

Multi-Scale Modeling and Rendering of Granular Materials

Johannes Meng^{2,1} Marios Papas^{1,3} Ralf Habel¹
 Carsten Dachsbaumer² Steve Marschner⁴ Markus Gross^{1,3} Wojciech Jarosz^{1,5*}

¹Disney Research Zürich ²Karlsruhe Institute of Technology ³ETH Zürich ⁴Cornell University ⁵Dartmouth College




Figure 1: We propose a multi-scale procedural approach for modeling granular materials. The user specifies the bounding shape for the aggregate material (top left), selects a pre-packed tile of grain bounding spheres (top middle), within which we instantiate randomly rotated copies of the selected exemplar grains (bottom left) according to the specified mixing ratios. The SANDCASTLE contains about 2 billion grains, each composed of approximately 200k triangles. We report the high-order / total render times in hours and the variance in parentheses. Our approach (top half) renders the high-order scattering over 12× (50 vs. 628 hrs) faster than explicitly path tracing (EPT) the individual grains (bottom half) while providing visually indistinguishable results. The insets on the right provide equal time and equal variance comparisons.

Abstract

We address the problem of modeling and rendering granular materials—such as large structures made of sand, snow, or sugar—where an aggregate object is composed of many randomly oriented, but discernible grains. These materials pose a particular challenge as the complex scattering properties of individual grains, and their packing arrangement, can have a dramatic effect on the large-scale appearance of the aggregate object. We propose a multi-scale modeling and rendering framework that adapts to the structure of scattered light at different scales. We rely on path tracing the individual grains only at the finest scale, and—by decoupling individual grains from their arrangement—we develop a modular approach for simulating longer-scale light transport. We model light interactions within and across grains as separate processes and leverage this decomposition to derive parameters for classical radiative transport, including standard volumetric path tracing and a diffusion method that can quickly summarize the large scale transport due to many grain interactions. We require only a one-time precomputation per exemplar grain, which we can then reuse for arbitrary aggregate shapes and a continuum of different packing rates and scales of grains. We demonstrate

*The work was done while the author was employed at Disney Research.

our method on scenes containing mixtures of tens of millions of individual, complex, specular grains that would be otherwise infeasible to render with standard techniques.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing;

Keywords: physically based rendering, granular media

1 Introduction

In this paper we consider rendering materials consisting of large assemblies of macroscopic granules. Such granular materials are common in our everyday environment: sand, gravel, and snow; sugar, salt, ground spices, laundry detergent; ocean spray or bubbles in a carbonated beverage—any large pile or aggregate object consisting of randomly oriented grains in which the individual scatterers are discernible (see Figure 2). Common among all these examples is the potential for detailed appearance at the scale of grains, but smooth large-scale appearance of the aggregate due to multiple scattering between grains. Rendering granular materials accurately and efficiently at arbitrary scales remains an open problem. Individual grains can have complex shapes and complex scattering behavior, while at the same time they can have high albedo, so that long paths with many scattering events can remain important (see Figure 3).

Treating each individual grain as explicit geometry and simulating global light transport using path tracing [Kajiya 1986] and its variants is a general solution, but is only practical for small collections of grains. At the other extreme, the aggregate object could be interpreted as a continuous medium, the smooth, large-scale appearance of which may be well expressed by participating media rendering techniques [Cerezo et al. 2005] derived from the radiative transfer equation (RTE) [Chandrasekar 1960]. Methods based on the diffu-



Figure 2: The appearance of real granular materials varies with the size, shape, packing rate, and scattering properties of individual grains.



Figure 3: A selection of real granular materials from Figure 2 illuminated from underneath with a cell phone’s flash light to reveal their translucency. Granular materials can exhibit long-range light transport that still retains complex small-scale structure. Each HDR capture was individually tonemapped for display.

sion approximation [Stam 1995; Jensen et al. 2001] could potentially further accelerate multiple scattering computation. Unfortunately, even if appropriate volumetric parameters could be determined for the continuous medium, such an approach cannot reproduce the fine-scale structure of individual visible grains.

Moon et al. [2007] were the first in computer graphics to consider these problems with classical volumetric light transport and proposed a data-driven technique to account for the effects of statistical dependence in the presence of closely packed scatterers. Their precomputed *shell transport functions* (STFs) store the large-scale spatio-angular light transport within a given discrete random medium, allowing faster rendering by enabling both shadow connections and larger steps during volumetric path tracing. When preceded with a few bounces of path tracing through the discrete objects, this two-scale approach provided both fine-scale details and accurate large-scale multiple scattering. While significantly faster than exhaustively path tracing the grain geometry, the approach requires a scene-dependent precomputation, making it expensive to explore the effects of using different packings or types of grains. Moreover, it does not consider level-of-detail for distant views.

We propose a procedural modeling approach for granular materials and a multi-scale framework for rendering these models that adapts to the structure of scattered light at different scales. Figure 1 illustrates our rendering pipeline. The input to our approach includes the geometry describing the shape of the aggregate object, the packing rate (density) and scale of the grains composing the aggregate, and the geometry and material properties of one or more types of individual grains.

We model a granular material as a procedurally defined, tiled as-

sembly of objects representing individual grains, and we use three models for the transport of light in the material. To capture the appearance of visible grains, the most detailed model explicitly path-traces the grain geometry (explicit path tracing, or EPT). To more efficiently capture larger-scale transport above the scale of grains, we approximate the granular material as a continuous medium and render using volumetric path tracing (VPT). To avoid the need to trace long paths, for scales above the mean free path of the volumetric medium we use a diffusion-based approximation.

Using these three levels of approximation in a practical rendering system requires solving two fundamental problems. First, we need to obtain parameters for each rendering technique that are visually consistent with the finest scale. We do this without any expensive scene-dependent precomputation that would preclude modifying the shape or composition of the aggregate object. To achieve this, we develop a stochastic “teleportation” model of light transport which accounts for *intra*-grain transport and combine it with analytic estimates for *inter*-grain propagation imported from the physics literature [Torquato 2001]. After a one-time precomputation for an individual grain, we can reuse the results to derive parameters for arbitrary aggregate shapes and a continuum of different packing rates and scales. Second, we need to determine when to use each of the three approximations to obtain maximum efficiency gain while remaining visually accurate, which we achieve using carefully designed heuristics based on how paths diverge in the medium and how deeply they penetrate the medium.

Our resulting multi-scale algorithm is able to efficiently render vast scenes consisting of millions to billions of individual grains (e.g. Figure 1 right).

2 Related Work

Aggregate scattering solutions. Condensing aggregate scattering behavior of complex geometric structures into more compact scattering functions has long been a research problem in computer graphics. Reflection from complex surfaces can be modeled using analytic micro-facet models [Torrance and Sparrow 1967], or by tabulating light scattered off real [Matusik et al. 2003] or virtual [Westin et al. 1992; Ashikhmin et al. 2000; Kimmel and Baranoski 2007; Sadeghi et al. 2012] surfaces. We also use a Monte Carlo approach to tabulate compact scattering functions for a single grain. Some work also expresses mesoscopic volumetric structure, propagation, and scattering using bidirectional texture functions [Dana et al. 1999; Filip and Haindl 2009], voxels [Kajiya and Kay 1989; Neyret 1998], or similar representations [Chen et al. 2004; Tong et al. 2005]. Pharr and Hanrahan [2000] formalized the notation of using aggregate

scattering functions and elaborated on their utility for a number of practical rendering problems including subsurface scattering.

Moon et al. [2007] precomputed scattering functions on spherical shells inside a random medium, accounting for the aggregate effect of transport within and between individual objects that make up the medium. In our work we precompute a scattering function from a single grain of our aggregate material, obtaining modularity by decoupling this from the arrangement, scale, and packing of grains.

