contributions to each evaluation point. For indirect (path) contributions, we use $N_P = 5 \cdot 10^5$ paths with path length 4. For each path, the volume and area term sample counts for the pseudo boundary are $N_V = N_A = 10$.

B WALK-ON-BOUNDARY FOR DIFFUSION

To get scalar diffusion equations from Eq. 18, we first absorb the viscosity coefficient ν into the variable to define $\mathbf{w}(\mathbf{x}, s) = \nu \bar{\mathbf{u}}(\mathbf{x}, s/\nu)$ and redefine the range of time s accordingly:

$$\begin{aligned} \frac{\partial \mathbf{w}(\mathbf{x}, s)}{\partial s} &= \nabla^2 \mathbf{w}(\mathbf{x}, s) & \text{for } \mathbf{x} \in \Omega, s \in (0, \nu \Delta t) \\ \mathbf{w}(\mathbf{x}, s) &= \mathbf{w}^*(\mathbf{x}, s) = \bar{\mathbf{u}}(\mathbf{x}, s/\nu) & \text{for } \mathbf{x} \in \partial\Omega, s \in (0, \nu \Delta t), \text{ and} \\ \mathbf{w}(\mathbf{x}, 0) &= \mathbf{w}^*(\mathbf{x}, 0) = \bar{\mathbf{u}}(\mathbf{x}, 0) & \text{for } \mathbf{x} \in \Omega, \end{aligned} \quad (25)$$

where the asterisk in $\mathbf{w}^*$ indicates that it is a given function. This is still a vector-valued equation, but because we only need to deal with simple Dirichlet boundary conditions, we can solve the vector diffusion equation Eq. 25 component-wise. Thus, we describe the scalar diffusion equation solver for one scalar component of $\mathbf{w}$, w below.

We will summarize the diffusion walk-on-boundary solver from the book by Sabelfeld and Simonov [1994, Chapter 4] here, focusing on its use in our context. For more general cases not described here, such as Neumann problems and nonhomogeneous problems, readers should refer to the book. The idea of the diffusion walk-on-boundary method is very similar to the one for the Poisson equation: we convert the partial differential equation into an integral equation and solve it using a ray-tracing-style solver. However, with the additional time dependency, we need to additionally consider the time variation in the diffusion equation and take walks in the space-time domain. Note that this solver does not discretize the time domain within each time step, unlike most traditional methods that discretize the time with a finite difference.

First, we define the fundamental solution for the diffusion equation Eq. 25, also known as the heat kernel,

$$Z(\mathbf{x}, s; \mathbf{y}, \tau) = \Theta(t') (4\pi t')^{-d/2} e^{-r^2/4t'}, \quad (26)$$

where $t' = s - \tau$ and $\Theta(\cdot)$ is the Heaviside step function.

Using the fundamental solution, we can write the solution to the diffusion equation in the form of the double layer potential and an additional initial condition term for $\mathbf{x} \in \Omega$ and $s \in [0, \nu\Delta t]$:

$$\begin{aligned} w(\mathbf{x}, s) = & - \int_0^s \int_{\partial\Omega} \frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}, \tau) \phi(\mathbf{y}, \tau) dA(\mathbf{y}) d\tau \\ & + \int_{\Omega} Z(\mathbf{x}, s; \mathbf{y}, 0) w^*(\mathbf{x}, 0) dV(\mathbf{y}), \end{aligned} \quad (27)$$

where $\phi(\cdot, \cdot)$ is an unknown density function defined for $\mathbf{x} \in \partial\Omega$ and $s \in [0, \nu\Delta t]$. We can also derive a recursive boundary integral equation for $\phi(\cdot, \cdot)$ by taking the limit $\mathbf{x} \rightarrow \partial\Omega$ in Eq. 27:

$$\begin{aligned} \phi(\mathbf{x}, s) = & \int_0^s \int_{\partial\Omega} 2 \frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}, \tau) \phi(\mathbf{y}, \tau) dA(\mathbf{y}) d\tau \\ & + 2w^*(\mathbf{x}, s) - 2 \int_{\Omega} Z(\mathbf{x}, s; \mathbf{y}, 0) w^*(\mathbf{x}, 0) dV(\mathbf{y}). \end{aligned} \quad (28)$$

Monte Carlo Estimation. We design a Monte Carlo estimator based on Eq. 27 and Eq. 28. First, we define an estimator for the initial condition term with N_I samples using PDF P_I as

$$\langle E_I(\mathbf{x}, s) \rangle = \frac{1}{N_I} \sum_{i=1}^{N_I} \frac{Z(\mathbf{x}, s; \mathbf{y}^i, 0)}{P_I(\mathbf{y}^i | \mathbf{x}, s)} w^*(\mathbf{y}^i, 0). \quad (29)$$

