Stochastic Barnes-Hut Approximation for Fast Summation on the GPU

ABHISHEK MADAN, University of Toronto, Canada NICHOLAS SHARP, NVIDIA, USA FRANCIS WILLIAMS, NVIDIA, USA KEN MUSETH, NVIDIA, USA DAVID I.W. LEVIN, University of Toronto and NVIDIA, Canada



Fig. 1. We compute the electrical potential induced by a power station on a GPU, by taking 2^{22} surface samples and evaluating the potential field on a 1000^2 slice plane. Using one sample per tree subdomain, our stochastic method computes the effect of over 4 trillion particle-particle interactions in 8ms with almost no visible artifacts, while a GPU-optimized Barnes-Hut implementation takes 14ms and brute force takes 2800ms. © Electrical substation model by CG_loma under TurboSquid 3D Model License — Standard.

We present a novel stochastic version of the Barnes-Hut approximation. Regarding the level-of-detail (LOD) family of approximations as control variates, we construct an unbiased estimator of the kernel sum being approximated. Through several examples in graphics applications such as winding

Authors' Contact Information: Abhishek Madan, University of Toronto, Canada, amadan@cs.toronto.edu; Nicholas Sharp, NVIDIA, USA, nsharp@nvidia.com; Francis Williams, NVIDIA, USA, fwilliams@nvidia.com; Ken Museth, NVIDIA, USA, kmuseth@nvidia.com; David I.W. Levin, University of Toronto and NVIDIA, Canada, diwlevin@cs.toronto.edu.

SIGGRAPH Conference Papers '25, Vancouver, BC, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada, https: //doi.org/10.1145/3721238.3730725. number computation and smooth distance evaluation, we demonstrate that our method is well-suited for GPU computation, capable of outperforming a GPU-optimized implementation of the deterministic Barnes-Hut approximation by achieving equal median error in up to 9.4x less time.

CCS Concepts: • Computing methodologies → Shape analysis.

Additional Key Words and Phrases: Monte Carlo methods, Barnes-Hut approximation

ACM Reference Format:

Abhishek Madan, Nicholas Sharp, Francis Williams, Ken Museth, and David I.W. Levin. 2025. Stochastic Barnes-Hut Approximation for Fast Summation on the GPU. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '25), August 10–14, 2025, Vancouver, BC, Canada.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3721238.3730725

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

2 • Madan et al.

1 Introduction

Fast summation lies at the heart of numerous practical algorithms in science, engineering, and computer graphics. For example, *N*-Body problems rely on fast summation schemes to approximate gravitational forces [Barnes and Hut 1986; Greengard and Rokhlin 1987], winding number computation for robust inside-outside testing can be accelerated using fast summation techniques [Barill et al. 2018], and solutions to linear partial differential equations (PDEs) via boundary element-type approaches require fast summation methods to evaluate the solution across the domain [Martinsson 2019].

In computer graphics, stochastic estimation of integrals, which use random sampling to form probabilistic estimates instead of deterministic approximations, have become increasingly popular mathematical and practical tools. Despite the noise artifacts produced by such methods at low sample counts, their ability to handle highly complex problems, and the tendency for the noise to flatten out with more samples, has been used to great effect in rendering and light transport simulation. More recently, these approaches have also been applied effectively to problems in geometry processing and to the solution of linear PDEs. Stochastic approaches have controllable running costs by user-defined sample budgets, and are often embarrassingly parallel, thus mapping well to modern GPU architectures and yielding algorithms which outperform classical approaches at large scales.

Unfortunately, stochastic methods for integration/summation can suffer from poor convergence rates in the number of samples. Control variates (approximations of the integral) are a variance reduction technique that can circumvent slow convergence by giving low-sample regimes more acceptable error [Owen 2013]. Introducing a control variate means that only the residual between the true integrand and the control variate needs to be estimated. The better the approximation provided by the control variate, the smaller this residual, which then allows fewer samples in the estimate for comparable error. However, it can be difficult to design a control variate by hand, limiting their use to problems with computationally efficient analytic approximations (although some neural approaches have recently been proposed to expand their applicability [Li et al. 2024; Müller et al. 2020]).

This work begins with the observation that in many of the new application domains for stochastic methods such as geometry processing, our integral/summation equations are much more structured, and many fast deterministic algorithms exist for their evaluation. Our motivating insight is that these fast, deterministic, approximate methods can serve as robust and accurate control variates for stochastic algorithms.

We apply this observation to computing fast approximate summations. Concretely, we propose an unbiased stochastic estimator for computing spatially varying fields which result from summation of a large number of globally supported kernel functions (e.g., gravitational potentials, winding numbers, smooth distances), using the traditional Barnes-Hut method as a control variate [Barnes and Hut 1986]. Standard Barnes-Hut is difficult to efficiently parallelize on the GPU due to the deep tree traversals required. We show that using a Barnes-Hut traversal as a control variate to a stochastic estimator for the the full summation yields an algorithm that can provide up to an **9.4x** performance improvement for identical median errors when running on an NVIDIA RTX 4090 GPU.

The contributions of this work include a new use of a classical numerical method as a control variate, a proof that the resulting stochastic numerical method is unbiased, and the demonstration of our approach on a number of graphics-centric applications. An open-source implementation is provided on the project website.¹

2 Related Work

Fast summation algorithms are crucial in several fields of science, and as a result, many computational approximation methods have stemmed from these fields. The classic example is N-body simulations: given N particles that each exert some force on all other particles, how does a physical simulation evolve? In computational physics, the Barnes-Hut approximation [Barnes and Hut 1986] was developed to speed up N-body gravitational dynamics simulations, leading to $O(N \log N)$ computational complexity for the entire simulation or $O(\log N)$ to compute the force on a single particle. Concurrently, the fast multipole method (FMM) [Greengard and Rokhlin 1987] was developed to solve similar problems, but instead achieves O(N) computational complexity, albeit at the cost of a significantly more complex preprocessing phase. At their core, these methods rely on similar principles: they both use spatial data structures to organize the particles and multipole approximations to aggregate contributions from points with little impact on the sum. However, the FMM takes it a step further and also computes polynomial approximations for each spatial region that approximate the total effect of distant multipole expansions, though higher-order approximations are required to obtain similar accuracy to Barnes-Hut [Martinsson 2019]. The analytic work in deriving the multipole and local approximations makes it challenging to use the FMM as a general-purpose fast summation accelerator, and so we build upon Barnes-Hut, which can use first-order approximations which do not require such work.

Neither the FMM nor Barnes-Hut are restricted to *N*-body problems — they can both be applied to many problems where a set of source points interacts with a set of query points (e.g., samples from a polyline interacting with a grid of image pixels, see inset), with interaction magnitudes that decay with distance. This generality has allowed



both methods to find several applications in computer graphics as well, including generalized winding number acceleration [Barill et al. 2018], point set surfaces [Alexa et al. 2001], fast linking number computation [Qu and James 2021], repulsive curves and surfaces [Yu et al. 2021a,b], and smooth distances [Madan and Levin 2022]. Our algorithm would be able to accelerate any of these methods as an alternative fast summation technique, and is particularly suitable for enabling or improving GPU implementations of these methods.

