Harmonic Caching for Walk on Spheres

ZIHONG ZHOU, Dartmouth College, USA EUGENE D'EON, NVIDIA, New Zealand ROHAN SAWHNEY, NVIDIA, USA WOJCIECH JAROSZ, Dartmouth College, USA




Fig. 1. We simulate heat conduction by solving a Laplace equation with Robin boundary conditions on a triangle mesh of a turkey leg containing one million primitives. The Robin conditions prescribe a radiant flux, precomputed via ray tracing on a texture to capture oven radiation (left). Our harmonic caching algorithm achieves robust temperature estimation with lower error and fewer correlation artifacts than alternative Monte Carlo solvers for PDEs.

We present a variance reduction technique for Walk on Spheres (WoS) that solves elliptic partial differential equations (PDEs) by combining overlapping harmonic expansions of the solution, each estimated using unbiased Monte Carlo random walks. Our method supports Laplace and screened Poisson equations with Dirichlet, Neumann, and Robin boundary conditions in both 2D and 3D. By adaptively covering the domain with local expansion regions and reconstructing the solution inside each region using an infinite Fourier series of the harmonic function, our method achieves over an order of magnitude lower error than traditional pointwise WoS in equal time. While low-order truncations of the series typically introduce limited bias, we also introduce a stochastic truncation scheme that eliminates this bias in the reconstructed solution. Compared to recently developed caching algorithms for WoS, such as Boundary and Mean Value Caching, our approach yields solutions with lower error and fewer correlation artifacts.

CCS Concepts: • Mathematics of computing \rightarrow Partial differential equations; Integral equations; Probabilistic algorithms.

Additional Key Words and Phrases: Monte Carlo, Walk on Spheres, Functional Expansion, FET

Authors' Contact Information: Zihong Zhou, zihong zhou.sr.gr@dartmouth.edu, Dartmouth College, USA; Eugene d'Eon, NVIDIA, New Zealand, edeon@nvidia.com; Rohan Sawhney, NVIDIA, USA, rsawhney@nvidia.com; Wojciech Jarosz, Dartmouth College, USA. wojciech.k.jarosz@dartmouth.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 1557-7368/2025/12-ART252

https://doi.org/10.1145/3763322

ACM Reference Format:

Zihong Zhou, Eugene d'Eon, Rohan Sawhney, and Wojciech Jarosz. 2025. Harmonic Caching for Walk on Spheres. *ACM Trans. Graph.* 44, 6, Article 252 (December 2025), 15 pages. https://doi.org/10.1145/3763322

1 Introduction

Monte Carlo algorithms for solving partial differential equations (PDEs), such as Walk on Spheres [Muller 1956; Sawhney and Crane 2020], avoid the meshing challenges inherent in traditional finite element methods (FEM). By sidestepping volumetric discretization, these grid-free methods provide simple and scalable solvers for elliptic PDEs–such as Laplace and screened Poisson equations—with applications across science, engineering, and computer graphics.

Yet WoS estimators often suffer from high variance, requiring many random walks to obtain low-noise solutions at individual points in the domain. This inefficiency limits their practical utility, particularly when dense or high-precision solution fields are desired. In this work, we introduce a new *functional-expansion* framework that substantially improves the efficiency of WoS and related methods in such settings. Prior variance reduction strategies, such as boundary value caching [Miller et al. 2023] and mean value caching [Bakbouk and Peers 2023], improve efficiency by spatially reusing samples. While effective in certain regimes, these methods may introduce artifacts and do not fully leverage the harmonic structure of the underlying PDE.

Our approach goes further by directly exploiting both the smoothness and harmonic structure of elliptic PDE solutions. Specifically,




Fig. 2. Left to right: For a harmonic function, the value at the center of a ball equals the average of u on its boundary (the mean value property, term a_0). More generally, the full solution inside the ball can be reconstructed as a weighted sum of separable basis functions. Each basis term is obtained by projecting the boundary data $u(R, \Theta)$ onto angular modes and scaling them radially inward using the generalized harmonic expansion.

within any ball fully contained in the domain $\Omega \subset \mathbb{R}^2$, the solution admits an analytic local Fourier expansion. For example, if Δ denotes the Laplace operator and g(x) the prescribed Dirichlet boundary data, then the solution u of the Laplace equation

$$\Delta u(x) = 0 \text{ on } \Omega, \quad u(x) = g(x) \text{ on } \partial\Omega,$$
 (1)

can be expressed in polar coordinates (r, θ) as (see Fig. 2)

$$u(r,\theta) = a_0 + 2\sum_{l=1}^{\infty} \left(\frac{r}{R}\right)^l \left(a_l \cos l\theta + b_l \sin l\theta\right), \tag{2}$$

where *R* is the radius of the ball, and a_l , b_l are the Fourier coefficients of the boundary function $u(R, \theta)$, with *l* denoting the order:

$$a_l = \frac{1}{2\pi} \int_0^{2\pi} u(R, \Theta) \cos(l\Theta) d\Theta, \quad l = 0, 1, 2, \cdots$$
 (3a)

$$b_l = \frac{1}{2\pi} \int_0^{2\pi} u(R, \Theta) \sin(l\Theta) d\Theta, \quad l = 1, 2, \cdots.$$
 (3b)

In standard WoS, the *mean value property* corresponds to retaining only the zeroth-order term in this expansion: it estimates u at the center of the ball (r=0) by averaging u over the surface at r=R, yielding $u(0)=a_0$. If we only required u at a single location, this would be all we could hope to exploit. However, the full harmonic expansion in Eq. (2) offers much more if our goal is to estimate u across a region of interest: once the boundary values $u(R,\theta)$ are known, the solution $u(r,\theta)$ at any point inside the ball follows from simple radial scaling of each harmonic. This expansion is essentially a *generalized mean value property* that reconstructs solutions within the entire ball (Fig. 2), not just the center. More generally, such Fourier expansions exploit the *separable* structure of harmonic functions to achieve a lossless form of dimensionality reduction: to reconstruct the solution within an n-dimensional ball using a harmonic series expansion, we require only a set of coefficients

encoding the boundary function on the (n-1)-dimensional sphere, while the radial component remains analytic and independent of these coefficients.

To exploit this property, we propose to estimate the coefficients (a_l,b_l) via Monte Carlo by launching random walks from the boundary of a ball in the domain. We then reconstruct an analytic estimate $u(r,\theta)$ inside the ball using a truncated harmonic expansion. This strategy falls under the general class of functional expansion tallies (FET) [Griesheimer et al. 2006], which are known to yield significantly lower variance than pointwise methods for smooth solutions. In our analysis, we show this approach is highly effective for elliptic PDEs, offering unbounded error reduction compared to pointwise WoS (assuming sufficient query resolution).

Our approach is inspired by Booth [1982, 1981], who introduced angular FETs for the screened Poisson equation in two dimensions. However, their formulation was not broadly evaluated or extended beyond a single ball, leaving open key questions about efficiency, accuracy, and applicability to more general problems. We revisit and generalize Booth's method, providing an in-depth study of its accuracy and efficiency (Sec. 3.1). Notably, while Booth [1982] discusses using the reconstructed solution both inside and outside the expansion ball (for locations within the same domain), we find that reliable reconstruction is achieved only within a subregion of the ball—typically up to r < 0.9R. Fortunately, due to the rapid decay of higher-order terms in Eq. (2), even low-order truncations suffice for high accuracy (Fig. 3). This property enables us to define compact harmonic cache records: local, continuous, differentiable, and reusable representations of the solution that can be efficiently queried and blended to approximate the solution over larger regions.

We extend spherical cache records to support Laplace and screened Poisson equations in 2D and 3D, as well as Neumann and Robin boundary conditions by invoking Walk on Stars [Miller et al. 2024b;




Fig. 3. Impact of boundary functions and truncation order on reconstruction accuracy inside a disk. We examine how reconstructed solutions of the Laplace equation vary with harmonic truncation order L under different Dirichlet boundary conditions on a unit circle. Top: For constant boundary values, the solution is recovered exactly using only the zeroth-order term (L = 0). Middle: For a bandlimited signal of order 8, perfect reconstruction requires all terms up to L = 8, with no further improvement beyond that point. Bottom: For the infinite-frequency Heaviside function, error is unavoidable but remains low across most of the disk even with modest L, due to the rapid decay of higher-order terms. Error plots (right) confirm that inaccuracies concentrate near the boundary as $r \rightarrow R$.

Sawhney et al. 2023] from the surface of each record. Building on this local construction, we then propose a practical global caching algorithm (Sec. 3.2) that covers the domain with a collection of overlapping maximal spheres. Each cache record estimates its local Fourier coefficients via unbiased Monte Carlo walks, and their combined reconstructions yield a continuous approximation of the solution field across the domain (Fig. 4). Finally, we show how to handle source terms by decomposing the problem into homogeneous and inhomogeneous components (Sec. 3.3).