Level-of-detail and prefiltering. Our approach combines several distinct representations and rendering methods to model the appearance and light transport within granular materials. Such multi-scale, level-of-detail approaches have a long history in graphics. Luebke et al. [2002] give a thorough treatment of early work, whereas Bruneton and Neyret [2012] provide a recent survey of non-linear prefiltering techniques for smoothly transitioning between multi-scale representations of appearance.

Our treatment of granular materials has conceptual similarity to work in hair [Moon and Marschner 2006; Zinke and Weber 2006; Zinke et al. 2008] and cloth rendering [Schröder et al. 2011; Zhao et al. 2013] where individual fibers are visible and therefore need to be modeled explicitly, but costly and smooth large-scale transport is approximated. We likewise leverage an explicit geometric representation for granules at the finest scale, but assume procedurally instantiated, randomly oriented grains with arbitrarily complex geometry.

Densely packed media. The packing density or packing rate f is the volume fraction that grains occupy relative to the background medium [Dullien 1991]. For suspended media, such as granules in water or bubbles in ice, arbitrarily low packing rates can be achieved. The classical RTE is a reasonable model for packing rates up to around 0.1 [Randrianalisoa and Baillis 2010]. However, non-point scattering effects start to significantly affect the aggregate appearance at higher packing rates. Unfortunately, many materials we are interested in rendering have significantly higher packing rates. Bagged white sugar has a packing rate of ~ 0.46 [Bubnik et al. 1998], dry sand ranges from 0.55–0.61 [McWhorter and Sunada 1977], and snow has packing rates from 0.1–0.7 [Cuffey and Paterson 2010].

A particularly well-studied sub-problem is that of densely packed hard spheres. So-called “poured random packings” result from pouring hard spheres into a bed, and induce packing rates in the range of 0.609–0.625; shaking the bed of spheres to densify can reach rates above 0.625 [Dullien 1991]. Song et al. [2008] recently derived an analytic maximum density of ~ 0.634 for this problem. We leverage work from sphere packings to inspire both a multi-scale modeling metaphor and to enable multi-scale rendering with minimal precomputation.

Non-classical RTE methods outside graphics. Understanding the optical and heat transfer properties of densely packed media is important in many fields, including thermal engineering, atmospheric sciences, and nuclear reactor physics. Accurate solutions can be obtained using a full wave approach [Foldy 1945] and solving Maxwell’s equations [Durant et al. 2007], but are computationally infeasible and not necessary for graphics-related problems.

A popular alternative is to make the “homogeneous phase approximation” [Randrianalisoa and Baillis 2009], where the densely packed medium (with statistically dependent scattering) is approximated using the standard RTE, which is valid at low volume fractions or for long-scale transport. Randrianalisoa and Baillis [2010] proposed a data-driven procedure to fit such “effective RTE properties” to

the results of a Monte Carlo random walk through a discrete granular medium. This requires a separate, scene-dependent simulation, akin to Moon et al., which we avoid. Instead, we first model radiative transport using a non-classical process with “teleportation” and derive compound parameters from that.

Singh and Kaviany [1992] model dependent radiative transfer by precomputing a scattering function with positional offsets for a large, smooth, dielectric spherical particle. Replacing the standard phase function in the Discrete Ordinates Method [Chandrasekar 1960] with it allows them to account for intra-grain transport effects. We also precompute a scattering function with position offsets, but consider grains of arbitrary geometry and material properties. We use this to augment volumetric path tracing with stochastic “teleportation” through grains. This, combined with analytic estimates for inter-grain distances [Torquato 2001], allows us to model the free-flight distribution of a granular medium.

Donovan et. al. [2003] used a two-step process where they sampled chord lengths for transport between grains from a statistical distribution and computed intra-grain transport by instantiating grain geometry. Our teleportation model is very similar, but uses a statistical representation for intra-grain transport. We currently rely on a standard diffusion approximation of exponentially-distributed paths, though recent diffusion derivations for non-exponentially distributed paths [d’Eon 2013] might further improve our accuracy.

3 Granular Model and Algorithm Overview

We propose a multi-scale modeling and rendering framework for granular materials that adapts to the structure of scattered light at different scales. In this section we introduce our stochastic model for granular materials (Section 3.1) and describe our three distinct rendering approaches in succession, which rely on explicitly path tracing the grain geometry (EPT: Section 3.2), volumetrically path tracing a homogeneous medium (VPT: Section 3.3), and applying the diffusion approximation (DA: Section 3.4). We initially assume suitable parameters are provided for each method. In the sections that follow we show how to obtain parameters that make VPT and DA consistent with EPT at larger scales (Section 4) without requiring a per-scene precomputation, and explain our criteria for switching between the methods to obtain the best performance while remaining visually accurate (Section 5). Figure 4 illustrates the various inputs and stages of our approach.

3.1 Stochastic Granular Model

The input to our random model for granular media (refer to green items in Figure 4) consists of a surface describing the boundary of the aggregate object, a model for the geometry and material properties of individual grains, numbers describing the size and packing rate of the grains composing the aggregate, and numbers defining the mixing weights if multiple grain types are used. Our basic assumption is that the interior of the bounding surface is filled with many non-overlapping grains that are randomly positioned and oriented. The goal of our granular model is to procedurally define a random but repeatable collection of individual grains that fill the medium and do not intersect one another, while keeping memory requirements modest.

Tiled sphere packings. To avoid explicitly modeling and storing the arrangement of millions of packed grains, we rely on a procedurally tiled infinite field of packed spheres to fill the volume of the aggregate shape. Each sphere contains one grain, and since the spheres do not overlap the grains cannot overlap. This approach greatly simplifies many parts of our system, but limits the maximum

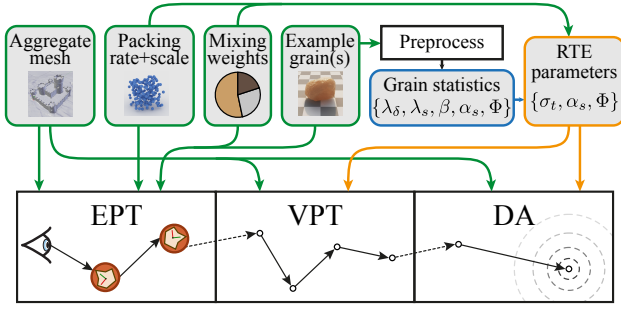


Figure 4: When rendering the granular material (bottom), primary rays from the eye start out by explicitly path tracing (EPT) grains, then after enough scattering events take place (Section 5), the path is continued using volumetric path tracing (VPT), and eventually terminated using a diffusion calculation (DA) to approximate the contribution of all further interactions. The input (top, green) to EPT consists of an aggregate mesh, packing rate & scale of spheres, example grains, and mixing ratios of these grains. We combine grain scattering statistics (blue)—calculated in a one-time, per-grain-type preprocess—with the packing rate and mixing ratios to obtain RTE parameters (orange) needed for the VPT and DA methods.




Figure 5: To intersect procedural granular media such as this bunny of cubical grains (left), we consider all bounding spheres, from an infinitely tiled space, that reside inside a given aggregate boundary mesh. We perform intersection tests with randomly rotated, instanced grain geometry (right) inside each bounding sphere.

attainable packing fraction to approximately 0.634 for spherical grains and less for non-spherical grains (see Section 2). This is a limitation especially if the grains are highly non-spherical, such as grains of rice, as we will discuss in Section 7.

The basic building block of our procedural sphere packing is a *tile*, a cuboid containing a number of non-intersecting spheres. We use the algorithm of Skoge et al. [2006] to generate a dictionary of periodic tiles for various packing ratios f . Figure 1 shows an example tile for $f = 0.63$. Though tilings generated with these blocks may exhibit some repetition artifacts, we found that in practice they hardly matter due to the randomization involved in our grain instantiation which we describe next. Aperiodic tiling [Peytavi et al. 2009] approaches could be used to remove tiling artifacts completely if necessary.