Fast summations on the GPU are a classic algorithm that can be easily implemented in an efficient and query-parallel fashion by parallel loads into shared memory [Nyland et al. 2007]. Despite the simplicity of the brute-force GPU algorithm, the fast summation

¹https://www.dgp.toronto.edu/projects/stochastic-barnes-hut/

algorithms that work well on the CPU like Barnes-Hut are more challenging to port to the GPU. The primary difficulty is in effectively parallelizing tree traversals, without introducing excessive thread execution divergence or being limited by memory access latency. Before the advent of modern GPUs, many techniques were designed for general parallel or vector architectures: Barnes [1990] groups particles with similar traversals together to perform a single vectorized traversal per group; Makino [1990] proposes a stackless tree traversal and a breadth-first vectorized tree construction algorithm; and Hernquist [1990] parallelizes a single query by a breadth-first traversal. More recently, Burtscher and Pingali [2011] proposed a complete pipeline for a Barnes-Hut implementation in CUDA using warp-voting traversals, and Bédorf et al. [2012] proposed a similar CUDA pipeline using breadth-first tree traversals. Many similar techniques have also been developed for ray tracing on the GPU, particularly for query-parallel [Wald et al. 2001] and memory-efficient tree traversals [Hapala et al. 2011; Museth 2021; Vaidyanathan et al. 2019], and parallel tree builders [Lauterbach et al. 2009; Ylitie et al. 2017], culminating in software [Parker et al. 2010] and hardware support for ray tracing on modern GPUs. We take inspiration from many of these techniques in our work.

Monte Carlo methods are a simple yet effective means of estimating integrals using samples of the integrand [Metropolis and Ulam 1949], at the cost of introducing some noise (see inset above). They have been used in physically-based rendering to estimate highdimensional integrals that would otherwise be intractable, from the seminal work of Kajiya [1986] to a plethora of techniques available today [Pharr et al. 2023]. More recently, Monte Carlo methods have also been used in other areas of graphics, such as solving volumetric PDEs [Sawhney and Crane 2020; Sugimoto et al. 2023] and fluid simulation [Rioux-Lavoie et al. 2022].

Although most of these applications involve estimating continuous quantities, Monte Carlo methods can also estimate sums using samples of the summand. This is particularly relevant for sampling the contributions of many lights in rendering, where there are a discrete number of light sources in the scene. Many works broadly involve building a tree around the light sources in order to classify their importance [Shirley et al. 1996], and cluster lights together in a manner similar to Barnes-Hut [Walter et al. 2005; Yuksel 2019] or through an online learning process [Wang et al. 2021]. However, these approximations must deal with significant discontinuities due to occlusion, as well as anisotropy that prevents spatial clustering (instead using pre-existing lights as cluster representatives); neither of these are present in our problem setting, so we can perform more effective aggregation to reduce variance, in line with multipole approximations used in Barnes-Hut. The Virtual Point Lights method [Dachsbacher et al. 2014; Keller 1997] does create lights that do not exist in the scene, but these exist along pre-computed light paths, and are primarily intended to unify complex lighting effects within a single framework, so they must be combined with the aforementioned many-light sampling methods to work well in a path tracer.



Fig. 2. An illustration of a tree traversal performed by Barnes-Hut. Spatially (left), the centers of mass $\tilde{\mathbf{p}}_a$ used in the approximation are shown in green, while the source points $\mathbf{p}_i \in S$ are shown in blue. On the tree data structure (right), the contribution nodes $D(\mathbf{q}, \beta)$ are highlighted. The binary tree is for illustrative purposes only and does not correspond to the quadtree on the left.

Departing from previous work in rendering and stochastic PDEs, we show that exploiting the structure of all-pairs summation problems and fusing fast, tree-based approximate methods with stochastic approaches yields a GPU-friendly algorithm that can outperform standard approaches in terms of error and performance in application-relevant regimes.

3 Method

For the rest of the paper, we will operate in the following framework: given a kernel function $f(\mathbf{p}, \mathbf{q})$, a set of source points $S = {\mathbf{p}_1, ..., \mathbf{p}_M}$ with associated masses ${m_1, ..., m_M}$, and a set of query points $Q = {\mathbf{q}_1, ..., \mathbf{q}_N}$, we want to evaluate

$$F(\mathbf{q}) = \sum_{i=1}^{M} m_j f(\mathbf{p}_i, \mathbf{q})$$
(1)

for all $\mathbf{q}_j \in Q$. In typical applications, $f(\mathbf{p}, \mathbf{q})$ is a function that depends on the distance between \mathbf{p} and \mathbf{q} , $\|\mathbf{p} - \mathbf{q}\|$, typically decreasing in magnitude as the distance increases (though there are some exceptions, such as the Green's function of the 2D Laplacian and the Green's function of the 3D biharmonic operator). Although this summation can be computed exactly for all \mathbf{q}_j , it scales linearly with M and N and so becomes impractically slow as both values increase. As a result, approximation methods are used in practice that take advantage of the distance-dependent behavior of f, such as the fast multipole method and the Barnes-Hut approximation. Since our method is an extension of Barnes-Hut, we will briefly review it before describing our work.

3.1 Overview of Barnes-Hut

The Barnes-Hut approximation greatly reduces the cost of evaluating Eq. 1, at the expense of some approximation error. The key insight behind it is that at large distances from a query point **q**, many points in S will have similar contributions to the total $F(\mathbf{q})$, which allows their contributions to be grouped together (Fig. 2). Many groups can be created, which greatly reduce the number of terms in the summation.

To facilitate the creation of such groups, the points in S are placed in a subdivision data structure, such as an octree or a bounding volume hierarchy (Fig. 2, left). Typically, spatial subdivision structures

are preferred for points (i.e., each tree node is characterized by the points it contains), so we will assume that we have an octree-like data structure, possibly with a wider per-dimension branching factor d than the standard 2, and that leaf nodes contain only one source point. We denote tree nodes by T_a and the set of children of T_a by $C(T_a)$ (or equivalently, the set of child node indices); we also use T_a to denote the subset of points (and equivalently, point indices) contained in the node, so $T_a = \bigcup_{T_b \in C(T_a)} T_b$. Each tree node has an associated axis-aligned bounding box containing all its points B_a , with diameter $|B_a|$, and also has an aggregate mass $\tilde{m}_a = \sum_{b \in T_a} m_b$ and a center of mass $\tilde{p}_a = \frac{\sum_{b \in T_a} \tilde{m}_b p_b}{\tilde{m}_a}$. To approximately evaluate $F(\mathbf{q})$, we traverse the tree starting from the root T_0 , and stop if either a *far field condition* $\frac{\|\mathbf{q}-\tilde{\mathbf{p}}_a\|}{|B_a|} \ge \beta$ is met (given a user-defined constant β) or if T_a is a leaf node, and add $\tilde{m}_a f(\tilde{\mathbf{p}}_a, \mathbf{q})$ to the total; otherwise we recursively continue the traversal at each of the node's children. (Higher-order approximations can be used instead [Martinsson 2019], but for ease of presentation we will not discuss them in detail.) Intuitively, the far field condition is a scaleindependent mechanism to check whether a cluster of particles is indistinguishable from a single, larger particle. This tree traversal produces a set of tree nodes $D(\mathbf{q}, \beta)$ that cover all the points in S, called the contribution nodes of the query, and the approximated summation is then

$$F_{BH}(\mathbf{q},\beta) = \sum_{T_a \in D(\mathbf{q},\beta)} \tilde{m}_a f(\tilde{\mathbf{p}}_a,\mathbf{q}).$$
(2)

Pseudocode for Barnes-Hut is provided in the supplemental.