We demonstrate that this harmonic caching (HC) method achieves over an order of magnitude lower error than traditional pointwise for equal computational cost, and outperforms recent variancereduction methods for WoS.

2 Related Work

Grid-Free Monte Carlo PDE Solvers. Monte Carlo methods such as Walk on Spheres [Muller 1956] avoid the meshing challenges of traditional grid-based solvers like the finite element method. WoS has gained popularity in computer graphics for its simplicity and robustness in handling complex geometries [Sawhney and Crane 2020]. Recent research has extended WoS to support more general equations [Sawhney et al. 2022] and domains [Nabizadeh et al. 2021], and to handle Neumann and Robin boundary conditions through the Walk on Stars (WoSt) algorithm [Miller et al. 2024b; Sawhney




Fig. 4. Overview of our harmonic caching algorithm for solving elliptic PDEs. We leverage the generalized mean value property of harmonic functions (Fig. 2) to reconstruct solutions from local boundary estimates. Within userspecified regions of interest (left), we adaptively place overlapping spherical cache records (middle), each storing a compact set of Fourier coefficients. Blending local reconstructions from each record yields a smooth, low-error approximation across the region (right).

et al. 2023], which is a strict generalization of WoS. These advances have enabled applications in fluid dynamics, thermal imaging [Bati et al. 2023: De Lambilly et al. 2023: Jain et al. 2024: Rioux-Lavoie et al. 2022], and inverse problems [Miller et al. 2024a; Yilmazer et al. 2024; Yu et al. 2024]. Walk on Boundary (WoB) [Sugimoto et al. 2023] provides an alternative grid-free method but generally exhibits higher bias and variance in non-convex domains.

Variance Reduction for Monte Carlo Solvers. Despite its advantages, WoS produces inherently noisy estimates, naturally motivating the development of variance reduction techniques. Mean Value Caching (MVC) [Bakbouk and Peers 2023] leverages a volumetric version of the mean value principle by reusing walks within a ball around each evaluation point. However, its efficiency is limited near boundaries where few cache samples can be gathered, leading to correlation artifacts. Boundary value cache (BVC) [Miller et al. 2023] computes solution and gradient estimates at boundary cache points and splats their contributions across the domain, enabling efficient sample reuse but introducing singular artifacts near cache points due to the absence of importance sampling. Reverse WoS (RWoS) [Qi et al. 2022] initiates walks "in reverse" from the boundary and source locations; it is effective when boundary and source functions are localized but becomes inefficient if solutions are needed only in a small region of interest (ROI).

Neural approaches for reducing noise have also emerged. Neural Caching [Li et al. 2023; Nam et al. 2024] trains a neural network to approximate PDE solutions from WoS estimates, though training can be time-consuming and predictions exhibit difficult-to-quantify bias. Orthogonal to sample reuse and caching, Li et al. [2024] train a network to learn a control variate and its antiderivative, achieving unbiased variance reduction for WoSt estimators. Along similar lines, Huang et al. [2025] adopt online learning to train a neural field representing a guiding distribution, enabling importance sampling during the walk and thereby reducing noise.

Our approach instead explicitly leverages the harmonic structure of PDE solutions to share computation and further reduce variance. Fourier reconstruction enables effective sample reuse even near domain boundaries (unlike MVC), produces robust reconstructions without singularity artifacts (unlike BVC), and focuses the computational budget on the region of interest (unlike RWoS). We provide a comprehensive evaluation of our method against these state-of-the-art variance-reduction techniques.

Finally, in unpublished concurrent work, Czekanski et al. [2024] propose a neighbor-reuse strategy using the off-centered Poisson kernel for unbiased caching of WoS estimates in 2D. Our harmonic series expansion reduces to their approach in special cases but generalizes to handle source terms, screening, and mixed boundary conditions in both 2D and 3D. We provide in-depth analysis and comparisons in the supplemental document.

Connections to (ir)radiance caching. In developing our algorithm, we take inspiration from irradiance [Ward et al. 1988] and radiance [Jarosz et al. 2008a; Křivánek et al. 2005b] caching methods. These methods compute expensive Monte Carlo estimates of illumination only at sparse locations across the scene and reuse them via extrapolation to gain efficiency. A key ingredient is estimating accurate gradients [Jarosz et al. 2012, 2008b; Křivánek et al. 2005a; Marco et al. 2018; Ramamoorthi et al. 2007; Schwarzhaupt et al. 2012; Ward and Heckbert 1992], which allows for local first-order Taylor expansions of the solution around each cached location. This idea can also be applied to WoS [Sawhney and Crane 2020], but at the cost of introducing bias. Unfortunately, while Taylor expansions provide a reasonable approximation for smooth functions, they do not exploit the harmonic structure of the PDEs we consider. In contrast, our method uses a harmonic expansion, which we show can (optionally) be made unbiased even when truncated to a finite number of terms.

3 Method

We outline our harmonic caching method for solving screened Poisson equations with mixed boundary conditions:

$$\Delta u_f(x) - \alpha u_f = -f(x) \quad \text{on } \Omega,$$

$$u_f(x) = g(x) \quad \text{on } \partial \Omega_D,$$

$$\frac{\partial u_f(x)}{\partial n_x} = h(x) \quad \text{on } \partial \Omega_N,$$


$$\frac{\partial u_f(x)}{\partial n_x} + \mu(x)u_f(x) = \ell(x) \quad \text{on } \partial \Omega_R,$$
(4)

where $\partial\Omega_D$, $\partial\Omega_N$, and $\partial\Omega_R$ denote the Dirichlet, Neumann, and Robin portions of the boundary, and n_x is the outward normal at x. The functions g(x), h(x), and $\ell(x)$ specify the boundary data, f(x) is the source term, $\mu(x)>0$ the Robin coefficient, and $\alpha\geq 0$ the screening parameter. Equation (4) covers several common PDEs, including Laplace ($\alpha=0$, f=0) and Poisson ($\alpha=0$) equations.

We first focus on homogeneous equations (f=0) in Sec. 3.1, denoting their solution by u, and describe how to reconstruct the solution inside a single ball fully contained in the domain. For clarity, we largely restrict our exposition to two dimensions and defer the necessary extensions to three dimensions to the appendix. We then extend this local construction into a global solver in Sec. 3.2 by covering the domain with overlapping balls and blending their local predictions to approximate solutions over arbitrary ROIs. Finally, we incorporate source terms to handle inhomogeneous equations ($f\neq 0$) in Sec. 3.3, with the resulting solution denoted by u_f .

3.1 Reconstruction Within a Single Ball

The core building block of our method is as follows: we place a ball entirely inside the domain Ω , estimate the solution at a set of points on its boundary using WoS(t), and then interpolate the solution into the interior using a harmonic expansion. As illustrated in the inset, we



estimate the unknown values $u(x_0)$ by launching walks from sampled points on the boundary of the ball centered at p. Rather than averaging these values to compute u(p)—as in traditional WoS—we project the boundary data onto a Fourier basis. The resulting harmonic expansion defines a *function* that allows the solution to be evaluated *everywhere* inside the ball.

For the homogeneous screened-Poisson equation ($f = 0, \alpha \ge 0$), the harmonic expansion generalizes the Laplace solution in Eq. (2) to the form [Booth 1982, 1981]:

$$u(r,\theta) = a_0 \mathcal{R}_{\alpha}^0(r,R) + 2 \sum_{l=1}^{\infty} \mathcal{R}_{\alpha}^l(r,R) \left(a_l \cos(l\theta) + b_l \sin(l\theta) \right), \quad (5)$$

where a_l and b_l are the Fourier coefficients in (3), and $\mathcal{R}^l_{\alpha}(r,R)$ is the radial basis function

$$\mathcal{R}_{\alpha}^{l}(r,R) = \begin{cases} (r/R)^{l} & \text{if } \alpha = 0, \\ \frac{I_{l}(\sqrt{\alpha}r)}{I_{l}(\sqrt{\alpha}R)} & \text{if } \alpha > 0, \end{cases}$$
 (6)

with I_l denoting the modified Bessel function of the 1st kind. In 3D (Appendix B), the harmonic expansion instead employs spherical harmonics for the angular terms and a different radial function when $\alpha > 0$

We estimate the coefficients a_l , b_l using Monte Carlo integration:

$$\widehat{a}_{l} = \frac{1}{N} \sum_{j=1}^{N} \frac{\widehat{u}(R, \theta_{j}) \cos(l\theta_{j})}{p(\theta_{j})2\pi}, \quad \widehat{b}_{l} = \frac{1}{N} \sum_{j=1}^{N} \frac{\widehat{u}(R, \theta_{j}) \sin(l\theta_{j})}{p(\theta_{j})2\pi}, \quad (7)$$

where $\widehat{u}(R,\theta_j)$ is the solution estimated at boundary points θ_j using WoSt, and $p(\theta_j)$ is the sampling density for θ_j . To improve the rate of convergence in Monte Carlo integration for these coefficients, we adopt a uniform-jittered distribution of equally spaced but randomly rotated points on the ball's boundary [Booth 1982; Pauly et al. 2000; Ramamoorthi et al. 2012; Singh et al. 2019], giving $p(\theta_j) = 1/2\pi$ in 2D. Importantly, we reuse the same boundary samples θ_j to estimate all orders l. Using WoSt rather than WoS enables support for more general mixed boundary conditions on the domain boundary.