Randomized instantiation. We treat each distinct sphere of the infinite field as a location for instantiating a randomly chosen grain from the set of provided exemplars. The exemplar grains are chosen according to the user-specified mixing probabilities (see Figure 1). The aggregate mesh provided by the user (assumed to be watertight with consistent normals) defines the region in space where grains are instantiated and each instance is randomly rotated within its bounding sphere. Figure 5 illustrates this idea. The random seed used for each individual grain is a deterministic function of the three-dimensional tile index and the sphere id. This ensures consistent grain instantiation for all rays when needed.

Algorithm 1 Procedural Grain Instantiation

```

1: function INTERSECTGRAINS(Ray  $r$ , Float  $t_{\min}$ , Float  $t_{\max}$ )
2:   while  $t_{\min} < t_{\max}$  do
3:      $i_v = \text{VoxelIndex}(r, t_{\min})$ 
4:     if PartiallyInside( $i_v$ ) then
5:        $o_v = \text{VoxelOrigin}(i_v)$ 
6:        $o = \text{Origin}(r) - o_v$ 
7:       if IntersectTile( $o$ , Direction( $r$ ),  $t_{\min}$ ,  $t_{\max}$ ) then
8:          $p = \text{BoundingSphereCenter}() + o_v$ 
9:         if SurelyInside( $i_v$ ) or PointInside( $p$ ) then
10:          if IntersectInstance( $r$ ,  $t_{\min}$ ,  $t_{\max}$ ,  $p$ ) then
11:            return True
12:          end if
13:        end if
14:         $t_{\min} = \text{BoundingSphereExit}()$ 
15:      end if
16:    end if
17:     $t_{\min} = \text{FindNextVoxelT}(r, t_{\min})$ 
18:  end while
19:  return False
20: end function

```

3.2 Explicit Path Tracing

For the finest level, we use an algorithm which explicitly path traces against the actual geometry of the individual grains generated by the procedural model just described. We dub this technique explicit path tracing (EPT) and illustrate the expected input in Figure 4. The core challenge with this approach is representing the vast collection of grains and efficiently tracing rays through it.

Ray tracing grains. Even with instancing, storing the transformation matrices for the collection of billions of grain instances is impractical. Instead, we procedurally instantiate and discard the grains on demand during ray traversal of the scene. To facilitate this, we voxelize the boundary mesh using two bits per voxel, marking voxels as either “fully outside,” “fully inside,” or “partially inside.” One voxel corresponds, in size and position, to one instance of the base tile.

During ray tracing, we intersect the ray with the voxel grid to obtain the first intersection with a voxel that is either “fully inside” or “partially inside.” We then intersect the bounding spheres overlapping the voxel. If the voxel is marked “fully inside,” we use the first intersected bounding sphere. If it is marked “partially inside,” we first have to determine if the bounding sphere’s center is contained in the aggregate shape. We do this by tracing a ray from the sphere center against the aggregate shape¹. If the bounding sphere is deemed to be outside the aggregate shape, we proceed to the next bounding sphere along the ray until a valid bounding sphere is found.

Upon hitting a bounding sphere, we select a random example grain (see *Randomized instantiation* above) according to our mixing weights and temporarily transform the ray into the (randomly rotated) local coordinate system of the grain. We then intersect the grain geometry to determine the final intersection point. If there is no intersection, we proceed to the next bounding sphere intersected by the ray (if any). Algorithm 1 lists pseudo-code for this procedure.

Rendering. With this grain intersection machinery in place, we employ path tracing using next event estimation and multiple importance sampling [Pharr and Humphreys 2010] as our finest level

¹Our procedure defines a grain to be inside the aggregate shape if its center is within the aggregate boundary mesh. Instantiating grains for spheres that lie entirely within the mesh is also possible with more costly overlap tests.

rendering technique. Note that while next event estimation may be useful outside the aggregate object and close to its boundary, there are virtually no unblocked direct connections to light sources from within the granular material. Furthermore, most of our granular materials consist of grains with specular boundaries, rendering shadow rays useless and underpinning the need to switch to other approaches.

3.3 Volumetric Path Tracing

While EPT retains the necessary grain-level details present in low-order scattering, it quickly becomes impractical for higher-order transport when rays become incoherent. Luckily, for incoherent rays it is less important to precisely track which rays hit which grains, so we transition (see Figure 4) to a continuous volumetric representation for medium-order transport, which we simulate using volumetric path tracing (VPT) [Kajiya 1986; Rushmeier 1988].

For this rendering approach we assume each bounding aggregate shape is filled with an index-matched homogeneous participating medium defined by its extinction coefficient σ_t , single-scattering albedo α_s , and the scattering phase function $\Phi(\cos \theta)$ (marked in orange in Figure 4). We discuss how to obtain these parameters from the properties of the granular packing in Section 4.

Apart from the fact that VPT does not need to instance or intersect grain geometry, transitioning to VPT allows us to leverage shadow connections from inside the granular medium, which was virtually impossible with EPT. As we will see in Section 6, this allows VPT to efficiently render scenes that are intrinsically hard for EPT, such as large aggregates of highly specular, non-absorbing grains.

3.4 Diffusion Approximation

While VPT has significantly lower variance than EPT, it can still be computationally expensive for highly scattering materials which require long transport paths. We therefore transition (see Figure 4) to a rendering technique based on a fast diffusion approximation [Stam 1995; Jensen et al. 2001] to account for large-scale transport due to high-order multiple scattering. Diffusion allows us to effectively short-circuit the recursion of VPT and approximate long scattering paths directly.

We switch to diffusion by sampling a location on the boundary mesh and estimating the diffusion transport. We use a technique inspired by the method of Li et al. [2005], though we use d'Eon and Irving's [2011] improved diffusion model instead of the classical dipole [Jensen et al. 2001], and integrate its contribution using Monte Carlo [Habel et al. 2013]. We also propose a different virtual source placement procedure which we found produces slightly better results than the approach described by Li et al., and we account for the finite thickness of the medium by using a multipole expansion [Donner and Jensen 2005]. We detail these changes in Appendix A. The parameters for this approach consist of the reduced medium parameters (α'_s and σ'_t) which we obtain from the RTE parameters (σ_t , α_s , Φ) in the usual way [Jensen et al. 2001] using similarity theory.

3.5 Importance of Using a Hybrid Method

While it would be possible to use any of the aforementioned methods in isolation, we aim to combine them into a hybrid approach since we wish to render granular materials at multiple scales, potentially in a single shot (e.g., a sand dune that is close to the camera at the bottom of the frame, and far away at the top). In such scenarios, using EPT alone would likely be prohibitively expensive, while the more approximate techniques would not faithfully reproduce fine surface details where individual grains are discernible (see Figure

11). One of our core contributions is showing how to systematically combine these disparate methods and representations to ensure visual consistency between grains visible at vastly different scales both across the image or across time in an animation.

4 Precomputation and RTE Parameters

Our goal in this section is to derive RTE parameters (σ_s , σ_t , Φ) for a classical homogeneous medium which match the statistical scattering behavior of the discrete granular material. This task is non-trivial since granular materials with high packing rates do not actually satisfy a core assumption of the RTE—namely, that scatterers are sufficiently small and well-separated so that each interaction may be considered statistically independent.

To handle this seeming incompatibility, we first (Section 4.1) introduce a statistical, non-classical model of light transport which we call *teleportation transport* (TT). This intermediate model allows us to reason about the non-point scattering effects present in granular materials and will provide a stepping stone to derive classical RTE parameters (Section 4.2) for our VPT and DA approaches without the need for a scene-dependent precomputation.

4.1 The Teleportation Transport Model

We wish to create a purely statistical model which characterizes paths constructed by path tracing the granular material as described previously in Section 3.2. While TT introduces several statistical independence assumptions, we design these to closely approximate the EPT behavior once rays have been sufficiently randomized by a series of explicit events.