3.2 Path Interpretation of Barnes-Hut

When $\beta = \infty$, the Barnes-Hut algorithm traverses the entire tree and visits every node to obtain the exact answer from Eq. 1. Effectively, for each source point \mathbf{p}_i , we traverse a path $\mathcal{P}_i = [T_{i,0}, \ldots, T_{i,d_i}]$ down the tree, where $T_{i,0} = T_0$ and d_i is the depth of the leaf node containing \mathbf{p}_i . Each tree node along a path represents an approximation of the true value of $m_i f(\mathbf{p}_i, \mathbf{q})$, and traversing to the next node in the path is equivalent to replacing this approximation with a more accurate one. The path traversal process can therefore be represented by a telescoping sum:

$$F_{\mathcal{P}_{i}}(\mathbf{q}) = m_{i}f(\tilde{\mathbf{p}}_{i,0},\mathbf{q}) + \sum_{k=1}^{d_{i}} m_{i}(f(\tilde{\mathbf{p}}_{i,k},\mathbf{q}) - f(\tilde{\mathbf{p}}_{i,k-1},\mathbf{q})), \quad (3)$$

and summing over all *M* paths (one path per source point) gives $F(\mathbf{q}) = \sum_{i=1}^{M} F_{\mathcal{P}_i}(\mathbf{q})$. With smaller values of β , only part of the tree is traversed, and as a result the paths to the leaves are truncated. Denoting a path prefix by $\mathcal{P}_{i,k} = [T_{i,0}, \ldots, T_{i,k}]$ ($k \le d_i$), and the Barnes-Hut-induced truncation length by $\ell_i(\beta)$, Barnes-Hut traverses path prefixes $\mathcal{P}_{i,\ell_i}(\beta)$, and we then have

$$F_{BH}(\mathbf{q},\beta) = \sum_{i} F_{\mathcal{P}_{i,\ell_i(\beta)}}(\mathbf{q}).$$
(4)

Many paths share common prefixes, and will thus be identical after truncation, after which they can be aggregated into a single path with the combined mass of all the constituent paths. Expanding the summands in Eq. 4 and grouping common prefixes gives Eq. 2, and





Fig. 3. In the near field, $\tilde{\mathbf{p}}_{i,k+1}$ is much closer to \mathbf{q}_1 than $\tilde{\mathbf{p}}_{i,k}$, so $R_i(\mathbf{q}_1, k) > 1$, making it an invalid probability (left). However, in the far field, the distance from \mathbf{q}_2 to $\tilde{\mathbf{p}}_{i,k}$ is similar to its distance to $\tilde{\mathbf{p}}_{i,k+1}$, so $R_i(\mathbf{q}_2, k) \approx \frac{1}{d}$ (i.e., the reciprocal of the per-dimension branching factor), which is useful convergent behavior for Russian roulette probabilities (right).

thus path truncation and aggregation is equivalent to the typical truncated tree traversal perspective of Barnes-Hut.

3.3 Path-Based Monte Carlo Estimator

At first glance, there is no computational benefit to the path interpretation of Barnes-Hut - computing telescoping sums is strictly more work than simply canceling terms and evaluating f at the contribution nodes. However, it opens up stochastic methods to estimating Eq. 1. The telescoping sum in Eq. 3, for example, resembles an estimator that uses an infinite series of biased estimators to form an unbiased estimator [Misso et al. 2022; Rhee and Glynn 2015]. Numerical approximations are in a sense biased estimators of the exact quantity - although Barnes-Hut approximations from Eq. 2 are deterministic, we can view them in a more general probabilistic sense, where it is clear that by construction their expected value does not match Eq. 1. Another way of viewing this is in terms of control variates, where the terms of the telescoping sum represent replacing a less accurate control variate for the estimator with a more accurate one. Monte Carlo methods are most commonly deployed in settings where deterministic approximations are challenging to construct, such as in physically-based rendering, but we flip this perspective and instead use deterministic approximations as control variates to bootstrap a Monte Carlo estimator. Not only does this produce an unbiased estimate of the exact quantity we wish to evaluate, but it also provides computational benefits on the GPU, as we will later demonstrate.

There are two primary random variables that we need to sample: the path index *I* (equivalent to a source point index), and the path length *K* (which depends on *I*). For simplicity, we uniformly sample the path index among source point indices, and leave importance sampling schemes (e.g., based on distance to **q**) as future work. The path length is sampled using a Russian roulette process, where at a node at depth *k* along the path \mathcal{P}_i , we decide to go to the next node with a probability $p_{i,k}(\mathbf{q})$. In general, any non-zero probability will allow for an unbiased estimate, but in practice they must be carefully selected to keep the variance low. To sample path lengths, we take further inspiration from Barnes-Hut and use the *far field ratio* $\tilde{\beta}_i(\mathbf{q}, k) = \frac{\|\mathbf{q} - \tilde{p}_{i,k}\|}{|B_{i,k}||}$ of a node $T_{i,k}$ along \mathcal{P}_i as a key building block of the Russian roulette probabilities. In Barnes-Hut, this ratio is compared with β in the far field condition to determine whether or not a tree traversal continues to the node's children. In the context of our algorithm, however, want to convert this deterministic decision into a stochastic decision that produces similar behavior on average (i.e., short paths in the far field, long paths in the near field). To achieve this, we use a ratio of far field ratios $R_i(\mathbf{q}, k) = \frac{\tilde{\beta}_i(\mathbf{q}, k)}{\tilde{\beta}_i(\mathbf{q}, k+1)}$ between the current and next node in the path to determine the Russian roulette probability.

To understand the significance of R_i , consider a query point \mathbf{q} that is far away from the source point \mathbf{p}_i . In an octree-like structure, the diameter of a child node $|B_{i,k+1}|$ is d times less than the diameter of its parent $|B_{i,k}|$, where d is the per-dimension branching factor. Also, as we move further down the path \mathcal{P}_i , the distance from \mathbf{q} to the parent's center of mass is similar to the distance from \mathbf{q} to the child's center of mass (i.e., $\|\mathbf{q} - \tilde{\mathbf{p}}_{i,k}\|$ and $\|\mathbf{q} - \tilde{\mathbf{p}}_{i,k+1}\|$ move closer together). Therefore, for long paths through an octree-like structure, $R_i(\mathbf{q}, k) \rightarrow \frac{1}{d}$ as k grows (Fig. 3, right). Thus, $R_i(\mathbf{q}, k)$ is a reasonable Russian roulette probability that disincentivizes deep paths in the far field without completely prohibiting them.