The expansion in Eq. (5) yields the exact solution inside the ball when the boundary values are known and the series is taken to infinity. In practice, we truncate the series at order L and estimate the boundary values via Monte Carlo, introducing both bias and variance. Fortunately, the radial basis functions $\mathcal{R}^l_\alpha(r,R)$ decay rapidly with decreasing r, especially for higher orders l. For $\alpha=0$, the weights reduce to $(r/R)^l$, while the Bessel-based functions for $\alpha>0$ have a similar shape. As a result, low-order truncations suffice to reconstruct the solution accurately throughout most of the ball. In our experiments, truncations at L=10 remain robust with low bias within a conservative subregion (r<0.9R), as illustrated in

If biased reconstruction is unacceptable, an unbiased prefix-sum estimator [Misso et al. 2022] can be formed by drawing L stochastically from a distribution over the positive integers. For example, sampling L from a geometric distribution with mean $\langle L \rangle$ corresponds to applying a Russian roulette [Arvo 1993] with termination probability $1/\langle L \rangle$ at each term of the sum in Eq. (5).

Performance. Before combining local reconstructions from multiple caches into a global solver, we first evaluate performance within a single ball compared to standard pointwise WoS. To this end, we place a ball at the center of a square domain and estimate the solution u over a grid of pixels using two approaches: (1) pointwise WoS, where each pixel launches its own k walks; and (2) our method, where the same total number of walks are uniformly distributed along the ball's boundary to estimate the Fourier coefficients and reconstruct the interior. By matching the total walk count, both methods achieve comparable runtime.

We varied the screening parameter, query resolution, reconstruction radius, and boundary function. Both approaches exhibited the expected 1/n squared-error scaling with respect to the number of walks n, confirming that pointwise estimation and function-expansion tallies converge at the same asymptotic rate until truncation bias dominates. To quantify the improvement of our method, we computed the ratio of the sum of squared errors between the two methods, considering only the reconstructed pixels.

We observed that resolution was a key factor in performance: improvement increased linearly with the number of pixels, implying that the harmonic approach can outperform pointwise estimation by an unbounded margin at high query resolutions. Put differently, unlike pointwise MC, harmonic reconstruction exhibits a *resolution-independent* sum of squared errors.

We further observed that the benefit of our approach decreases with increasing harmonic complexity of the boundary data: the error ratio falls approximately as 1/L with truncation order L. The screening parameter had little effect on relative performance, and reconstruction radii in the range [0.7R, 0.9R] provided the best balance between spatial coverage and truncation error. Additional experimental details are provided in the supplemental material.

3.2 Harmonic Caching for General Domains


Inspired by irradiance caching [Ward et al. 1988], we populate a cache database with overlapping spherical expansions to handle arbitrary regions of interest. Each cache record corresponds to a ball centered at \mathbf{p}_i with radius R_i , storing the Fourier coefficients $\{\widehat{a}_i^i, \widehat{b}_l^i\}$. As described in Sec. 3.1, the coefficients are always estimated on the surface of the maximal inscribed ball with radius $R_{\partial\Omega}(\mathbf{p}_i)$, but–based on the analysis in the previous section–we set $R_i = 0.9R_{\partial\Omega}(\mathbf{p}_i)$ to reduce truncation error for each cache record.

Reconstruction. We estimate the solution at any location within an ROI by blending contributions from overlapping cache records. For a query location \mathbf{x} , we identify the set of records $S(\mathbf{x})$ that cover

x and compute the weighted average:

$$\widehat{u}(\mathbf{x}) = \frac{\sum_{i \in S(\mathbf{x})} w(d_i) \, \widehat{u}_i(\mathbf{x})}{\sum_{i \in S(\mathbf{x})} w(d_i)},\tag{8}$$

where $\widehat{u}_i(\mathbf{x})$ is the Fourier reconstruction from the ith record (evaluated by inserting $\{\widehat{a}_l^i, \widehat{b}_l^i\}$ into Eq. (5)), and $w(d_i)$ is a weight based on the normalized distance $d_i = 1 - \frac{\|\mathbf{x} - \mathbf{p}_i\|}{R_i}$ between the query location and record center. Since reconstruction error grows with distance



from the record center (Fig. 3), constant weights can cause discontinuity artifacts and are therefore be avoided. To ensure smooth reconstruction, we use the smoothstep function $w(d_i) := 3d_i^2 - 2d_i^3$ which assigns weight one at the record center \mathbf{p}_i and falls smoothly to zero at the record boundary, without any singularities. We also experimented with alternative weighting functions that satisfy these properties, but found little difference in reconstruction quality.

Populating the cache. To populate the cache database, we adopt a lazy evaluation scheme inspired by irradiance caching [Ward et al. 1988]. For each query location in the ROI, we first identify the overlapping records $S(\mathbf{x})$ and evaluate the denominator of Eq. (8). If this sum exceeds a user-defined threshold w_{\min} , the solution is reconstructed from the existing records. Otherwise, a new cache record is generated at \mathbf{x} , inserted into the database, and included in $S(\mathbf{x})$ for reconstruction using Eq. (8). This procedure yields an adaptive, Poisson-disk-style distribution of cache records over the ROI, with w_{\min} controlling their overall density.

Choosing walk counts for coefficient estimation. Creating a new record via Eq. (7) requires choosing the number of walks N launched from the surface of the largest contained ball. This choice strongly impacts reconstruction error, as it determines the variance of the coefficient estimates. A simple option is to let the user fix N, but we found that varying N across the domain provides significant benefits. Each cache record i therefore selects its own N_i . Records with larger radii contribute to more query locations, while smaller balls produce smoother, nearly constant solutions—both suggesting that N_i should increase with R_i . After testing different strategies, we found that setting $N_i = \lambda R_i^{d-1}$, where d is the domain dimension, provides a good balance. The user controls λ , which specifies the number of walks per arc length (2D) or per solid angle (3D). We clamp $N_i \geq 32$ for all records, and in 3D, where WoSt queries are more expensive, we cap N_i at 10,000.

Additional refinement pass. When a new cache record is created, walks launched from the surface of the largest ball often pass through regions already covered by existing records. This raises a natural question: at the k-th step of a walk, should we evaluate $u(\mathbf{x}_k)$ recursively with WoSt, or reconstruct it from existing cache records via Eq. (8)? The latter is attractive because, if successful, it provides a much lower-variance estimate that effectively aggregates many prior walks.

To enable such sample reuse while ensuring thread-safe cache construction, we adopt a two-pass scheme. When populating the





Fig. 5. Our biased harmonic caching using a truncated expansion (L=10, bottom) can be debiased with a prefix-sum estimator employing stochastic truncation ($\langle L \rangle = 10$, top). We compute and visualize the bias and variance of both estimators for a Dirichlet problem [Qi et al. 2022, Figure 8] by running independent trials with different random numbers for WoS (and for stochastic truncation), each with the same cache (bottom right). Stochastic truncation yields an unbiased reconstruction but with higher variance, visible as stronger correlation artifacts in a single run (top left).

cache initially, each walk used to estimate Fourier coefficients proceeds independently without cache lookups. Once the cache is built, we duplicate it and perform an "overture" pass to refine the coefficients: for each record in the duplicated cache, we expand the largest inscribed ball and recompute the coefficients via Eq. (7). In this pass, however, $\widehat{u}(R,\theta_i)$ is obtained by looking up the previous cache via Eq. (8), rather than through recursive WoSt. Because these lookups provide lower variance estimates, we use only $0.5N_i$ samples to recompute the i-th record, reducing lookup overhead. During this pass, we also set $w_{\min}=0$: whenever a valid record overlaps the query location, the solution is taken directly from the cache.

Incorporating contributions from neighboring cache records enables substantial sample reuse, which we found highly effective for reducing noise in high-order coefficients. This strategy is conceptually similar to the recursive reuse in MVC [Bakbouk and Peers 2023] and could, in principle, be applied repeatedly to extend reuse across larger distances. In practice, however, we found a single overture pass to be sufficient.