With this analogy in mind, our TT model consists of two core steps which occur in alternation. The first step considers *inter-grain* transport by randomly deciding how far along a ray to move before the next interaction with a grain bounding sphere occurs. Upon hitting a bounding sphere, the second step considers *intra-grain* transport by scattering and “teleporting” the ray directly to a point where it exits the sphere. A deflection angle and spatial offset from the incoming point and direction are randomly selected to produce an outgoing point and direction from which the path is continued. This is done in such a way that the distribution of continuations is the same as is observed with the grain geometry, when averaged over all possible entry points and grain orientations. Note that this model is similar to a volumetric path tracing view of the RTE, but with the introduction of non-point scattering effects due to teleportation.

Inter-grain transport. To characterize inter-grain transport, we need to determine the distribution of free-flight distances from the point where a path exits one grain's bounding sphere to the point where it enters the next grain's. We could tabulate this distribution as Moon et al. [2007] did, by accumulating free-flight distances from many paths traced through a particular packing. We would, however, like to avoid a precomputation that depends on a particular packing.

We instead make use of prior work on transport in sphere packings—this is appropriate since the distances we are interested in are between interactions with *bounding spheres*, not with the *grains* themselves. Dixmier [1978] derived a simple model for free-flight distances in sphere packings, known in that context as outer chord-length distributions. He showed that these outer chord-lengths can be modeled by an exponential distribution

$$p_b(z) = \sigma_b e^{-\sigma_b z}, \quad \text{with} \quad \sigma_b = \frac{3}{4R} \frac{f}{1-f}, \quad (1)$$

where R is the radius of the spheres, f is the sphere packing rate, and $\lambda_b = 1/\sigma_b$ is the average distance between bounding sphere

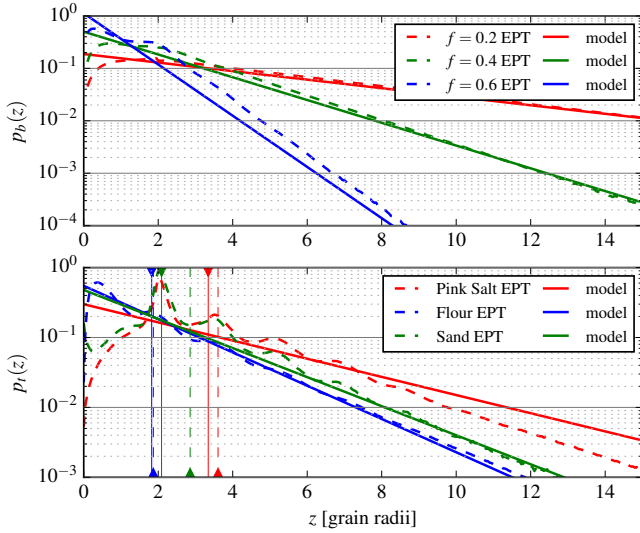


Figure 6: We compare (top) Monte Carlo measurements against analytic estimates [Dixmier 1978] of the path segment lengths $p_b(z)$ between packed dielectric spheres in an infinite volume. We also compare (bottom) our full free-flight distribution $p_t(z)$ (Eq. 4) for three of our grain (mixtures) against Monte Carlo measurements. The packing rate is $f = 0.63$ for all grain types. The hit probability is $\beta = 0.53$ (Pink Salt), $\beta = 0.66$ (Flour), and $\beta = 0.63$ (Sand). The mean free paths are indicated by vertical lines.

interactions. As Torquato and Lu [1993] showed, this formula is only an approximation, albeit one that works remarkably well. Several authors [Levitz 1993; Donovan et al. 2003; Olson et al. 2006] have shown that this formula is incorrect for very small distances (< 1 grain radius), but this is usually ignored in favor of the simple analytical model. We also found the same effect in our experiments (see Figure 6, top), but did not find that accounting for this difference by using tabulations or non-exponential analytical models provided a substantial benefit. We thus also opted for the simple exponential model, which yields plausible results for various packing rates.

Intra-grain transport. The goal of intra-grain transport is to statistically model the interaction of a random ray with a randomly oriented grain inside a bounding sphere. We could fully summarize the effect of a single grain interaction on a path by utilizing a virtual scatterometer setup akin to Westin et al.’s [1992] approach. This involves placing a single grain in a bounding sphere and tracing a collection of rays, distributed over directions and the cross section of the sphere. After interacting with the grain zero or more times, a path would either exit the bounding sphere or be absorbed. The points and directions upon entering $(\mathbf{x}_i, \vec{\omega}_i)$ and exiting $(\mathbf{x}_o, \vec{\omega}_o)$ the bounding sphere (see illustration) characterize the scattering function $S(\mathbf{x}_i, \vec{\omega}_i \rightarrow \mathbf{x}_o, \vec{\omega}_o)$ of the grain (in the sense of scattering functions used in layer-oriented transport calculations [Pharr and Hanrahan 2000]). We call this function the *teleportation scattering distribution function* (TSDF), because it models scattering on a single grain by “teleporting” a photon from \mathbf{x}_i to \mathbf{x}_o .

Like a BSSRDF [Jensen et al. 2001], the TSDF is parameterized over the complete set of rays entering and exiting a volume of space; in this way it is also similar to the BFSDF [Zinke and Weber 2007] used to describe scattering from fibers.

One could create a full rendering algorithm around the teleportation

transport model, which would require storing the TSDF and sampling it and Equation (1) in alternation. We provide more details and an analysis of one prototype implementation in the supplemental material. In practice however, we found that we can obtain similar visual fidelity (see Figure 7) with significantly less implementation complexity by using the TT model only to obtain classical RTE parameters. This allows us to leverage standard RTE rendering methods, such as VPT and DA, in our approach. Additionally, deriving classical RTE parameters will require us to gather only a compact set of statistics of the TSDF from the virtual scatterometer, instead of exhaustively tabulating the full 8D function, as we discuss next.

4.2 Classical RTE & Diffusion Parameters

Given the teleportation transport model just described, we wish to use it as an intermediate model to obtain RTE parameters (σ_t , α_s , Φ) describing an “equivalent” homogeneous participating media. Since our teleportation transport is a generalization of classical transport (with the addition of non-point scattering), there will inherently be some information loss in this conversion. In particular, classical transport has only one mechanism—propagation with exponential free-flight sampling dictated by σ_t —to affect the positions (distances) along rays. In TT, however, these distances are influenced not only by exponential inter-grain transport, but also by the intra-grain teleportation encoded in the TSDF. Our goal is therefore to extract all directional effects of the TSDF into the phase function Φ and combine the positional influences of the TSDF and the inter-grain transport into an effective extinction coefficient σ_t .

4.2.1 The Phase Function & Albedo

In the classical RTE, scattering interactions are assumed to happen at a single point. This amounts to assuming that outgoing direction $\vec{\omega}_o$ depends only on the incident direction $\vec{\omega}_i$. We can therefore extract the directional component of the TSDF and use this directly as a phase function Φ . Due to symmetry of the random grain rotations, we additionally know that the phase function will depend solely on the cosine of the deflection angle $\cos \theta = \langle \vec{\omega}_i, \vec{\omega}_o \rangle$. This allows us to easily tabulate just a 1D distribution in the scatterometer setup by recording the dot product between the incident and outgoing directions of any ray that interacts with the grain. We found that tabulating this 1D distribution is practical, and gives superior results to fitting a simple parametric model like the Henyey-Greenstein [1941]

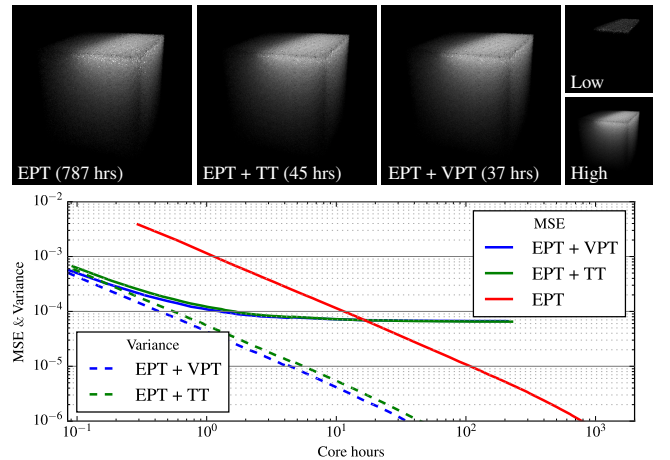


Figure 7: Equal quality comparison (top) of a cube of glass beads, as shown in Moon et al. [2007, Fig. 9]. The last column shows individual low- and high-order contributions computed with EPT and VPT, respectively. MSE and variance vs. core hours for all three techniques (bottom).