The near field (i.e., source points near **q**) requires a bit more care, since $\tilde{\beta}_i(\mathbf{q}, k)$ may be quite small even at deep levels, and $R_i(\mathbf{q}, k)$ may be greater than 1 at shallow levels if $\mathbf{q} \in B_{i,k}$, making it an invalid probability (Fig. 3, left). To solve these problems, we clamp the parent's far field ratio to be at least 1, to ensure that **q** is outside of $B_{i,k}$ and $\tilde{\beta}_i(\mathbf{q}, k)$ exhibits far field behavior before we use it in R_i , and we clamp the ratio to be at most 1 to ensure it is a valid probability. (We could use R_i to simultaneously define a Russian roulette probability and a splitting factor when $R_i > 1$ [Szirmay-Kalos et al. 2005], but due to the difficulty of efficiently implementing path splitting on the GPU, we opted to exclusively use Russian roulette.) The final Russian roulette probability is then

$$p_{i,k}(\mathbf{q}) = \min\left(1, \frac{\max(1, \tilde{\beta}_i(\mathbf{q}, k))}{\tilde{\beta}_i(\mathbf{q}, k+1)}\right).$$
(5)

In the supplemental material, we compare our Russian roulette probabilities with fixed probabilities of $p_{i,k} = \frac{1}{2}$ and completely disabling Russian roulette (i.e., $p_{i,k} = 1$).

After deciding to traverse to the next node in the path, we must include the corresponding term of the telescoping series in our running estimate. Directly using Eq. 3 is inefficient because it only accounts for the contribution of a single \mathbf{p}_i to Eq. 1, and has high variance because a node's center of mass can vary quite significantly from any individual source point in the node. The tree structure allows a solution to both of these problems: it makes it possible to aggregate several identical telescoping terms from different $F_{\mathcal{P}_i}$ together in a single estimator term, and enables *contribution swaps* between a parent node and the sum of all of its children:

$$\Delta_{i,k} = \left(\sum_{c \in C(T_{i,k})} \tilde{m}_c f(\tilde{\mathbf{p}}_c, \mathbf{q})\right) - \tilde{m}_{i,k} f(\tilde{\mathbf{p}}_{i,k}, \mathbf{q}).$$
(6)



Fig. 4. The first two contribution swaps along a path \mathcal{P}_i (bold). The parent node (green) subtracts its contribution from the sum of its child contributions (purple).

Essentially, this swap is a form of antithetic sampling, because the contribution of the child node along the current path is counterbalanced by the contributions of its sibling nodes. Some example contribution swaps are illustrated in Fig. 4.

The final estimator for a single truncated path sample (I, K) is

$$\hat{F}_{1}(\mathbf{q}) = \tilde{m}_{I,0} f(\tilde{\mathbf{p}}_{I,0}, \mathbf{q}) + \sum_{k=1}^{K} \frac{\Delta_{I,k-1}}{p(I \in T_{I,k-1})p(K \ge k)},$$
(7)

where $p(I \in T_{i,k}) = \sum_{j \in T_{i,k}} p(I = j)$, and $p(K \ge k) = \sum_{\ell=k}^{d_i} p(K = \ell) = \prod_{j=0}^{k-1} p_{i,j}(\mathbf{q})$. The *S*-sample estimator is $\hat{F}_S(\mathbf{q}) = \frac{1}{S} \sum_{i=1}^{S} \hat{F}_1(\mathbf{q})$.

THEOREM 3.1. Eq. 7 is an unbiased estimator of Eq. 1; that is, $E\left[\hat{F}_1(\mathbf{q})\right] = F(\mathbf{q}).$

The proof of Theorem 3.1 is in the supplemental material.

3.3.1 Domain Stratification. Stratified sampling is an important variance reduction technique in Monte Carlo estimators that reduces the effect of samples clumping in similar regions. Analogously, stratified sampling of different paths helps us achieve good coverage of the terms in Eq. 1. Spatial data structures already provide a convenient means of domain stratification with the different levels of the tree that represent nested subdomains (e.g., Fig. 2, left), so we can easily achieve stratified samples by modifying the path roots. In other words, rather than starting all paths from the tree root T_0 , we can instead define a set of descendant nodes, or subdomains, as path roots (as long as they cover all of S), and sample paths starting from these nodes. For simplicity, we simply use the direct children of the root T_0 as the subdomains, but other choices are also possible, such as the Barnes-Hut contribution nodes $D(\mathbf{q}, \beta)$ for a low β , though we leave the GPU-efficient implementation of such schemes as future work.

Algorithm pseudocode is provided in the supplemental.

4 Results

4.1 Implementation Details

We implemented our method and Barnes-Hut using C++ and CUDA, using the Eigen linear algebra library [Guennebaud et al. 2010]. To our knowledge, there is no open-source, general-purpose GPU

6 • Madan et al.



Fig. 5. Convergence of our method vs Barnes-Hut, for 1 million random and grid query points, and 2^{15} and 2^{20} source points sampled from the Stanford bunny. Each point on the curves for our method (green) correspond to the number of samples per subdomain *S* from 1 to 32; each point on the curves for Barnes-Hut (orange) correspond to the accuracy parameter β from 1 to 40. The *S* = 1 point on each curve for our method is outlined in black, and the point on the Barnes-Hut curve that achieves the same error is also marked (found by linear interpolation on the log-log plot). Across all source point sets and query distributions, and for both mean and median error, our method is roughly 2-9 times faster than Barnes-Hut. The Barnes-Hut curves turn upwards for large β due to floating-point error for computing sums with many terms in single-precision.

implementation of Barnes-Hut, so we opted to implement and optimize our own baseline. Both methods use a common octree-like spatial data structure, that supports variable branching factors. The tree builder is currently CPU-only, and we leave incorporating fast GPU tree builders (e.g., the IndexGrid builder from fVDB [Williams et al. 2024]) as future work. We implemented Barnes-Hut with a stackless BVH traversal [Hapala et al. 2011; Makino 1990] and warp voting [Barnes 1990; Burtscher and Pingali 2011] to simultaneously reduce thread divergence and improve accuracy, using a tree with a per-dimension branching factor of d = 2 (i.e., a total branching factor of 8). The implementation of our method maintains two separate RNG states for drawing path index samples and Russian roulette samples independently, and seeds them both identically across warps to reduce thread divergence. Our method also uses a tree with a wider per-dimension branching factor of d = 4 (i.e., a total branching factor of 64) compared to Barnes-Hut, as we found that a narrow branching factor for Barnes-Hut allowed the algorithm to make efficiently make fine-grained far field decisions, while our algorithm performs better with wider trees that balance path depth with the cost of scanning over all of a node's children for contribution swaps. To maximize performance, all operations are implemented using single-precision floating point. Performance testing was done using a Dell Precision Workstation equipped with an AMD Threadripper CPU (Model 7955WX), 64 GB of RAM, and an NVIDIA RTX 4090 GPU with 24 GB of memory.

All geometric data used in the rest of this section is isotropically rescaled to fit in a $[-1, 1]^3$ grid.

4.2 Comparison to Barnes-Hut

To start, we used a Coulomb-like potential $f(\mathbf{p}, \mathbf{q}) = -\frac{1}{\|\mathbf{p}-\mathbf{q}\|}$ to investigate the performance of our algorithm compared to Barnes-Hut. We took 2¹⁵ (32,768) and 2²⁰ (1,048,576) samples on the Stanford bunny, and evaluated both our method and Barnes-Hut on a grid of 100³ points in $[-1, 1]^3$, as well as 1,000,000 randomly distributed points in the same cube for a variety of parameter values. In Barnes-Hut, the main parameter is the far field threshold β , and our method provides the number of samples per subdomain (spsd), denoted *S*. In Fig. 5, we plotted the mean and median absolute error versus execution time for both source sizes and both query distributions, using parameter sweeps $\beta \in [1, 2, ..., 40]$ and $S \in [1, 2, ..., 32]$.