3.3 Reconstruction with Source Term Contributions

To estimate u_f for the inhomogeneous equation in Eq. (4), we extend the harmonic expansion in Eq. (5) by incorporating source term contributions via Green's functions [Duffy 2015]. Specifically, in 2D we integrate the offcentered Green's function of the ball against the source term f over the entire ball:



$$u_f(r,\theta) := u(r,\theta) + \int_0^{2\pi} \int_0^R f(R',\Theta) G(r,\theta;R',\Theta) R' dR' d\Theta, \quad (9)$$

where $G(r, \theta; R', \Theta)$ denotes the contribution from a unit source at (R', Θ) to the evaluation point (r, θ) . While this source integral could, in principle, be further expanded using a multipole expansion [Müller and Steinmetz 1995], doing so introduces an additional




Fig. 6. Top: Increasing the truncation order L of the harmonic expansion reduces reconstruction error across all boundary condition types: Dirichlet (left), mixed Dirichlet/Neumann (middle), and Robin (right). Very small L (e.g., L=1) yields poor results, while large L increases runtime with diminishing returns. Middle: Increasing w_{\min} produces denser cache records and improves accuracy, with cost scaling roughly linearly. Bottom: Increasing λ raises N_i , the number of walks per cache record. This suppresses noise in the coefficients and reduces reconstruction error, shown here as a function of the average N_i across all cache records.

set of radial coefficients, substantially increasing both estimation cost and cache storage. Instead, we evaluate the integral directly using Monte Carlo integration:

$$\widehat{u}_f(r,\theta) = \widehat{u}(r,\theta) + \frac{1}{M} \sum_{j=1}^M \frac{f(R_j,\Theta_j) G(r,\theta;R_j,\Theta_j)}{p(R_j,\Theta_j)}, \quad (10)$$

where $\widehat{u}(r,\theta)$ is the harmonic reconstruction from overlapping records using Eq. (8), $p(R_j,\Theta_j)=\frac{1}{\pi R^2}$ is the uniform sampling density for the source variable (R_j,Θ_j) in the disk, and M is the number of source samples. Importantly, we introduce no additional modifications or cache storage to the harmonic caching algorithm in Sec. 3.2 (the underlying WoSt estimator already supports source terms [Sawhney et al. 2023, Sec. 4.6]); we simply extend the series expansion in Eq. (5) with the estimator in Eq. (10). This estimator is also invoked during the refinement pass for the coefficients, using the same number of source samples as in the final reconstruction.

4 Results and Evaluation

In this section, we evaluate harmonic caching, beginning with implementation details and biased versus unbiased variants, then analyzing parameter choices, and finally comparing against WoSt, RWoS, BVC, and MVC under equal-time and progressive settings.

Fig. 7. Computation time depends on both the number of cache records and the walks per record used for coefficient estimation. For equal runtime, it is generally better to use fewer records with more walks (left) than more records with fewer walks (right).

4.1 Implementation

We implemented our method on the CPU using the $Zombie^1$ library with single-precision floating point arithmetic [Sawhney and Miller 2023]. The method is straightforward to implement and deploy, functioning as a meta-integrator over the underlying WoSt solver. Pseudocode is provided in Algorithms 1–3.

We use a multi-reference octree [Pharr and Humphreys 2010] for efficient cache lookups. During initial cache population, concurrent access to the octree is protected by a reader—writer lock. The refinement pass, however, is lock-free: we require only read-access to the old records, while each new set of cache coefficients is updated concurrently and independently.

4.2 Biased vs Unbiased Caching

In Fig. 5, we compare biased harmonic caching (series truncation at L=10) with an unbiased variant using a prefix-sum estimator and stochastic truncation drawn from a geometric distribution with mean $\langle L \rangle = 10$. In the unbiased case, we apply the same stochastic truncation order across all evaluation points within a cache record, trading reduced variance for additional correlation. In the biased case, truncation bias becomes more pronounced as $r \to R$ (see Fig. 3). Nevertheless, even in the worst-case scenario—reconstructing up to the largest ball rather than a conservative subregion (r < 0.9R)—the biased approach remains preferable due to its lower variance and fewer correlation artifacts. Consequently, we adopt the biased harmonic caching algorithm with series truncation as our default strategy for the remaining experiments (referred to simply as HC). Misso et al. [2022] may provide a useful pathway for deriving variance-optimal unbiased series reconstructions in the future.

4.3 Parameter Ablation

Our biased caching algorithm has three main parameters: L, w_{\min} , and λ . For the Poisson equation, an additional parameter M controls the number of independent source samples used in reconstruction.

Increasing the truncation order L generally reduces error, but also requires estimating, storing, and evaluating more Fourier coefficients, with diminishing returns (Fig. 6, top row). We found that




Fig. 8. Harmonic caching shows fewer correlation artifacts in the interior than RWoS (top-left inset) and less noise near the boundary than MVC (bottom-left inset). Unlike BVC (top-right inset), it also free of singularity issues near the boundary. Results shown for a screened 3D Laplace problem.

L=10 strikes a good balance between accuracy and cost, and use this value as the default unless stated otherwise.

Increasing w_{\min} produces more cache records and an approximately linear rise in computation time. A denser cache allows evaluation points to draw higher-quality contributions from nearby records, which carry greater weight and thus reduce overall reconstruction error (Fig. 6, middle row).

Increasing λ raises the number of walks used to estimate Fourier coefficients, reducing their noise and thereby lowering reconstruction error (Fig. 6, bottom row).


Given a fixed time budget, w_{\min} and λ must be balanced. As shown in Fig. 7, allocating more walks per cache record with a sparser cache generally yields better results than the reverse.

4.4 Comparison with Walk on Stars

We evaluate harmonic caching against WoSt on 2D and 3D (screened) Laplace equations under various boundary conditions (Figs. 1, 9 and 15). Performance is strongly influenced by boundary type: pure Dirichlet problems generally produce shorter WoSt walks, while Neumann or Robin problems require longer walks and thus exhibit higher variance. As shown in Fig. 10, harmonic caching achieves the lowest error across all test models and runtimes in Fig. 9. For the Neumann-dominated 3D problem in Fig. 15, the underlying WoSt estimator, and thus the Fourier coefficients, are considerably noisier, yet harmonic caching still provides a clear reduction in error.

In Fig. 1, we further demonstrate harmonic caching on a 3D thermal simulation of a turkey leg, represented as a triangle mesh with over one million primitives. Robin boundary conditions are derived by simulating radiative heat transfer via ray tracing, while convective transfer outside the turkey leg is ignored. Solving the 3D Laplace equation yields the equilibrium temperature distribution across several vertical slices of the leg (Fig. 14). Although WoSt is well-suited to this problem—providing localized temperature previews without meshing the complex domain—it suffers from significant noise. Harmonic caching reduces this noise substantially, producing smoother and more accurate reconstructions under equal-time comparisons and achieving 1–3 orders of magnitude lower error than WoSt.

¹Our implementation is available at https://github.com/rohan-sawhney/zombie.



(d) Screened Teapot (α = 20): mixed Dirichlet and Robin BC.

Fig. 9. Comparison of variance-reduction techniques for solving a Laplace equation across diverse 2D domains and boundary conditions under equal-time constraints. Dirichlet boundaries are shown in black, Neumann in green, and Robin in violet. For each result, we show the solution (top) and a zoomed-in error region (bottom), using the inferno and viridis colormaps, respectively. Relative MSE values (scaled by 100) are reported with 1% outlier rejection for both the full image and the inset. Reference solutions are computed with WoSt using 6M walks per evaluation point. Harmonic caching consistently achieves lower error, strongly suppressing Monte Carlo noise and correlation artifacts compared to prior approaches.





Fig. 10. We compare the convergence of WoSt, RWoS, BVC, MVC, and harmonic caching on the 2D scenes (α = 0) from Fig. 9. Each algorithm runs for about 30 minutes, progressively improving as more samples are taken, with relative MSE reported at each iteration. Unlike Fig. 9, which compares performance at equal time, this plot shows the full convergence history. HC consistently achieves the lowest error across all test scenes, independent of runtime.

4.5 Comparison with Reverse Walk on Spheres

RWoS launches walks from the boundary and interior sources, splatting contributions to multiple evaluation points at each walk step while estimating the Green's function of the domain [Miller et al. 2024b; Qi et al. 2022]. This reverse formulation is particularly effective for equations with localized delta sources and boundary conditions. However, as shown in Fig. 11, harmonic caching can outperform RWoS even in this favorable setting, since integrating the off-centered Green's function against delta sources admits an analytic solution that avoids Monte Carlo sampling.

Unlike harmonic caching, RWoS is not output-sensitive, as reverse walks are not guided toward regions of interest. In the 3D domain shown in the inset, the star-shaped regions formed by reverse walks often fail to intersect the planar ROI, yielding few effective splats and stronger correlation artifacts. In contrast, although harmonic caching performs walks over the entire domain, cache records are generated only within the ROI.