Figure 8: The expected free-flight distance between grain interactions λ_t is a potentially infinite sum λ_β of sphere-to-sphere distances λ_b and delta scattering distances λ_δ due to the possibility $(1 - \beta)$ of passing through spheres unobstructed, followed by a scattering teleportation distance $\alpha_s \lambda_s$.

phase function. The integral of 1D distribution over the unit sphere is the effective grain albedo α_s (i.e. the ratio of rays that hit the grain and subsequently escape the bounding sphere).

4.2.2 Combined Free-Flight Distribution

To summarize the overall free-flight distribution between scattering events, we need to augment the free-flight distribution between bounding spheres, Equation (1), in two ways. Firstly, we need to account for the possibility that some rays pass completely unscattered through bounding spheres; and secondly, we need to account for the potential teleportation that occurs to rays that do intersect grains within their bounding spheres. We handle each effect in turn.

Accounting for delta scattering. As shown in Figure 8, a grain may not fill its bounding sphere completely. This means that only a certain fraction β of rays incident on the bounding sphere will hit the grain, while the remainder will exit the bounding sphere without ever interacting with the grain. We call β the *hit probability*, and estimate it directly from our scatterometer simulations. Associated with it is a distance λ_δ , the mean spatial offset of an unscattered ray.

Consider a ray that is just exiting a grain bounding sphere after interacting with the grain contained within (refer to Figure 8). We are now interested in the expected distance λ_β to the next bounding sphere in which an interaction will take place. This length generally differs from λ_b , the expected distance to the next bounding sphere, in Equation (1) because of the hit ratio β . We can express it as:

$$\begin{aligned} \lambda_\beta &= \underbrace{\beta}_{P_1} \lambda_b + \underbrace{\beta(1 - \beta)}_{P_2} [(\lambda_b + \lambda_\delta) + \lambda_b] + \dots \\ &= \beta \sum_{i=0}^{\infty} (1 - \beta)^i [i(\lambda_b + \lambda_\delta) + \lambda_b], \end{aligned} \quad (2)$$

where $P_i = \beta(1 - \beta)^{i-1}$ is the probability of scattering on the i^{th} grain whose bounding sphere is intersected. Using $\sum_{i=0}^{\infty} i(1 - \beta)^i = (1 - \beta)/\beta^2$ and $\sum_{i=0}^{\infty} (1 - \beta)^i = 1/\beta$, we obtain

$$\lambda_\beta = (\lambda_b + \lambda_\delta) \frac{1 - \beta}{\beta} + \lambda_b. \quad (3)$$

Note that for $\beta = 1$ (the grains fill their bounding spheres completely), this formula collapses to λ_b . The other extreme of $\beta = 0$ (no ray will ever scatter) results in an infinite mean free path.

Account for teleportation. So far, λ_β only considers inter-grain transport and we must still account for intra-grain teleportation. A ray may be absorbed in the interaction with a grain, in which case it never exits the grain bounding sphere, or it may experience a teleport step. During preprocessing, we estimate in our scatterometer

simulation the mean teleport vector $(\mathbf{x}_o - \mathbf{x}_i)$ of all rays that scatter with (and are not absorbed by) the grain. Due to symmetry of the random rotations, this mean teleport vector is always aligned with the original ray direction, and we denote its length λ_t . By accounting for the probability of teleportation vs. absorption (the albedo α_s), our final exponential free-flight distribution between exit points of grain interactions (refer to Figure 8) becomes:

$$p_t(z) = \sigma_t e^{-\sigma_t z}, \quad \text{with} \quad \sigma_t = \frac{1}{\lambda_t} = \frac{1}{\lambda_\beta + \alpha_s \lambda_s}. \quad (4)$$

In Figure 6 (bottom), we compare $p_t(z)$ for three grains from our example scenes to results obtained from explicitly path tracing grains in an infinite medium. Notably, the means of the distributions—indicated by vertical lines—match well, indicating that our analytic model is able to capture the average behavior of granular media, even though the true free-flight distribution is not an exponential.

4.2.3 Effective RTE and Diffusion Parameters

In summary, we extract just two average distances, two probabilities, and a 1D distribution from our grain scatterometer simulations: $\{\lambda_\delta, \lambda_s, \alpha_s, \beta, \Phi(\cos \theta)\}$. In Table 1 we present the scalar statistics for the grains used in our results, and we visualize the phase functions for two grains in the supplemental. From these, we obtain our effective RTE parameters by reusing the grain albedo α_s and phase function $\Phi(\cos \theta)$ directly, and by computing the extinction coefficient σ_t using Equation 4. We compute effective RTE parameters for mixtures of different grains by computing a weighted average of their individual statistics. We then derive diffusion parameters by computing the mean scattering cosine g from $\Phi(\cos \theta)$, and from that the reduced scattering coefficient $\sigma'_s = (1 - g)\alpha_s \sigma_t$, which gives a reduced extinction coefficient $\sigma'_t = \sigma'_s + (\sigma_t - \sigma_s)$ and reduced albedo $\alpha'_s = \sigma'_s / \sigma'_t$.

Table 1: Extracted statistics (averaged over wavelength and rounded to two digits) for the grains used in our results.

	glass beads	flour	pink salt	sand	snow	cane sugar	brown sugar
λ_δ	0.00	0.87	0.97	0.85	0.86	0.77	0.77
λ_s	1.44	0.18	1.00	0.50	1.40	1.17	0.88
α_s	1.00	1.00	0.99	0.68	1.00	1.00	0.76
β	1.00	0.66	0.53	0.63	0.65	0.71	0.71

Validation. In Figure 9, we validate our derivations by comparing the fluence and radiance fields produced within granular materials composed of dielectric spheres or snow grains. We visualize the illumination for various inclinations on a spherical detector surrounding a beam source inside the medium.

We obtain the data visualized in Figure 9 using a numerical simulation by tracing photons from a beam source in an infinite medium of unit radius grains. To ensure sufficient initial randomization, we initialize each photon at a random location and orientation in the infinite medium, and trace it until it has experienced 10 grain interactions. We take the resulting position and direction of the photon as the origin of the beam source, around which we center a detector sphere of a predefined radius (in our case, 7.5 or 15 mean free paths, which corresponds to 10 or 20 grain radii). We then continue tracing the photon until it exits the sphere.

Upon exiting the sphere, we parametrize the photon's position and direction in the coordinate system shown in Figure 10: The inclination θ describes how much the photon deviates from its initial direction, and the exitant direction ω_o is expressed as its projection




Figure 9: The top plots visualize fluence $[\text{W mm}^{-2}]$ due to a unit beam source in an infinite packing of dielectric spheres (left) and snow grains (right). The bottom plots visualize the corresponding radiance field $[\text{W mm}^{-2} \text{sr}^{-1}]$. There are results for detector radii of 7.5 and 15 mean free paths. Simulations used $4 \cdot 10^8$ samples.

(x, y) onto the tangent plane of the detector sphere at the point of exit, p_e . Each photon is then accumulated in a 3D-histogram.