Across all configurations, Barnes-Hut can only match the mean and median error achieved by our method with S = 1 after 2-8 times more execution time. Furthermore, our relative speedup is faster on random query sets, where Barnes-Hut is



slowed down by significant thread divergence (though warp voting allows it to use that time to drive down the overall error). The relative lack of sensitivity of our method to the query set's spatial coherence is a remarkable property that is achieved through probabilistic path traversals, which change far less across query points than overall tree traversals. Our method also achieves higher relative speedups under median error compared to mean error, but this suggests that much of the error in our method comes from isolated outliers rather than spread out across all queries. In rendering, this is caused by "fireflies" due to importance sampling schemes giving high weight to some outlier samples that are far from the true expected value, so improving our sampling method and even using a specially designed denoiser for these noise patterns [Vicini et al. 2019] are potential ways to address this. Another aspect of our method is that we converge at a slower than Barnes-Hut — although we have a performance advantage at lower sample counts, this advantage disappears as *S* increases because the error converges at a rate the Monte Carlo rate of $O(S^{-1/2})$ (though since we measure the mean and median absolute error rather than the root mean squared error, this rate is not precise in our case, see inset). Our method is therefore most useful with low sample counts, though regardless the low-sample regime is the main regime of interest, since we are estimating a quantity that can be evaluated in finite time.

We also tested our method on a variety of source distributions, using recommended parameter configurations for both Barnes-Hut and our method. Papers in the Barnes-Hut literature in graphics recommend using $\beta = 2$ [Barill et al. 2018; Madan and Levin 2022; Qu and James 2021], as this usually achieves a good tradeoff between efficient computation and acceptable error in graphics applications, so we make the same choice. As for our method, approximating a discrete summation puts a hard limit on the sample count, as there is a sample threshold where our method becomes more expensive than a full evaluation, and as seen in Fig. 5, our absolute error converges slowly; therefore, we pick S = 1 as a reasonable default. We evaluated our method with these choices across a mesh dataset from Myles et al. [2014], where we drew 2¹⁵, 2¹⁷, and 2²⁰ uniformly distributed samples from each surface to use as kernel source points (Table 1). Although our performance is comparable with Barnes-Hut for a small number of sources (2^{15}) , our method scales much better, becoming 2x faster at the chosen settings with 2²⁰ source samples, while maintaining lower mean and median error across all source sizes M. Our maximum error is always an order of magnitude larger than Barnes-Hut, but the significantly lower median error suggests that these points are outliers.

A useful theoretical property of our approach is that it is an unbiased estimator of the true summation value, as opposed to a deterministic but biased estimate of the quantity. Fig. 6 looks at the same slice as Fig. 1 and examines the log error of both our method and Barnes-Hut. Barnes-Hut error has ring-like artifacts that indicate where the far field condition caused the set of contribution nodes $D(\mathbf{q}, \beta)$ to change in space; our method, meanwhile, does not have such artifacts, but error is more highly concentrated in regions with larger gradients. Monte Carlo noise is also visible in our error.

4.3 Applications

We applied our method to a variety of different kernel functions, and compared the visual result quality from our method to the ground truth results computed with an exact GPU-optimized algorithm [Nyland et al. 2007]. In all examples, slice planes are evaluated as 1000² grids for 1 million query points. The identically-seeded RNG states across warps cause visually displeasing spatial coherence, so we shuffle the grid positions before passing them to our algorithm. Although this induces a slight performance hit due to additional thread divergence, it is much smaller than a comparable hit to performance would be to Barnes-Hut (Fig. 5), since we retain most of the benefits of identical RNG states.

Stochastic Barnes-Hut Approximation for Fast Summation on the GPU • 7



Fig. 6. In the example from Fig. 1, Barnes-Hut (left) exhibits discontinuous error patterns when the far field condition causes a change in contribution nodes, while our method (right) does not exhibit such artifacts and instead is more concentrated in regions where the field rapidly changes, while being lower on average than Barnes-Hut.



1 Sample/Subdomain 16 Samples/Subdomain Ground Truth

Fig. 7. Winding numbers computed via fast stochastic summation from 2^{20} sources. A full evaluation via brute force takes 887ms, while our method takes 7.80ms with 1 sample per subdomain, and 119ms with 16 samples per subdomain.

4.3.1 Coulomb Potentials. In Fig. 1, we show a visualization of an electric (Coulomb) potential field induced by 2^{22} (over 4 million) points on the surface of an electrical substation, with some points given a higher mass (i.e., charge) in the cables above the main units in the center. With just one sample per subdomain, we produce a smooth result on all evaluation slices in roughly 8 milliseconds for the large slice plane and 6 milliseconds for the smaller planes.

4.3.2 Winding Numbers. In Fig. 7, we compute winding numbers [Jacobson et al. 2013] using our method from 2^{20} sources, with 1 sample per subdomain and 16 samples per subdomain, as well as the ground truth. Our method mostly succeeds in classifying inside and outside cells with S = 1, but contains noise near the boundary where the winding number changes rapidly; using 16 samples per subdomain mostly alleviates this issue at a significant performance cost.

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

Table 1. Timings and errors of our method and Barnes-Hut on a dataset of 116 meshes, run with typical parameter settings. Barnes-Hut is evaluated with $\beta = 2$, and our algorithm is evaluated with S = 1 (i.e., one sample per subdomain). We vary the number of source samples *M* to be 2^{15} (32,768), 2^{17} (131,072), and 2^{20} (1,048,576), draw them from each mesh surface in the dataset, and their total gravitational potential kernel on a grid of 100^3 points. The mean and standard deviation for each error statistic across the dataset are reported, and the better metric is highlighted in bold at each source set size. At typical algorithm parameter settings, we match our outperform an optimized Barnes-Hut implementation by over 2x, while always achieving roughly 5x lower mean error and 17x lower median error.

Algorithm (Typical Config.)	Time (ms)	Mean Abs. Err.	Median Abs. Err	Max Abs. Err.
Barnes-Hut $(\beta = 2, M = 2^{15})$	$\textbf{3.48} \pm \textbf{0.93}$	$5.59e-04 \pm 2.83e-04$	$4.50e-04 \pm 2.37e-04$	$3.72e03 \pm 1.77e03$
Ours $(S = 1, M = 2^{15})$	3.63 ± 0.97	$\textbf{9.91e-05} \pm \textbf{2.28e-05}$	$2.51e05 \pm 6.73e06$	$5.61e-02 \pm 3.05e-02$
Barnes-Hut $(\beta = 2, M = 2^{17})$	5.00 ± 1.46	$5.60e-04 \pm 2.82e-04$	$4.50e-04 \pm 2.36e-04$	$3.73\text{e-}03 \pm 1.75\text{e-}03$
Ours $(S = 1, M = 2^{17})$	$\textbf{3.92} \pm \textbf{1.04}$	$9.81e\text{-}05 \pm 2.21e\text{-}05$	$\textbf{2.43e-05} \pm \textbf{6.22e-06}$	$5.62e-02 \pm 3.23e-02$
Barnes-Hut $(\beta = 2, M = 2^{20})$	8.61 ± 2.89	$5.60e-04 \pm 2.82e-04$	$4.50e-04 \pm 2.36e-04$	$3.74\text{e-}03 \pm 1.76\text{e-}03$
Ours $(S = 1, M = 2^{20})$	$\textbf{4.04} \pm \textbf{1.12}$	$9.77e05 \pm 2.16e05$	$2.40e05 \pm 5.66e06$	$5.60e-02 \pm 3.13e-02$



1 Sample/Subdomain 64 Samples/Subdomain Ground Truth

Fig. 8. Smooth distances computed via fast stochastic summation from 2^{20} source points. A full evaluation via brute force takes 710ms, and our method takes 7.10ms with 1 sample per subdomain 605ms with 64 samples per subdomain.