RWoS also suffers from additional bias under Dirichlet and mixed boundary conditions due to its finite-difference approximation of the Poisson kernel [Qi et al. 2022, Section 5.3]. Moreover, performance degrades in domains dominated by reflecting boundaries, since each reverse walk step requires a ray–intersection test to avoid splatting into disjoint (invisible) parts of the domain within starshaped regions [Sawhney et al. 2023, Fig. 10]. As a result, under equal-time comparisons in Neumann-dominated problems (Fig. 15), RWoS exhibits more severe correlation artifacts than other caching algorithms due to insufficient walks.

4.6 Comparison with Boundary Value Caching

To reduce redundant computation and suppress noise in WoSt, BVC [Miller et al. 2023] first estimates solution values u and normal derivatives $\partial u/\partial n$ at random boundary points using WoSt. These




Fig. 11. Comparison of reverse walk on spheres and harmonic caching for a Poisson equation with sparse source terms. The scene is intentionally designed to favor RWoS: the source consists of three Dirac delta spikes with zero Dirichlet boundary conditions. Unlike WoS, RWoS estimates the solution by launching walks directly from the source locations, yet our method still achieves lower error at equal time. Because the Green's function integral over a delta function is exact, HC incurs no error in source integration here—only the usual error from constructing the Fourier representation of the unknown boundary solution remains.

cached values are then used to evaluate a boundary integral equation for the PDE solution at interior points via Monte Carlo.

BVC, however, suffers from artifacts caused by singularities in the free-space Green's function and Poisson kernel (Fig. 8). Estimating $\partial u/\partial n$ on Dirichlet boundaries is particularly challenging: Miller et al. [2023] propose offsetting the boundary and using WoSt to estimate both u and $\partial u/\partial n$, but obtaining noise-free gradients this way requires an excessive number of samples. Moreover, because BVC relies on a boundary integral equation with a signed integrand—even when the solution itself is strictly positive—it can reconstruct negative solution values when only a limited number of cache points are used, sometimes yielding higher relative MSE than the baseline WoSt estimator (Figs. 9c and 12).





Fig. 12. For general Poisson problems, error in harmonic caching is often dominated by regions near the source term, due to the off-centered Green's function in the Monte Carlo estimator (Eq. (10)). However, allocating additional source samples is relatively inexpensive since the source term is not recursively defined.

Nevertheless, BVC can outperform our method in practice for Neumann boundary conditions, where $\partial u/\partial n$ is prescribed and need not be estimated. With non-smooth Neumann data, WoSt suffers from high variance near the boundary due to long reflecting walks and signed contributions, which introduce noise into the harmonic coefficients of nearby cache records. While BVC still requires estimating u on Neumann boundaries, its estimates can benefit from splatted contributions across the entire boundary, whereas harmonic caching may lack sufficient overlapping records near the boundary and too few walks per record to estimate coefficients robustly. As a result, harmonic caching requires a large λ and dense cache records, degrading performance. However, as shown in Fig. 9b and Fig. 15 (red zoomed-in crop), our method still achieves the lowest error near smooth Neumann data.

4.7 Comparison with Mean Value Caching

We implemented MVC [Bakbouk and Peers 2023] as a similar *meta-integrator*. In the first pass, WoSt is used to estimate the solution at a set of stratified cache points. At each evaluation point, all cached estimates within the largest ball are then averaged using the volumetric mean value property (Appendix E). MVC suffers from strong correlation artifacts near the boundary (see the cropped insets in Figs. 8 and 9) because ball radii are small and few cache points can be gathered there (inset, top). This remains



true even when employing the recursive reuse passes of Bakbouk and Peers [2023] to share cached information more broadly across the domain (Fig. 13). In contrast, our algorithm is less sensitive to the proximity of evaluation points to the boundary, since distant cache records can still contribute (inset, bottom).

MVC and HC share a key similarity: both rely on forward walks (in contrast to the reverse walks from Sec. 4.5) and employ a more localized caching strategy compared to BVC and RWoS. As a result, when estimates of u(x) inside the domain are very noisy, these methods can become more susceptible to correlation artifacts (Fig. 15).

4.8 Progressive Evaluation

We implemented progressive versions of the caching algorithms and evaluated their convergence in Fig. 10 for the models in Fig. 9. Each progressive algorithm computes a running average over multiple iterations, starting each iteration with an empty cache. For harmonic caching, we keep all parameters except w_{\min} fixed across iterations, which we increase by 0.5 per iteration (starting from 0) to reduce truncation bias from the series expansion. Relative MSE converges at the expected $O(N^{-1})$ rate for WoSt, and harmonic caching consistently achieves the lowest error across the full 30-minute runtime. For the 2D Gear and Teapot setups, BVC and RWoS sometimes incur higher error than WoSt, likely due to stronger correlation artifacts and singularity issues that inflate the reported relative MSE.

5 Conclusion and Future Work

We introduced an efficient, easy-to-implement caching algorithm for WoSt that leverages local harmonic series expansions to reduce variance. The method applies to second-order linear elliptic PDEs with Dirichlet, Neumann, Robin, and mixed boundary conditions. By truncating the expansion at low orders, we obtain a practical reconstruction scheme with few artifacts, and we further show how stochastic truncation via Russian roulette can eliminate the resulting bias. Our approach achieves orders-of-magnitude lower error compared to WoSt, while generally outperforming existing caching algorithms in both accuracy and robustness. In addition to lower error and fewer correlation artifacts, it retains key advantages of prior methods, including progressiveness (all), output sensitivity (MVC/B-VC/ours), singularity-free estimates (MVC/ours), and resilience to domain mesh quality (compared to grid-based PDE solvers).

Limitations and Future Work. In this paper, we adopt WoSt as the underlying pointwise estimator, which can suffer from high variance in problems with large reflecting Neumann boundaries; as shown in Fig. 15, such variance reduces the efficiency of our method. Combining harmonic caching with complementary strategies such as BVC, MVC, RWoS, or path-guided WoSt estimators [Huang et al. 2025] could further reduce error and correlation artifacts. A key limitation of harmonic caching is that it reduces error only within the domain, not directly on the boundary—a restriction shared by all existing caching-based variance-reduction methods for WoSt.




Fig. 13. Reusing the harmonic cache to recompute Fourier coefficients in a refinement pass reduces correlation artifacts and lowers error. This approach resembles the recursive reuse pass in mean value caching but achieves higher accuracy, particularly near boundaries.




Fig. 14. We show results of the 3D Laplace problem from Fig. 1 on three additional slice planes, comparing harmonic caching with state-of-the-art variancereduction algorithms for WoSt under an equal time budget (~1500s) and reporting relative MSE with 1% outlier rejection. The results demonstrate that harmonic caching is competitive, producing smooth reconstructions with the lowest error.

On the theoretical side, our current treatment of Neumann and Robin boundary conditions confines reconstruction to the largest inscribed ball. Extending the reconstruction to general star-shaped regions could significantly reduce the number of required cache records, but poses substantial challenges, namely constructing series expansions for such regions under suitable parametrizations while simultaneously enforcing boundary conditions.

Harmonic caching is effective for Poisson equations with sparse and delta source terms (Fig. 11). However, for problems with smooth, continuously distributed sources, as shown in Fig. 12, MVC can outperform harmonic caching (red inset). This advantage arises because Bakbouk and Peers [2023, Section 8] derived a volumetric mean value property that enables unbiased reuse of source samples. In contrast, our method relies on Monte Carlo integration of the off-centered Green's function without importance sampling, which can lead to higher variance in practice. Extending the harmonic series expansion to incorporate the source term may further reduce variance, which we leave for future work.




Fig. 15. For non-smooth Neumann boundary conditions, harmonic caching can exhibit stronger correlation artifacts than BVC; as shown in the red inset, noisy estimates from the underlying WoSt solver can introduce significant error into the Fourier coefficients. However, our method achieves the lowest error in regions near smooth Neumann boundaries (blue inset).

Finally, while our method applies to the Laplace and screened Poisson equations, extending it to variable-coefficient PDEs [Sawhney et al. 2022] remains an interesting direction for future work.

6 Acknowledgments

The grilled turkey leg mesh in Fig. 1 is courtesy of Quixel Megascans under a standard license, and the oven scene is courtesy of Oven Crystal Black Dct Cr 982 by Franke from Blender Market. The "Engine" model in Fig. 9a is from Zombie [Sawhney and Miller 2023] under the MIT license. The model in Fig. 5 is from Qi et al. [2022], also under the MIT license. All other 2D models are based on images from Adobe Stock under the Education license. The fixed Stanford bunny mesh in Fig. 8 is courtesy of VertexDon from Sketchfab under CC BY-NC-SA 4.0, and the gear mesh in Fig. 15 is courtesy of Ali107 YT from Sketchfab under CC BY 4.0.