The fluence (Figure 9, top) computed using VPT with our RTE parameters generally matches the EPT ground truth very well at various inclinations, albeit with some over-estimation of the forward scattering component for short distances. The bottom plots of Figure 9 show exitant radiance fields (in essence, this is the information contained in Moon et al.’s [2007] STF) for the same grain simulations. As with fluence, we find that VPT overestimates forward scattering at small distances, but otherwise closely resembles the EPT result. This indicates that VPT could be used to precompute Moon et al.’s STF, or, conversely, STF could further accelerate our algorithm, albeit at the cost of a scene-dependent precomputation.

Figure 11 compares EPT to VPT with and without DA for a more traditional rendering scenario. The results show that our derived RTE parameters closely approximate the large-scale transport within granular materials. In isolation, however, neither VPT nor DA can retain the high-frequency details of discernible grains, so next we describe our automatic criterion for transition between these




Figure 10: The coordinate system used for plots in Figure 9. A photon is traced from a beam source at location p_i with direction ω_i . The photon follows a path (shown as a dotted line) until it exits the detector sphere at p_e . The exitant direction ω_o is then projected onto the blue tangent plane with normal ω_e .

rendering methods.

5 Switching between Rendering Techniques

We assume that primary rays are spawned outside the granular material and begin rendering using explicit path tracing of the grain geometry. We start with EPT since there are high-frequency details at the grain level which we can only recover with EPT. To motivate why it could be both beneficial and safe to switch from EPT to VPT and DA, we consider the bundle of rays belonging to a pixel. As long as these rays are somewhat coherent, EPT will perform reasonably well. However, once the rays become statistically independent by scattering in different directions, the variance from EPT would become enormous since the chance of randomly hitting small light sources tends to zero. Luckily, when rays are incoherent, there is also no reason to keep track of which ray hits which grain and we can safely switch to a smooth model like VPT or the DA. Our goal in this section, is therefore to develop an automatic switching criterion which allows us to use EPT when we need to—when rays are still coherent on the pixel scale—and which switches to the smoother VPT and DA techniques once rays become incoherent. Figure 11 shows that by intelligently combining our techniques with EPT, we can retain small-scale details while also accurately approximating the long-scale transport within granular materials.

EPT → VPT. To put this intuition into practice, we spawn a bundle of $N = 16$ paths per pixel which are traced in lock-step through the grains. Individual paths might be terminated (e.g. by escaping the medium or due to Russian roulette), while others are traced further. In this case, we simply continue with the bundle of remaining paths. At each bounce k , we gauge the level of ray coherence by computing the standard deviation σ_k of the untruncated N_k vertex positions $(\mathbf{x}_1^k, \dots, \mathbf{x}_{N_k}^k)$. We switch to VPT when:

$$\sigma_k > \tau \frac{N_k}{N} \quad (5)$$




Figure 11: Diagonal comparison of proposed techniques (top left) with EPT (bottom right) on a sphere composed of snow grains. We compare four combinations of our methods: The first two combinations show how VPT with and without diffusion compare against EPT. The latter two show how preceding these methods with a few bounces of EPT can accurately recover high-frequency detail.

where τ is a user-specified multiple of the maximum grain radius r ; the N_k/N factor allows bundles with few surviving rays to switch more quickly. For all our scenes we use a threshold of 4 grain radii. When σ_k exceeds this threshold, all paths in a bundle switch upon hitting the next grain. If only one ray of a bundle remains, we switch immediately at the next grain interaction.

To account for the fact that the VPT and DA techniques may converge faster than EPT, we can optionally dedicate it more computation time by probabilistically pruning paths when they attempt to switch. We achieve this using Russian roulette with an acceptance threshold P_a . Later in this section, we show how to automatically compute the value of P_a which optimally balances the variance of the different techniques.

VPT \rightarrow DA. While diffusion can estimate multiple scattering very efficiently, it also introduces a number of approximations and assumptions (most notably the assumption of isotropic scattering and planar bounding geometry) that limit its accuracy in the general setting. For switching from VPT to DA, we adopt a criterion adapted from the work of Li et al. [2005] which aims to allow DA only when the approximations would not be too noticeable.

As in that work, our primary criterion for switching to diffusion is a minimum distance between the VPT path vertex \mathbf{x}_i^k and the surface of the boundary mesh. While Li et al. used a threshold of one reduced mean free path, we use a threshold of $d_{DA} = \min(1/\sigma'_t, 0.5/\sigma_{tr})$ —that is, we switch if \mathbf{x}_i^k is at least 1 reduced mean free path away from the boundary or at least half a diffuse mean free path. This allows us to accept diffusion connections more frequently for lower albedo materials which can be better represented using d'Eon and Irving's improved diffusion model [2011]. Li et al. accelerate this distance calculation using a kD-tree filled with points uniformly sampled over the boundary mesh, but we found this to be impractical for optically thick granular media since the sampling would need to be excessively dense. Instead, we reuse the kD-tree acceleration structure of the aggregate mesh geometry, and perform a depth-first search to check whether any triangles overlap a sphere of radius d_{DA} centered at \mathbf{x}_i^k . The appendix details our diffusion source placement.

High Order Acceptance Rate. Since convergence of low-order paths computed with EPT is usually slower than convergence of high-order paths from VPT or DA, we can often reduce the total variance of the image by dedicating more sample budget to low-order paths without requiring each to also spend effort estimating high order scattering (which may already be converged). To accomplish this, we perform Russian roulette with an acceptance rate of $P_a \leq 1$ when switching between EPT and VPT. While we could set this manually as done by Moon et al. [2007], the optimal acceptance probability is scene and lighting dependent, making it difficult to estimate by hand. Instead, we propose to automatically compute the value of P_a that would minimize the variance of the final image.

Our approach requires a relatively inexpensive preprocess (taking about 1% of the total render time) during which we render a down-scaled image 1% of the original size (reduced by $10\times$ in both dimensions) with the same total number of samples per pixel as the final image. To ensure that the statistics of the switching criteria are not affected, we set the footprint of the camera ray packets to the size of the original (smaller) pixels. For each pixel in the low resolution preview we incrementally estimate the sample variance of the low-order $V_L(x, y)$ and high-order $V_H(x, y)$ contributions, and we also accumulate the total CPU time for computing these contributions as $t_L(x, y)$ and $t_H(x, y)$. The variance of the combined image as a function of the number of samples n and the acceptance rate P_a is

approximately:

$$V \approx \frac{1}{n} \left(V_L + \frac{V_H}{P_a} \right), \quad (6)$$

where V_L and V_H are the averages across the image of $V_L(x, y)$ and $V_H(x, y)$, respectively. The total time needed for the image is approximately:

$$t = n(t_L + P_a t_H), \quad (7)$$

where t_L and t_H are the total CPU time needed for low order and high order paths respectively. By solving for the value of P_a which minimizes t , we obtain:

$$P_a = \sqrt{\frac{V_H t_L}{V_L t_H}}. \quad (8)$$

In our results we report the render times both with $P_a = 1$ and with the optimal acceptance rate computed using Equation (8).

6 Results

We implemented our method in Mitsuba [Jakob 2010] as a new *Integrator* for our rendering algorithms and a new *Shape* primitive for our granular volumes. We rendered all results on a homogeneous cluster with nodes containing two 12-core Intel Xeon E5-2697v2 processors at 2.7 GHz with 64 GB RAM and report all render times in core-hours. For time related comparisons, we independently render many low sample count images across the machines and average the resulting floating-point images for the final result. In all results, RGB channels are rendered separately and then combined and tonemapped with an sRGB gamma curve.

For all our figures we estimate variance (in parentheses) by dividing the *time to unit variance* (ttuv) by the core hours used for the specific scene and method. Time to unit variance for each method is an estimate of the time needed to achieve a variance of 1, assuming $1/N$ variance reduction. We compute ttuv by rendering low sample count versions of the image, and then multiplying the average sample variance over all pixels with the average time needed to render a low sample count image.