4.3.3 Smooth Minimum Distances. In Fig. 8, we compute smooth distance fields [Madan and Levin 2022] from 2^{20} sources, using our method with 1 sample per subdomain and 64 samples per subdomain, and the ground truth. More precisely, we estimate the sum of exponentials and apply a logarithm to the estimate as one would do for the full summation; however, unlike previous examples, the result is also biased for our method, because the expectation operator cannot be directly "passed through" the logarithm. Unlike Madan and Levin [2022], we do not use off-center bounding boxes to guarantee the field is a conservative estimate of the true distance. At S = 1, our method has difficulty resolving the band near the zero level set, but can more effectively resolve it at S = 64, though this is almost as expensive as a full evaluation.

5 Conclusions and Future Work

We presented a method that stochastically estimates kernel sums using a novel unbiased algorithm using a Barnes-Hut-esque hierarchy of control variates. Our method already achieves excellent results on Coulomb potential evaluation problems, but there is significant room to improve the employed sampling techniques to achieve comparable results on other problems such as winding number and smooth distance evaluation. We only use simple uniform sampling for paths, so there is an immediate opportunity to improve the sampled paths with importance sampling (though this could create thread divergence without care). The Russian roulette probabilities also deserve further investigation using more principled approaches such as zero-variance theory [Vorba and Křivánek 2016] and estimator efficiency optimization [Rath et al. 2022] - since there are a finite number of paths, it may be possible to use the aforementioned techniques to compute all probabilities for a given query point. A unified Russian roulette and splitting framework on the GPU is another related direction of future work. Recently, the ReSTIR algorithm [Wyman et al. 2023] has revolutionized realtime path tracing by reusing samples across pixels in space and time; integrating such an approach into our framework is another compelling direction. Beyond discrete summations, applying this method to continuous integrals could be fruitful, potentially leading to an alternative to boundary element methods in the same manner Walk on Spheres provides an alternative to finite elements [Sawhney and Crane 2020].

Our method uses more compute than traditional tree-based approximations like Barnes-Hut, but as a result, it reduces thread divergence in a way that is relatively insensitive to query order -aproperty that, to our knowledge, is unique to our method, providing it with unique advantages on the GPU. Furthermore, stackless tree traversals require more memory traffic than their recursive counterparts to access nodes, so being able to perform simpler path traversals means that we access fewer tree nodes overall. However, part of our computation involves a memory scan over an uncompressed buffer of node-aggregated data such as centers of mass and aggregated masses, which still increases memory traffic in conjunction with compute and can cause scaling issues at higher sample counts, though it is possible to reduce this through aggressive quantization and compression, as long as leaf data remains uncompressed to keep the estimator unbiased. Compressing the tree nodes would also help Barnes-Hut, and further pursuing optimizations of both methods would change the performance tradeoff between the two. More radical changes such as completely different data structures could also enable the use of "single-term" estimators [McLeish 2011; Misso et al. 2022] that require cheap random access to internal tree nodes (as opposed to following child pointers in typical tree data structures), and would completely eliminate thread divergence on the GPU. Overall, following the trend of other recent Monte Carlo methods, we believe there is a lot of potential in stochastic methods to improve upon deterministic methods, or at the very least, provide a dialogue between the two classes of methods to offer different choices for users depending on their needs, making it an exciting time to explore the use of stochasticity in traditionally deterministic computational domains.

Acknowledgments

This project was funded in part by an NSERC Discovery Grant (RGPIN-2023-05120) and an Ontario Early Researchers Award. The first author was funded by an NSERC Canada Graduate Scholarship – Doctoral. We thank the attendees of the GraphQUON workshop for useful early discussions, and Mark Harris and Lars Nyland for proofreading.

References

- Marc Alexa, Johannes Behr, Daniel Cohen-Or, Shachar Fleishman, David Levin, and Claudio T Silva. 2001. Point set surfaces. In *Proceedings Visualization, 2001. VIS'01.* IEEE, 21–29.
- Gavin Barill, Neil G Dickson, Ryan Schmidt, David IW Levin, and Alec Jacobson. 2018. Fast winding numbers for soups and clouds. ACM Transactions on Graphics (TOG) 37, 4 (2018), 1–12.
- Josh Barnes and Piet Hut. 1986. A hierarchical O (N log N) force-calculation algorithm. *nature* 324, 6096 (1986), 446–449.
- Joshua E Barnes. 1990. A modified tree code: Don't laugh; it runs. J. Comput. Phys. 87, 1 (1990), 161–170.
- Jeroen Bédorf, Evghenii Gaburov, and Simon Portegies Zwart. 2012. A sparse octree gravitational N-body code that runs entirely on the GPU processor. J. Comput. Phys. 231, 7 (2012), 2825–2839.
- Martin Burtscher and Keshav Pingali. 2011. An efficient CUDA implementation of the tree-based barnes hut n-body algorithm. In GPU computing Gems Emerald edition. Elsevier, 75–92.
- Carsten Dachsbacher, Jaroslav Křivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable realistic rendering with many-light methods. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 88–104.
- Leslie Greengard and Vladimir Rokhlin. 1987. A fast algorithm for particle simulations. Journal of computational physics 73, 2 (1987), 325–348.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. http://eigen.tuxfamily.org.
- Michal Hapala, Tomáš Davidovič, Ingo Wald, Vlastimil Havran, and Philipp Slusallek. 2011. Efficient stack-less bvh traversal for ray tracing. In Proceedings of the 27th Spring Conference on Computer Graphics. 7–12.
- Lars Hernquist. 1990. Vectorization of tree traversals. J. Comput. Phys. 87, 1 (1990), 137–147.
- Alec Jacobson, Ladislav Kavan, and Olga Sorkine. 2013. Robust Inside-Outside Segmentation using Generalized Winding Numbers. ACM Trans. Graph. 32, 4 (2013).
- James T Kajiya. 1986. The rendering equation. In Proceedings of the 13th annual conference on Computer graphics and interactive techniques. 143–150.
- Alexander Keller. 1997. Instant radiosity. In Proceedings of the 24th annual conference on Computer graphics and interactive techniques. 49–56.
- Christian Lauterbach, Michael Garland, Shubhabrata Sengupta, David Luebke, and Dinesh Manocha. 2009. Fast BVH construction on GPUs. In *Computer Graphics Forum*, Vol. 28. Wiley Online Library, 375–384.
- 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 2024 Conference Papers. 1–9.
- Abhishek Madan and David I.W. Levin. 2022. Fast Evaluation of Smooth Distance Constraints on Co-Dimensional Geometry. ACM Trans. Graph. 41, 4 (2022).
- Junichiro Makino. 1990. Vectorization of a treecode. J. Comput. Phys. 87, 1 (1990), 148-160.
- Per-Gunnar Martinsson. 2019. Fast direct solvers for elliptic PDEs. SIAM.
- Don McLeish. 2011. A general method for debiasing a Monte Carlo estimator. Monte Carlo Methods and Applications 17, 4 (2011), 301–315. doi:doi:10.1515/mcma.2011.013