We thank the reviewers for their valuable feedback and suggestions. We are grateful to Aaron Lefohn and Ken Museth for supporting this research, and for fruitful discussions with Thomas Booth and Yang Qi. We also thank the Dartmouth Research Computing team for support on the Discovery cluster. Zihong thanks members of the Dartmouth Visual Computing Lab for their support during the project. ChatGPT was used solely to improve grammar and readability, and did not contribute to the technical content or ideas in this paper or supplemental document. This work was partially supported by the National Science Foundation under award numbers 1844538, 2403122, and 2440472.

References

James Richard Arvo. 1993. Transfer Functions in Global Illumination. In Global Illumination (ACM SIGGRAPH Courses). 1–28.

Ghada Bakbouk and Pieter Peers. 2023. Mean Value Caching for Walk on Spheres. In Proceedings of the Eurographics Symposium on Rendering (EGSR), Tobias Ritschel and Andrea Weidlich (Eds.). Eurographics Association. https://doi.org/10/gsk2p7

Mégane Bati, Stéphane Blanco, Christophe Coustet, Vincent Eymet, Vincent Forest, Richard Fournier, Jacques Gautrais, Nicolas Mellado, Mathias Paulin, and Benjamin Piaud. 2023. Coupling Conduction, Convection and Radiative Transfer in a Single Path-Space: Application to Infrared Rendering. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 42, 4 (July 2023), 79:1–79:20. https://doi.org/10/gsk4f5

T.E. Booth. 1982. Regional Monte Carlo Solution of Elliptic Partial Differential Equations. J. Comput. Phys. 47, 2 (1982), 281–290. https://doi.org/10/d3wcch

Thomas E Booth. 1981. Exact Monte Carlo Solution of Elliptic Partial Differential Equations. J. Comput. Phys. 39, 2 (Feb. 1981), 396–404. https://doi.org/10/dchv97 Rustum Choksi. 2022. Partial Differential Equations: A First Course. Number volume 54 in
Pure and Applied Undergraduate Texts. American Mathematical Society, Providence,
Rhode Island.

Michael Czekanski, Benjamin Faber, Margaret Fairborn, Adelle Wright, and David Bindel. 2024. Walking on Spheres and Talking to Neighbors: Variance Reduction for Laplace's Equation. arXiv:physics/2404.17692

Auguste De Lambilly, Gabriel Benedetti, Nour Rizk, Chen Hanqi, Siyuan Huang, Junnan Qiu, David Louapre, Raphael Granier De Cassagnac, and Damien Rohmer. 2023. Heas Simulation on Meshless Crafted-Made Shapes. In ACM SIGGRAPH Conference on Motion, Interaction and Games (MIG '23). ACM Press, 1–7. https://doi.org/10/g57648

Dean G. Duffy. 2015. Green's Functions with Applications (2 ed.). Chapman and Hall/CRC, New York. https://doi.org/10.1201/9781315371412

David P Griesheimer, William R Martin, and James Paul Holloway. 2006. Convergence properties of Monte Carlo functional expansion tallies. J. Comput. Phys. 211, 1 (2006), 129–153. https://doi.org/10/bv5b4h

Tianyu Huang, Jingwang Ling, Shuang Zhao, and Feng Xu. 2025. Guiding-Based Importance Sampling for Walk on Stars. In Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIG-GRAPH Conference Papers '25). Association for Computing Machinery, New York, NY, USA, Article 3, 12 pages. https://doi.org/10/g9x2ms

Pranav Jain, Ziyin Qu, Peter Yichen Chen, and Oded Stein. 2024. Neural Monte Carlo Fluid Simulation. In ACM SIGGRAPH Conference Papers (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, 1–11. https://doi.org/10/gt48wt

Wojciech Jarosz, Craig Donner, Matthias Zwicker, and Henrik Wann Jensen. 2008a. Radiance Caching for Participating Media. ACM Transactions on Graphics 27, 1 (March 2008), 7:1–7:11. https://doi.org/10/cwnw78

Wojciech Jarosz, Volker Schönefeld, Leif Kobbelt, and Henrik Wann Jensen. 2012. Theory, Analysis and Applications of 2D Global Illumination. ACM Transactions on Graphics 31, 5 (Aug. 2012), 125:1–125:21. https://doi.org/10/gbbrkb

Wojciech Jarosz, Matthias Zwicker, and Henrik Wann Jensen. 2008b. Irradiance Gradients in the Presence of Participating Media and Occlusions. Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering) 27, 4 (June 2008), 1087–1096. https://doi.org/10/bg8nww

Jaroslav Křivánek, Pascal Gautron, Kadi Bouatouch, and Sumanta Pattanaik. 2005a. Improved Radiance Gradient Computation. In Proceedings of the Spring Conference on Computer Graphics (SCCG). ACM Press, New York, NY, USA, 155–159. https://doi.org/10/c8pb3t

Jaroslav Křivánek, Pascal Gautron, Sumanta Pattanaik, and Kadi Bouatouch. 2005b. Radiance Caching for Efficient Global Illumination Computation. IEEE Transactions on Visualization and Computer Graphics 11, 5 (2005), 550–561. https://doi.org/10/ csf2sw

Zilu Li, Guandao Yang, Xi Deng, Christopher De Sa, Bharath Hariharan, and Steve Marschner. 2023. Neural Caches for Monte Carlo Partial Differential Equation Solvers. In ACM SIGGRAPH Asia Conference Papers. Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10/gs82vd

Zilu Li, Guandao Yang, Qingqing Zhao, Xi Deng, Leonidas Guibas, Bharath Hariharan, and Gordon Wetzstein. 2024. Neural Control Variates with Automatic Integration. In ACM SIGGRAPH Conference Papers (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, 1–9. https://doi.org/10/gt48wv

Julio Marco, Adrian Jarabo, Wojciech Jarosz, and Diego Gutierrez. 2018. Second-Order Occlusion-Aware Volumetric Radiance Caching. ACM Transactions on Graphics 37, 2 (July 2018), 1–14. https://doi.org/10/gdv86k

Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2023. Boundary Value Caching for Walk on Spheres. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 42, 4 (July 2023), 82:1-82:11. https://doi.org/10/gsk4f8

Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2024a. Differential Walk on Spheres. *ACM Trans. Graph.* 43, 6, Article 174 (Nov. 2024), 18 pages. https://doi.org/10/g8vg2t

Bailey Miller, Rohan Sawhney, Keenan Crane, and Ioannis Gkioulekas. 2024b. Walkin' Robin: Walk on Stars with Robin Boundary Conditions. ACM Trans. Graph. 43, 4 (July 2024), 41:1–41:18. https://doi.org/10/gt48vf

Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and Consistent Rendering Using Biased Estimators. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 41, 4 (July 2022), 48:1–48:13. https://doi.org/10/gqjn66

Mervin E. Muller. 1956. Some Continuous Monte Carlo Methods for the Dirichlet Problem. *Annals of Mathematical Statistics* 27, 3 (Sept. 1956), 569–589. https://doi.org/10/cpxd3d

Ewald Müller and Matthias Steinmetz. 1995. Simulating self-gravitating hydrodynamic flows. Computer Physics Communications 89, 1 (1995), 45–58. https://doi.org/10/ cnbb2s Numerical Methods in Astrophysical Hydrodynamics.

Mohammad Sina Nabizadeh, Ravi Ramamoorthi, and Albert Chern. 2021. Kelvin Transformations for Simulations on Infinite Domains. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 40, 4 (July 2021), 97:1–97:15. https://doi.org/10/gnmwgs

Hong Chul Nam, Julius Berner, and Anima Anandkumar. 2024. Solving Poisson Equations using Neural Walk-on-Spheres. In Proceedings of the 41st International Conference on Machine Learning (ICML'24). JMLR.org, Article 1513, 16 pages. https://dl.acm.org/doi/10.5555/3692070.3693583

Mark Pauly, Thomas Kollig, and Alexander Keller. 2000. Metropolis Light Transport for Participating Media. In Rendering Techniques (Proceedings of the Eurographics Workshop on Rendering). Springer-Verlag, Vienna, 11–22. https://doi.org/10/gfzm93

Matt Pharr and Greg Humphreys. 2010. Physically Based Rendering: From Theory to Implementation (2 ed.), Morgan Kaufmann, San Francisco, USA.

Yang Qi, Dario Seyb, Benedikt Bitterli, and Wojciech Jarosz. 2022. A Bidirectional Formulation for Walk on Spheres. Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering) 41, 4 (2022). https://doi.org/10/jgzr

Ravi Ramamoorthi, John Anderson, Mark Meyer, and Derek Nowrouzezahrai. 2012. A Theory of Monte Carlo Visibility Sampling. ACM Transactions on Graphics 31, 5 (Sept. 2012). https://doi.org/10/gbbrnz

Ravi Ramamoorthi, Dhruv Mahajan, and Peter Belhumeur. 2007. A First-Order Analysis of Lighting, Shading, and Shadows. *ACM Transactions on Graphics* 26, 1 (Jan. 2007), 2:1–2:21. https://doi.org/10/dhkp9c

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 Transactions on Graphics (Proceedings of SIGGRAPH Asia) 41, 6 (Nov. 2022), 240:1–240:16. https://doi.org/10/gscnnz

Rohan Sawhney and Keenan Crane. 2020. Monte Carlo Geometry Processing: A Grid-Free Approach to PDE-based Methods on Volumetric Domains. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 39, 4 (2020). https://doi.org/10/ghw7t7

Rohan Sawhney and Bailey Miller. 2023. Zombie: A Grid-Free Monte Carlo Solver for PDEs.