In Table 2 we provide a summary of render times and analysis of variance for the main scenes shown in Figures 1, 12, and 13. We quantify the efficiency of each algorithm using ttuv. On the left half of the table we report ttuv and the resulting speedup compared to EPT for the overall render time, and on the right for only the high-order transport which we aim to accelerate with our approximations.

EPT+VPT provides a significant overall speedup over EPT in all our scenes, ranging from $2.1\text{--}30\times$. This is largely due to VPT's ability to create shadow connections to the light sources at various depths inside the granular medium, whereas EPT must rely on random chance for any grain with a dielectric boundary. Enabling our automatic acceptance rate calculation (P_a) provides further improvements, with speedups now ranging from $2.2\text{--}101\times$. The optimal acceptance rate balances the computation between EPT and higher-order methods to minimize ttuv. Enabling diffusion provides significant additional speedup in scenes with low-absorption grains and long transport paths such as SNOWMAN (Figure 13, $259\times$ vs. $101\times$ of EPT+VPT) at the cost of some visible bias. We attain these speedups in total render time primarily by reducing the computation time for the high-order scattering component. Measuring just the computation time spent on high-order scattering, our full method obtains dramatic speedups ranging from $241\text{--}7394\times$.

In the rest of this section we will evaluate the accuracy and efficiency of our methods on our main scenes and scientific experiments.




Figure 12: Our SPICES scene showcases flour, pink salt, brown sugar, and white sugar. We model flour using diffuse grains and the other spices with dielectric grains filled with homogeneous media (top right). We provide approximate equal variance (top left) and equal-time comparisons (bottom left) between pure EPT and our successively faster techniques incorporating VPT and DA, reporting the high-order / total render time in hours and the (variance). Our full approach computes the high-order scattering over $120\times$ faster (34 vs. 4148 hrs) than EPT, resulting in an $8\times$ overall speedup for equal variance (top left). Note that all six images on the left include all light transport, just simulated using successive subsets of our full approach. Any differences are due to bias and/or variance. We visualize our automatic switching depth (bottom right) and the individual contributions (middle right) of each technique, which sum to our final image for the EPT+VPT+DA example.

Sandcastle. Figure 1 shows a sand castle scene containing about 2 billion grains, each modeled as a homogeneous medium enclosed in a dielectric boundary. The top half of the castle image is rendered using EPT+VPT+DA and the bottom half using EPT. For equal variance, we find that both our approximate methods render this scene about $2.2\times$ faster than EPT with indistinguishable results. The cropped insets compare our approximate methods (middle) to the progress of EPT for the same CPU time (top) and for similar variance (bottom). The supplemental video contains a zoom sequence on this scene accompanied with visualizations of the switch depth as the camera approaches the grains.

While the high-order scattering speedup in the SANDCASTLE is over $200\times$, this results in only a modest overall speedup of $2.2\times$. This is primarily because the sand grains composing the castle have a low albedo of 0.68 (see Table 1), and due to Russian roulette few paths survive to our accelerated high-order approaches. This is compounded by the fact that our implementation currently enforces at least one grain interaction in EPT mode before switching to VPT. This allows us to reproduce the fresnel-like BRDF effects that occur at grazing angles with the aggregate, but also prevents us from completely side-stepping costly EPT for distant views. We hope to address this in future work.

Spices. In Figure 12 we render a scene with four very different granular materials: flour, pink salt, brown sugar and cane sugar. Our

approximate methods accelerate the high-order scattering computation by over $120\times$ compared to EPT, while introducing only barely perceptible bias. The visualization of the per-pixel switch depth from EPT to VPT shows that our threshold performs reasonably well: the highly coherent and high-albedo cane sugar switches later than the absorbing brown sugar or mildly absorbing pink salt. For the case of flour our method usually switches fairly quickly, except when the packet rays are trapped between the tightly packed opaque flour grains. Our visualization of individual components shows that our single, automatic switching criterion allows each rendering method to make a reasonable energy contribution to the final image.

Snowman. In Figure 13 we render a 1 m (top) and 10 cm (bottom) tall snowman. To achieve the miniaturization effect we increased the grain radius on the bottom row by a factor of 10. In the right-most images we show a closeup, rendered with our full method (EPT+VPT+DA), focused on the rim around the left eye. We render each snowman with three variations of our method. We show equal quality renders of the full light transport with EPT, EPT+VPT and EPT+VPT+DA. DA allows for a significant speedup over volumetric path tracing at the cost of bias which is most visible near creases. The black bordered insets (top right) show zoom-ins of the fallen button of the full resolution images and the green border insets (top left) show the progress of each method for the same computation time. The somewhat darker appearance of EPT at 107 and 89 hours

Table 2: Time to unit variance in seconds for all methods. The value in parenthesis is the speedup relative to EPT alone. We report times for 100% acceptance probability (P_a) as well as our automatically computed optimal P_a . On the left we report times for rendering both low order and high order transport whereas on the right we report the timings and speedup only for the high order component.

Scene	Time to unit variance, combined (sec)					Time to unit variance, high order (sec)				
	EPT	EPT+VPT		EPT+VPT+DA		EPT	EPT+VPT		EPT+VPT+DA	
		$P_a = 1$	$P_a = \text{opt}$	$P_a = 1$	$P_a = \text{opt}$		$P_a = 1$	$P_a = \text{opt}$	$P_a = 1$	$P_a = \text{opt}$
SPICES	1617.3	400.3 (4.00 \times)	249.7 (6.50 \times)	219.8 (7.40 \times)	201.9 (8.00 \times)	688.8	5.3 (130.8 \times)	5.0 (137.8 \times)	0.3 (2554. \times)	0.4 (1968. \times)
SANDCASTLE	30.7	14.4 (2.10 \times)	13.7 (2.20 \times)	13.9 (2.20 \times)	13.7 (2.20 \times)	11.8	0.1 (206.8 \times)	0.1 (230.0 \times)	0.4 (26.50 \times)	0.1 (242.0 \times)
SNOWMAN LG.	4000.8	132.0 (30.3 \times)	39.6 (101. \times)	16.3 (246. \times)	15.4 (260. \times)	3548.6	8.8 (404.6 \times)	9.2 (385.3 \times)	0.4 (8473. \times)	0.5 (7394. \times)
SNOWMAN SM.	2868.4	171.1 (16.8 \times)	52.7 (54.4 \times)	28.1 (102. \times)	25.2 (114. \times)	2338.0	6.1 (381.5 \times)	8.6 (271.4 \times)	0.4 (6431. \times)	0.5 (4997. \times)