- Nicholas Metropolis and Stanislaw Ulam. 1949. The monte carlo method. Journal of the American statistical association 44, 247 (1949), 335-341.
- Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and consistent rendering using biased estimators. ACM Transactions on Graphics (TOG) 41, 4 (2022), 1–13.
- Thomas Müller, Fabrice Rousselle, Alexander Keller, and Jan Novák. 2020. Neural control variates. ACM Transactions on Graphics (TOG) 39, 6 (2020), 1–19.
- Ken Museth. 2021. NanoVDB: A GPU-Friendly and Portable VDB Data Structure For Real-Time Rendering And Simulation. In ACM SIGGRAPH 2021 Talks (Virtual Event, USA) (SIGGRAPH '21). Association for Computing Machinery, New York, NY, USA, Article 1, 2 pages. doi:10.1145/3450623.3464653
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust Field-aligned Global Parametrization. ACM Trans. on Graphics - Siggraph 2014 33, 4 (2014), Article No. 135. http://vcg.isti.cnr.it/Publications/2014/MPZ14
- Lars Nyland, Mark Harris, and Jan Prins. 2007. Fast n-body simulation with cuda. GPU gems 3 (2007), 62–66.
- Art B. Owen. 2013. Monte Carlo theory, methods and examples. https://artowen.su. domains/mc/.
- Steven G Parker, James Bigler, Andreas Dietrich, Heiko Friedrich, Jared Hoberock, David Luebke, David McAllister, Morgan McGuire, Keith Morley, Austin Robison, et al. 2010. Optix: a general purpose ray tracing engine. Acm transactions on graphics (tog) 29, 4 (2010), 1–13.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. Physically based rendering: From theory to implementation. MIT Press.
- Ante Qu and Doug L James. 2021. Fast linking numbers for topology verification of loopy structures. ACM Trans. Graph. 40, 4 (2021), 106–1.
- Alexander Rath, Pascal Grittmann, Sebastian Herholz, Philippe Weier, and Philipp Slusallek. 2022. EARS: Efficiency-Aware Russian Roulette and Splitting. ACM Transactions on Graphics (Proceedings of SIGGRAPH 2022) 41, 4, Article 81 (jul 2022), 14 pages. doi:10.1145/3528223.3530168
- Chang-han Rhee and Peter W Glynn. 2015. Unbiased estimation with square root convergence for SDE models. Operations Research 63, 5 (2015), 1026–1043.
- 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 (TOG) 41, 6 (2022), 1–16.
- Rohan Sawhney and Keenan Crane. 2020. Monte Carlo geometry processing: A gridfree approach to PDE-based methods on volumetric domains. ACM Transactions on Graphics 39, 4 (2020).
- Peter Shirley, Changyaw Wang, and Kurt Zimmerman. 1996. Monte carlo techniques for direct lighting calculations. ACM Transactions on Graphics (TOG) 15, 1 (1996), 1–36.
- 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 (TOG) 42, 4 (2023), 1–16.
- László Szirmay-Kalos, György Antal, and Mateu Sbert. 2005. Go with the Winners Strategy in Path Tracing. J. WSCG 13, 2 (2005), 49–56. http://wscg.zcu.cz/wscg2005/ Papers_2005/Journal/!WSCG2005_Journal_Final.pdf
- Karthik Vaidyanathan, Sven Woop, and Carsten Benthin. 2019. Wide BVH traversal with a short stack. In Proceedings of the Conference on High-Performance Graphics. 15–19.
- Delio Vicini, David Adler, Jan Novák, Fabrice Rousselle, and Brent Burley. 2019. Denoising deep monte carlo renderings. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 316–327.
- Jiří Vorba and Jaroslav Křivánek. 2016. Adjoint-driven Russian Roulette and Splitting in Light Transport Simulation. ACM Trans. Graph. 35, 4, Article 42 (July 2016), 11 pages. doi:10.1145/2897824.2925912
- Ingo Wald, Philipp Slusallek, Carsten Benthin, and Markus Wagner. 2001. Interactive rendering with coherent ray tracing. In *Computer graphics forum*, Vol. 20. Wiley Online Library, 153–165.
- Bruce Walter, Sebastian Fernandez, Adam Arbree, Kavita Bala, Michael Donikian, and Donald P Greenberg. 2005. Lightcuts: a scalable approach to illumination. In ACM SIGGRAPH 2005 Papers. 1098–1107.
- Yu-Chen Wang, Yu-Ting Wu, Tzu-Mao Li, and Yung-Yu Chuang. 2021. Learning to cluster for rendering with many lights. ACM Transactions on Graphics (TOG) 40, 6 (2021), 1–10.
- Francis Williams, Jiahui Huang, Jonathan Swartz, Gergely Klar, Vijay Thakkar, Matthew Cong, Xuanchi Ren, Ruilong Li, Clement Fuji-Tsang, Sanja Fidler, Eftychios Sifakis, and Ken Museth. 2024. fVDB: A Deep-Learning Framework for Sparse, Large-Scale, and High-Performance Spatial Intelligence. ACM Transactions on Graphics (SIGGRAPH) 43, 4, Article 133 (July 2024), 15 pages. doi:10.1145/3658226
- Chris Wyman, Markus Kettunen, Daqi Lin, Benedikt Bitterli, Cem Yuksel, Wojciech Jarosz, Pawel Kozlowski, and Giovanni De Francesco. 2023. A Gentle Introduction to ReSTIR: Path Reuse in Real-time. In ACM SIGGRAPH 2023 Courses (Los Angeles, California). doi:10.1145/3587423.3595511

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

- Henri Ylitie, Tero Karras, and Samuli Laine. 2017. Efficient incoherent ray traversal on GPUs through compressed wide BVHs. In *Proceedings of High Performance Graphics*. ACM, 1–13.
- Chris Yu, Caleb Brakensiek, Henrik Schumacher, and Keenan Crane. 2021a. Repulsive Surfaces. ACM Trans. Graph. 40, 6 (2021).
- Chris Yu, Henrik Schumacher, and Keenan Crane. 2021b. Repulsive curves. ACM Transactions on Graphics (TOG) 40, 2 (2021), 1-21.
- Cem Yuksel. 2019. Stochastic lightcuts. In Proceedings of the Conference on High-Performance Graphics. 27–32.

A Proof of Unbiasedness

Here we prove Theorem 3.1. A restatement of the theorem is as follows:

THEOREM. Eq. 7 is an unbiased estimator of Eq. 1; that is, $E\left[\hat{F}_1(\mathbf{q})\right] = F(\mathbf{q})$.