Using this estimator, we can define the estimator for Eq. 27 with N_D path samples using PDF P_D as

$$\langle w(\mathbf{x}, s) \rangle = \langle E_I(\mathbf{x}, s) \rangle + \frac{1}{N_D} \sum_{j=1}^{N_D} - \frac{\frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}^j, \tau^j)}{P_D(\mathbf{y}^j, \tau^j | \mathbf{x}, s)} \langle \phi(\mathbf{y}^j, \tau^j) \rangle \quad (30)$$

and the one for Eq. 28 with 1 sample with PDF P_E as

$$\langle \phi(\mathbf{x}, s) \rangle = 2 \frac{\frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}, \tau)}{P_E(\mathbf{y}, \tau | \mathbf{x}, s)} \langle \phi(\mathbf{y}, \tau) \rangle + 2w^*(\mathbf{x}, s) - 2\langle E_I(\mathbf{x}, s) \rangle, \quad (31)$$

which is a recursive estimator: $\langle \phi(\mathbf{y}, \tau) \rangle$ included on the right-hand side should be estimated recursively. We use a backward estimator and start generating paths for this recursive estimator from each evaluation point $(\mathbf{x}, \nu\Delta t)$ toward time 0. Note that P_D and P_E only need to sample points with time $\tau < s$ because the integral kernel $\frac{\partial Z}{\partial \mathbf{n}_y}$ is zero for $\tau \geq s$. Intuitively, this is because the solution to the heat equation depends only on previous times. Using this property, we sample points in the space-time domain so that times for the sampled sequence of points strictly decrease. We terminate the recursion when the sampled time τ is negative because the original integral domain does not contain the negative part. While the walk-on-boundary method for the Poisson equation is biased, the walk-on-boundary method for the diffusion equation gives an unbiased solution estimate.

As for the specific sampling strategy, for the initial condition sampling $P_I(\mathbf{y}^i | \mathbf{x}, s)$, we sample the points by $\mathbf{y}_i \leftarrow \mathbf{x} + \sqrt{t} \gamma_{d/2} \omega$, where $\gamma_{d/2}$ is a sample drawn from the Gamma distribution with shape parameter $d/2$ and scale parameter 1, and ω is a uniformly

random direction sampled on a unit sphere. With this sampling strategy, we get

$$\frac{Z(\mathbf{x}, s; \mathbf{y}^i, 0)}{P_I(\mathbf{y}^i | \mathbf{x}, s)} = 1. \quad (32)$$

Similar to Section 2.3.2, we set $Z(\mathbf{x}, s; \mathbf{y}^i, 0) = 0$ for all $\mathbf{y}$ inside of solid obstacles in Eq. 29.

For $P_D(\mathbf{y}, \tau | \mathbf{x}, s)$ and $P_E(\mathbf{y}, \tau | \mathbf{x}, s)$, we sample $\mathbf{y}$ using the uniform line intersection sampling as in Section 2.3.2 and sample time τ , given t , by $\tau \leftarrow s - \frac{\|\mathbf{y} - \mathbf{x}\|}{4\gamma_{d/2}}$ to get

$$\begin{aligned} \frac{\frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}, \tau)}{P_D(\mathbf{y}, \tau | \mathbf{x}, s)} &= -\kappa(\mathbf{x}, \mathbf{y}) \cdot \text{sgn}(\mathbf{n}(\mathbf{y}) \cdot (\mathbf{y} - \mathbf{x})), \quad \text{and} \\ 2 \frac{\frac{\partial Z}{\partial \mathbf{n}_y}(\mathbf{x}, s; \mathbf{y}, \tau)}{P_E(\mathbf{y}, \tau | \mathbf{x}, s)} &= -\kappa(\mathbf{x}, \mathbf{y}) \cdot \text{sgn}(\mathbf{n}(\mathbf{y}) \cdot (\mathbf{y} - \mathbf{x})). \end{aligned} \quad (33)$$

In the above, the left-hand side scaling factors of the two expressions differ by 2 because the PDF P_D and P_E differ even though we use the same strategy: point $\mathbf{x}$ for P_D lies within the domain, whereas point $\mathbf{x}$ for P_E lies on a boundary.

To draw samples from the Gamma distribution, we use $\gamma_1 \leftarrow -\ln(\alpha)$ for 2D, where α is a uniformly random sample in the range $(0, 1)$, and $\gamma_{3/2} \leftarrow \gamma_1 + \xi^2/2$ for 3D, where ξ is a standard normal sample.

For problems without boundaries, we drop the second term in Eq. 30, and the remaining first term estimates the convolution of the input field and a Gaussian function.

In our implementation, we use $N_I = 5 \cdot 10^5$ initial conditions samples at each evaluation point, with $N_D = 5 \cdot 10^5$ paths with $N_I = 10$ initial condition samples per bounce. Unlike the walk-on-boundary method for the projection step, we do not share the subpaths among evaluation points as we found that the diffusion walk-on-boundary is typically cheaper, and there was not much need to improve its efficiency relative to the projection.