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 Transactions on Graphics (Proceedings of SIGGRAPH) 42, 4 (July 2023), 80:1–80:20. https://doi.org/10/gsk4f6

Rohan Sawhney, Dario Seyb, Wojciech Jarosz, and Keenan Crane. 2022. Grid-Free Monte Carlo for PDEs with Spatially Varying Coefficients. ACM Transactions on Graphics (Proceedings of SIGGRAPH) 41, 4 (July 2022), 53:1–53:17. https://doi.org/10/gqjn65

Jorge Schwarzhaupt, Henrik Wann Jensen, and Wojciech Jarosz. 2012. Practical Hessianbased Error Control for Irradiance Caching. ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 31, 6 (Nov. 2012), 1. https://doi.org/10/gbb6n4

Gurprit Singh, Cengiz Öztireli, Abdalla G.M. Ahmed, David Coeurjolly, Kartic Subr, Oliver Deussen, Victor Ostromoukhov, Ravi Ramamoorthi, and Wojciech Jarosz. 2019. Analysis of Sample Correlations for Monte Carlo Rendering. Computer Graphics Forum (Proceedings of Eurographics State of the Art Reports) 38, 2 (April 2019). https://doi.org/10/gforzc

Ryusuke Sugimoto, Terry Chen, Yiti Jiang, Christopher Batty, and Toshiya Hachisuka. 2023. A Practical Walk-on-Boundary Method for Boundary Value Problems. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 42, 4 (July 2023), 81:1–81:16. https://doi.org/10/gsk4f7

Gregory J. Ward and Paul S. Heckbert. 1992. Irradiance Gradients. In CE_EGWR93, Alan Chalmers, Derek Paddon, and François X. Sillion (Eds.). Consolidation Express Bristol, Bristol, UK, 85–98.

Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. 1988. A Ray Tracing Solution for Diffuse Interreflection. Computer Graphics (Proceedings of SIGGRAPH) 22, 4 (Aug. 1988), 85–92. https://doi.org/10/dk6rt5

Ekrem Fatih Yilmazer, Delio Vicini, and Wenzel Jakob. 2024. Solving Inverse PDE Problems Using Grid-Free Monte Carlo Estimators. ACM Trans. Graph. 43, 6, Article 175 (Nov. 2024), 18 pages. https://doi.org/10.1145/3687990

Zihan Yu, Lifan Wu, Zhiqian Zhou, and Shuang Zhao. 2024. A Differential Monte Carlo Solver for the Poisson Equation. In ACM SIGGRAPH Conference Papers (SIGGRAPH '24). Association for Computing Machinery, New York, NY, USA, 1–10. https://doi.org/10/gt48ww

A Pseudocode

Pseudocode for harmonic caching is provided in Algorithms 1-3.

Algorithm 1 Harmonic caching algorithm for WoSt

Input: evalPts (evaluation points over ROI) **Output:** updated solution in evalPts

: function HarmonicCachingForPoissonEquation(evalPts)

```
S \leftarrow \text{NULL}
                                              ▶ Initialize an empty cache S
2:
        ▶ Cache Population Pass
3:
        LAZYCACHEUPDATE(evalPts, S, False)
 4:
        ▶ Additional Refinement Pass
5:
6:
        S_{\mathbb{R}} \leftarrow S
                                                   ▶ Deep copy of the cache
        for each Harmonic Cache Record s in S_R do
7:
             s.\widehat{a_0},\cdots,s.\widehat{b_L}\leftarrow 0
                                                           ▶ Reset coefficients
8:
             s.Project(s.p, S)
                                                      ▶ Enable cache lookup
9:
10:
        ▶ Reconstruction Pass
        LAZYCACHEUPDATE (evalPts, S<sub>R</sub>, True)
```

B Harmonic Series Expansion in 3D

Appendix $\mathbb C$ derives the harmonic series expansion of the homogeneous screened Poisson equation in spherical coordinates:

$$u(r,\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} a_l^m \mathcal{R}_{\alpha}^{l,\text{3D}}(r,R) y_l^m(\theta,\phi), \tag{11}$$

where $\mathcal{R}_{\alpha}^{l,\mathrm{3D}}(r,R)$ is the corresponding radial basis function in 3D:

$$\mathcal{R}_{\alpha}^{l,3D}(r,R) = \begin{cases} (r/R)^{l} & \text{if } \alpha = 0, \\ \frac{i_{l}(\sqrt{\alpha}r)}{i_{l}(\sqrt{\alpha}R)} & \text{if } \alpha > 0, \end{cases}$$
(12)

with i_l denoting the modified spherical Bessel function of the first kind. The coefficients a_l^m represent the spherical harmonic decomposition of the boundary function $u(R,\omega)$. They are independent of r and indexed by the spherical harmonics of degree l and order m:

$$a_l^m = \int_{S^2} u(R, \omega) y_l^m(\omega) \, d\omega.$$
 (13)

The spherical harmonics $y_I^m(\theta, \phi)$ are defined as:

$$y_l^m(\theta,\phi) = \begin{cases} K_l^m P_l^{|m|}(\cos\theta)\cos(|m|\phi) & \text{if } m \ge 0, \\ K_l^m P_l^{|m|}(\cos\theta)\sin(|m|\phi) & \text{if } m < 0, \end{cases}$$
(14)

where P_l^m are the associated Legendre polynomials and K_l^m are normalization constants:

$$K_l^m = \sqrt{\frac{(2l+1)}{4\pi} \frac{(l-|m|)!}{(l+|m|)!}}.$$
 (15)

Algorithm 2 Harmonic Cache Record in 2D

1: class HarmonicCacheRecord

22 23

24:

25:

26

28:

```
2: attributes:
  3: p ← 0
                                                                                                          ▶ Cache center
  4: R \leftarrow 0
                                                                          ▶ Closest distance to boundary
  5: \widehat{a_0}, \dots, \widehat{a_L}, \widehat{b_0}, \dots, \widehat{b_L} \leftarrow 0
                                                                                              ▶ Fourier coefficients
       Input: p_0 (cache center), S (cache used for reconstruction)
       Output: Updated cache record, solution estimate at \mathbf{p}_0
  6: function PROJECT(\mathbf{p}_0, S)
  7:
              \mathbf{p} \leftarrow \mathbf{p}_0
               R \leftarrow \text{QueryLargestSphere}(\mathbf{p})
  8:
  9:
              N \leftarrow \lambda R
                                                         ▶ Adaptive number of walks per record
10:
              for i = 1 to N do
                      \Theta_i \sim p^{\Theta_i}
11:
                                                              ▶ Stratified sampling over unit circle
12:
                      \omega \leftarrow (\cos \Theta_i, \sin \Theta_i)
                      \mathbf{p}_1 \leftarrow \mathbf{p}_0 + R \cdot \omega
13:
                      \hat{u} \leftarrow 0
14:
                      if S \neq NULL then
15
                             \hat{u} \leftarrow \text{Reconstruct}(\mathbf{p}_1) \Rightarrow \textit{Used in refinement pass}
16
17:
                      else
                             \hat{u} \leftarrow \text{WalkOnStars}(p_1) \Rightarrow \textit{Used in first population}
18:
                      \widehat{a_0} \leftarrow \widehat{a_0} + \frac{1}{N} \frac{\widehat{u}}{p^{\Theta_i} 2\pi}
19:
                     \begin{aligned} & \mathbf{for} \ l = 1 \ \text{to} \ L \ \mathbf{do} \\ & \widehat{a_l} \leftarrow \widehat{a_l} + \frac{1}{N} \frac{\widehat{u} \cos l\Theta_l}{p^{\Theta_l} 2\pi R^l} \\ & \widehat{b_l} \leftarrow \widehat{b_l} + \frac{1}{N} \frac{\widehat{u} \sin l\Theta_l}{p^{\Theta_l} 2\pi R^l} \end{aligned}
20:
21:
```

⊳ Eq. (7)

▶ Stratified sampling over disk

Input: p_x (evaluation point for reconstruction) **Output:** Solution \hat{u} reconstructed at \mathbf{p}_x with weight w