PROOF. Expanding $E\left[\hat{F}_1(\mathbf{q})\right]$:

Ε

$$\begin{split} \left[\hat{F}_{1}(\mathbf{q})\right] &= \sum_{i=1}^{M} p(I=i) \sum_{\ell=0}^{d_{i}} p(K=\ell) \left(\tilde{m}_{i,0}f(\tilde{\mathbf{p}}_{i,0},\mathbf{q}) + \\ &= \sum_{k=1}^{\ell} \frac{\Delta_{i,k-1}}{p(I\in T_{i,k-1})p(K\geq k)}\right) \\ &= \sum_{i=1}^{M} \sum_{\ell=0}^{d_{i}} p(I=i)p(K=\ell) \left(\sum_{j\in T_{i,0}} m_{j}f(\tilde{\mathbf{p}}_{I,0},\mathbf{q}) + \\ &= \sum_{k=1}^{\ell} \sum_{c\in C(T_{i,k-1})} \sum_{j\in T_{c}} \frac{m_{j}(f(\tilde{\mathbf{p}}_{c},\mathbf{q}) - f(\tilde{\mathbf{p}}_{i,k-1},\mathbf{q}))}{p(I\in T_{i,k-1})p(K\geq k)}\right) \\ &= \sum_{i=1}^{M} m_{i}f(\tilde{\mathbf{p}}_{0},\mathbf{q}) + \sum_{i=1}^{M} \sum_{\ell=0}^{d_{i}} \sum_{k=1}^{\ell} \sum_{c\in C(T_{i,k-1})} \sum_{j\in T_{c}} \\ &\frac{p(I=i)p(K=\ell)m_{j}(f(\tilde{\mathbf{p}}_{c},\mathbf{q}) - f(\tilde{\mathbf{p}}_{i,k-1},\mathbf{q}))}{p(I\in T_{i,k-1})p(K\geq k)} \\ &= \sum_{i=1}^{M} \left(m_{i}f(\tilde{\mathbf{p}}_{0},\mathbf{q}) + \\ &\sum_{k=1}^{d_{i}} \frac{\left(\sum_{j\in T_{i,k-1}} p(I=j)\right)\left(\sum_{\ell=k}^{d_{i}} p(K=\ell)\right)}{p(T\in T_{i,k-1})p(K\geq k)} \\ &m_{i}(f(\tilde{\mathbf{p}}_{i,k},\mathbf{q}) - f(\tilde{\mathbf{p}}_{i,k-1},\mathbf{q}))\right) \\ &= \sum_{i=1}^{M} F_{\mathcal{P}_{i}}(\mathbf{q}) \\ &= F(\mathbf{q}) \end{split}$$

The third equality comes from the fact that every path has a common root (i.e., $T_{i,0} = T_0$ for all *i*) and $\sum_{i=1}^{M} \sum_{\ell=0}^{d_i} p(I = i)p(K = \ell) = 1$, and the fourth equality comes from rearranging and grouping like terms from contribution swaps along different paths with common prefixes.

B Pseudocode

Pseudocode for the Barnes-Hut approximation, and our path sample estimator, are given below.

Algorithm 1: Barnes-Hut approximation, barnesHut				
Inputs : Tree node T_a , kernel function f , query point q ,				
Outputs : Approximation of Eq. 1 (i.e., Eq. 2)				
- ··· F ····· ··· · · F ···············				
1 T_a has total mass $ ilde{m}_a$ and center of mass $ ilde{\mathbf{p}}_a;$				
$_{2} F \leftarrow 0;$				
$3 \tilde{\beta} \leftarrow \frac{\ \mathbf{q} - \tilde{\mathbf{p}}_a\ }{ B_a };$				
4 if T_a is a leaf or $\tilde{\beta} \ge \beta$ then				
5 $F \leftarrow \tilde{m}_a f(\tilde{\mathbf{p}}_a, \mathbf{q});$				
6 else				
7 for $T_b \in C(T_a)$ do				
8 $F \leftarrow F + \text{barnesHut}(T_b, f, \mathbf{q}, \beta);$				
9 end				
10 end				
11 Return F;				
Algorithm 2: Stratified path sample estimator,				
pathSampleEstimator				

Inputs : Tree root T_0 , kernel function f, query point q, samples per subdomain S Outputs: Approximation of Eq. 1 1 $F \leftarrow 0;$ ² for $T_a \in C(T_0)$ do T_a has total mass \tilde{m}_a and center of mass $\tilde{\mathbf{p}}_a$; 3 $CV_a \leftarrow \tilde{m}_a f(\tilde{\mathbf{p}}_a, \mathbf{q});$ 4 $F_a \leftarrow 0;$ 5 for $s \in 1 \dots S$ do 6 $I \leftarrow uniformSamplePath(T_a);$ 7 $k \leftarrow 0;$ 8 $p_{\rm rr} \leftarrow 1;$ 9 $T_{I,k} \leftarrow T_a;$ 10 while $T_{I,k}$ is not a leaf do 11 12 $p_{I,k} \leftarrow \text{Eq. 5};$ if rand() $\geq p_{I,k}$ then 13 break; 14 end 15 $\Delta_{I,k} \leftarrow \text{Eq. 6};$ 16 $\begin{array}{l} \begin{array}{c} \hline P_{agg} \leftarrow \hline |T_{I,k}| \\ F_{a} \leftarrow F_{a} + \hline |T_{a}| \\ \hline P_{agg} \cdot p_{rr} \\ \hline P_{rr} \leftarrow P_{rr} \cdot p_{I,k}; \end{array}$ 17 18 $T_{I,k} \leftarrow \text{child of } T_{I,k} \text{ containing } I ; k \leftarrow k + 1;$ 19 end 20 end 21 $F \leftarrow F + CV_a + \frac{1}{S}F_a;$ 22 23 end 24 return F;

SIGGRAPH Conference Papers '25, August 10-14, 2025, Vancouver, BC, Canada.

Stochastic Barnes-Hut Approximation for Fast Summation on the GPU • 11



Fig. 9. Compared to the ground truth (top, left), an estimator with Russian roulette disabled (top, middle left) and our estimator (top, right) produce virtually identical Coulomb potential fields from a source set sampled from a lemur polyline, while using a fixed Russian roulette probability (top, middle right) has significant errors. The log error of each estimator is shown in the bottom row. Despite our method having slightly higher log error than a disabled Russian roulette estimator, it is roughly 5 times faster to evaluate on a GPU, only being slightly slower than the simple fixed-probability Russian roulette scheme.

C Russian Roulette Comparison

Russian roulette schemes are a technique used to improve an estimator's efficiency (the reciprocal of variance times computational cost), with the tradeoff of increasing variance. In other words, good Russian roulette schemes should indirectly lead to variance reduction by allowing more samples to be taken for the same computational budget, or equivalently, a slight increase in variance in less time for the same sample budget. We investigated the quality of our Russian roulette algorithm described in the main paper by comparing it to alternative estimators without Russian roulette (i.e., all paths are traversed to the leaves), and using a fixed probability of $\frac{1}{2}$. The results are shown in Fig. 9. Overall, our method achieves results comparable to an estimator without Russian roulette while being 5 times faster.