▶ Account for the source term in standard WoS

 $\widehat{u_0} \leftarrow \widehat{u_0} + \frac{1}{M} \frac{f(y)G^{B(p,R)}(p,y)}{p^{B(p,R)}}$

for i = 1 to M **do**

 $\mathbf{y} \sim p^{B(\mathbf{p},R)}$

```
29: function RECONSTRUCT(\mathbf{p}_x)
               \theta \leftarrow \text{Atan2}(\mathbf{p}_x.y - \mathbf{p}.y, \mathbf{p}_x.x - \mathbf{p}.x)
30:
31:
               r \leftarrow |\mathbf{p}_x - \mathbf{p}|
               d \leftarrow 1 - \frac{r}{R}
32:
               w \leftarrow 3d^2 - 2d^3
                                                                                                    ▶ Weighting kernel
33
               > Reconstruct solution via harmonic series expansion
34:
               \hat{u} \leftarrow \widehat{a_0}
35:
               for l = 1 to L do
36:
                       \hat{u} \leftarrow \hat{u} + (\frac{r}{R})^l (\widehat{a_l} \cos l\theta + \widehat{b_l} \sin l\theta)
                                                                                                                          ⊳ Eq. (2)
37:
               \overline{\mathbf{for}}\ i = 1\ \mathrm{to}\ M\mathbf{do}
38:
                      \mathbf{y} \sim p^{B(\mathbf{p},R)} \qquad \triangleright St
\hat{u} \leftarrow \hat{u} + \frac{1}{M} \frac{f(\mathbf{p}_x) G^{B(\mathbf{p},R)}(\mathbf{p}_x,\mathbf{y})}{p^{B(\mathbf{p})}}
                                                                            ▶ Stratified sampling over disk
39
                                                                                                                        ⊳ Eq. (10)
40:
                                                                                                                          ⊳ Eq. (8)
41:
               return \hat{u}, w
42:
```

Algorithm 3 Populate cache on-the-fly and update solution

Input: evalPts (evaluation points), S (harmonic cache), updateSolution (bool)

Output: Populate S, and update solution if requested

```
function LazyCacheUpdate(evalPts, S, updateSolution)
           for p in evalPts do
2:
                 \hat{u}_{\text{sum}} \leftarrow 0
3:
                 w_{\text{sum}} \leftarrow 0
4:
 5
                 for each HarmonicCacheRecord s in S do
 6:
                       if |s.p - p| \ge 0.9s.R then
                            continue
                                                      \triangleright Skip records outside 0.9R_{\partial\Omega}(\mathbf{p})
 7:
                       ▶ Reconstruct solution from current cache record
8:
9:
                       \hat{u}, w \leftarrow s.\text{Reconstruct}(\mathbf{p})
                       \hat{u}_{\text{sum}} \leftarrow \hat{u}_{\text{sum}} + \hat{u}
10:
                       w_{\text{sum}} \leftarrow w_{\text{sum}} + w
11:
                 if w_{\text{sum}} \leq w_{\text{min}} then
12:
                       HarmonicCacheRecord s \leftarrow NULL
13:
                       \hat{u} \leftarrow s.\text{Project}(\mathbf{p}, \text{NULL})
14:
                       ▶ Collect solution at the cache center
15:
                       \hat{u}_{\text{sum}} \leftarrow \hat{u}_{\text{sum}} + \hat{u}
16:
                       w_{\text{sum}} \leftarrow w_{\text{sum}} + 1
                                                               ▶ weighting kernel at 0 is 1
17:
                       S.ADD(s)
18:
19:
                 if updateSolution then
                       UPDATESOLUTION (p, \frac{\hat{u}_{\text{sum}}}{w_{\text{sum}}})
                                                                                             ▶ Eq. (8)
20:
```

Derivation of the 3D Harmonic Series Expansion

We derive the 3D harmonic series expansion in Eq. (11) for the homogeneous screened Poisson equation using the Fourier method (separation of variables). For the classical case $\alpha = 0$, we refer to Choksi [2022] for a detailed derivation.

The screened Laplace equation in spherical coordinates is:

$$\frac{1}{r^2}\frac{\partial}{\partial r}\left(r^2\frac{\partial u}{\partial r}\right) + \frac{1}{r^2\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial u}{\partial\theta}\right) + \frac{1}{r^2\sin^2\theta}\frac{\partial^2 u}{\partial\phi^2} - \alpha u = 0.$$
(16)

Using separation of variables with $u(r, \theta, \phi) = R(r)\Theta(\theta)\Phi(\phi)$, sub-

$$\frac{1}{\Theta \sin \theta} \frac{\mathrm{d}}{\mathrm{d}\theta} \left(\sin \theta \frac{\mathrm{d}\Theta}{\mathrm{d}\theta} \right) + \frac{1}{\Phi \sin^2 \theta} \frac{\mathrm{d}^2 \Phi}{\mathrm{d}\phi^2} = -\frac{1}{R} \frac{\mathrm{d}}{\mathrm{d}r} \left(r^2 \frac{\mathrm{d}R}{\mathrm{d}r} \right) + \alpha r^2, \tag{17}$$

which holds if both sides equal a separation constant λ . Equating the angular components, we have:

$$\frac{1}{\Theta \sin \theta} \frac{\mathrm{d}}{\mathrm{d}\theta} \left(\sin \theta \frac{\mathrm{d}\Theta}{\mathrm{d}\theta} \right) + \frac{1}{\Phi \sin^2 \theta} \frac{\mathrm{d}^2 \Phi}{\mathrm{d}\phi^2} = -\lambda, \tag{18}$$

which is identical to the Laplace case [Choksi 2022]. The solutions to this equation are:

$$\Phi(\phi) = e^{im\phi} \tag{19}$$

$$\Theta(\theta) = P_I^m(\cos \theta),\tag{20}$$

where P_i^m are the associated Legendre polynomials, indexed by degree l and order m ($|m| \le l$). The separation constant satisfies $\lambda = l(l+1).$

The radial equation from Eq. (17) becomes:

$$\frac{1}{R}\frac{\mathrm{d}}{\mathrm{d}r}\left(r^2\frac{\mathrm{d}R}{\mathrm{d}r}\right) - \alpha r^2 - \lambda = 0. \tag{21}$$

With the substitutions $x = \sqrt{\alpha}r$ and f(x) = R(r), it simplifies to:

$$x^{2}f''(x) + 2xf'(x) - (x^{2} + l(l+1))f(x) = 0,$$
 (22)

which is known as the modified spherical Bessel differential equation. Its solutions are the modified spherical Bessel functions $i_l(x)$ and $k_l(x)$. Requiring finiteness at r=0 excludes $k_l(x)$, leaving $R(r)=i_l(\alpha r)$.

Finally, the general solution of Eq. (16) is given by superposition:

$$u(r,\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} a_l^m i_l(\sqrt{\alpha}r) y_l^m(\theta,\phi), \tag{23}$$

where y_l^m are spherical harmonics defined in Eq. (14), and a_l^m are spherical harmonic coefficients. To restrict the expansion to a ball of radius R, we introduce the radial functions $\mathcal{R}_a^{l,\mathrm{3D}}(r,R)$:

$$u(r,\theta,\phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} a_l^m \mathcal{R}_{\alpha}^{l,\mathrm{3D}}(r,R) \, y_l^m(\theta,\phi). \tag{24}$$

Matching $u(R, \theta, \phi)$ on the sphere with spherical harmonics yields the coefficients a_I^m as given in Eq. (13).

D Green's Functions

For a Poisson equation, the off-centered Green's functions in 2D and 3D for a ball centered at c with radius R are:

$$G^{\text{2D}}(x,y) = \frac{1}{2\pi} \log \left(\frac{R^2 - \langle x - c, y - c \rangle}{Rr} \right), \tag{25}$$

$$G^{3D}(x,y) = \frac{1}{4\pi} \left(\frac{1}{r} - \frac{R}{R^2 - \langle x - c, y - c \rangle} \right), \tag{26}$$

where r := ||y-x|| and $\langle \cdot, \cdot \rangle$ denotes the dot product. For the screened case, see Sawhney et al. [2022, Eq. 5, Supplemental].

E Volume Mean Value Property with Screening

We extend the volume mean value property of Bakbouk and Peers [2023] to handle non-zero screening coefficients $\alpha > 0$, a result we use in Fig. 8. The derivation follows directly from the classical case $\alpha = 0$, and we summarize the results here. In this case, the screened solution satisfies the following mean value property:

$$u(x) = \frac{C}{|B(x)|} \int_{B(x)} u(y) \mathrm{d}y, \tag{27}$$

where |B(x)| denotes the area (in 2D) or volume (in 3D) of a disk or ball of radius R, and C is a normalization factor:

$$C^{\text{2D}} := \frac{R\sqrt{\alpha}}{2I_1(R\sqrt{\alpha})} \text{ and } C^{\text{3D}} := \frac{(R\sqrt{\alpha})^3}{3(R\sqrt{\alpha}\cosh(R\sqrt{\alpha}) - \sinh(R\sqrt{\alpha}))},$$
(28)

with I_1 denoting the modified Bessel function of the first kind.