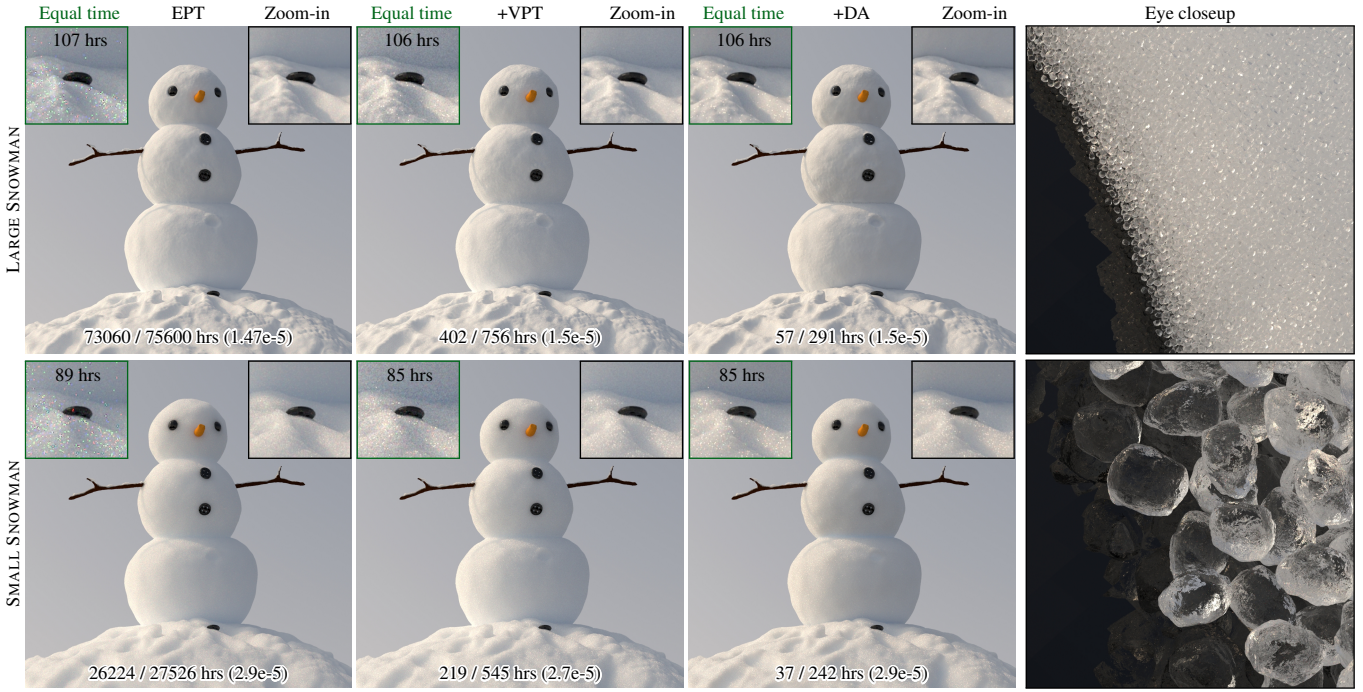


Figure 13: For the tall (1 m, top) and tiny (10 cm, bottom) SNOWMAN we report the high-order / total render time in hours and the (variance). The equal time insets (green) show unconverged results at roughly 106 and 85 hours for the large and small snowman, respectively. All renderings in this figure simulate the full light transport. Any differences within the same row are due to bias and/or variance.

is due to the large amount of energy concentrated in colored bright firefly noise. These scenes are dominated by long transport since snow is optically thick and has very high albedo, so as expected, DA provides a huge speedup. We accelerate high-order scattering by a factor of over 5000–7000 \times , leading to an overall speedup of 108–299 \times .

7 Conclusion & Discussion

In this paper we have shown how to approximate the large-scale behavior of a granular medium with a smooth radiative transport model. Normally, path tracing a high-albedo granular medium like snow is completely impractical because very long, very high-variance paths contribute much of the illumination. Our method uses smooth approximations to reduce both the variance (by enabling shadow connections) and the length (by using diffusion) of paths, greatly reducing the time spent on high-order scattering. The result is that render times are no longer dominated by long paths; instead the computation is spent rendering the visible structure of grains due to low-order scattering. As shown by our results, this makes it feasible to render many practically relevant scenes that would otherwise take unacceptably long to converge.

The speedups of our method are limited by the continued need to path trace the low-order contributions that create visible grain structure; in many cases we succeed in reducing the cost of high-order paths essentially to zero, so that in order to obtain further speedup one must turn to a new problem, that of rendering the glittery, structured low-order contributions faster without smoothing out the appearance.

Limitations in packing density. Relying on the well-studied properties of spherical packings allowed us to decouple inter-grain and intra-grain transport effects; however, it also imposes some practical limitations. We assume packings of spheres of a single radius and instantiate individual grains of smaller physical size within these bounding spheres. This inherently places a limit on how closely

packed the actual grains may be. One partial workaround would be to use an “aggregate grain” which is itself a tight packing of irregularly shaped grains. While this could reduce the packing limitations for non-spherical grains such as rice, the empty space between aggregate grains would still limit the overall packing density. Leveraging work on ellipsoid packings [Donev et al. 2004] could prove fruitful for such randomly oriented oblong grains.

Generality and mixtures. While our method allows mixtures of different grain types in a single medium, we currently assume the mixing ratios are homogeneous at the scale of the aggregate. For some scenes, mid-scale heterogeneity, in which the ratios of particle types vary spatially, is desirable. Our method could be extended to this case by modulating the background medium’s extinction coefficient (Equation (1)) using a volumetric texture and computing correspondingly varying free-flight distances.

Appearance editing. We have proposed a bottom-up approach for specifying the appearance of granular materials where the grain properties and their packing rate dictate the large-scale appearance. An interesting alternative (akin to inverse bi-scale appearance design [Wu et al. 2013]) would be a top-down approach where the user specifies, or measures, the desired large-scale appearance, and the system proposes the individual grain properties and their arrangement from a dictionary of possibilities. Editing one scale independently while maintaining a fixed appearance at other scales would be a valuable appearance design tool for scenes containing granular media. Truly useful material design would require further improvements to rendering speed to obtain interactivity.

8 Acknowledgements

The authors are grateful to Maurizio Nitti, who modeled the Sandcastle, Spices, and Snowman scenes along with individual grains, and Romain Prévost, who helped with the schematic in Appendix A.

The authors would like to thank Andrew Selle, Brent Burley and Chuck Tappan for fruitful discussions about granular media, as well as Nelson Max and Jan Novák for providing valuable feedback on early versions of the paper.

References

- ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. S. 2000. A microfacet-based BRDF generator. In *Proc. of ACM SIGGRAPH*, 65–74.
- BRUNETON, E., AND NEYRET, F. 2012. A survey of non-linear pre-filtering methods for efficient and accurate surface shading. *IEEE Trans. on Visualization and Computer Graphics* 18, 2, 242–260.
- BUBNIK, Z., KADLEC, P., URBAN, D., AND BRUHNS, M. 1998. *Sugar Technologists Manual*. Verlag Dr. Albert Bartens.
- CEREZO, E., PÉREZ, F., PUEYO, X., SERON, F. J., AND SILLION, F. X. 2005. A survey on participating media rendering techniques. *The Visual Computer* 21, 5, 303–328.
- CHANDRASEKAR, S. 1960. *Radiative Transfer*. Dover Publications.
- CHEN, Y., TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2004. Shell texture functions. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 23, 3, 343–353.
- CHRISTENSEN, P. H., HARKER, G., SHADE, J., SCHUBERT, B., AND BATALLI, D. 2012. Multiresolution radiosity caching for global illumination in movies. In *ACM SIGGRAPH Talks*.
- CUFFEY, K., AND PATERSON, W. 2010. *The physics of glaciers*. Academic Press.
- DANA, K. J., VAN GINNEKEN, B., NAYAR, S. K., AND KOENDERINK, J. J. 1999. Reflectance and texture of real-world surfaces. *ACM Trans. on Graphics* 18, 1, 1–34.
- D’EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 30, 4, 56:1–56:14.
- D’EON, E. 2013. Rigorous asymptotic and moment-preserving diffusion approximations for generalized linear boltzmann transport in arbitrary dimension. *Transport Theory and Statistical Physics* 42, 6-7, 237–297.
- DIXMIER, M. 1978. Une nouvelle description des empilements aléatoires et des fluides denses. *Le Journal de Physique* 39, 873–895.
- DONEV, A., CISSE, I., SACHS, D., VARIANO, E. A., STILLINGER, F. H., CONNELLY, R., TORQUATO, S., AND CHAIKIN, P. M. 2004. Improving the density of jammed disordered packings using ellipsoids. *Science* 303, 5660 (Feb.), 990–993.
- DONNER, C., AND JENSEN, H. W. 2005. Light diffusion in multi-layered translucent materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 24, 3, 1032–1039.
- DONOVAN, T., SUTTON, T., AND DANON, Y. 2003. Implementation of chord length sampling for transport through a binary stochastic mixture. In *Nuclear Mathematical and Computational Sciences: A Century in Review, A Century Anew*.
- DULLIEN, F. A. L. 1991. *Porous Media: Fluid Transport and Pore Structure*, 2nd ed. Academic Press Inc.
- DURANT, S., CALVO-PEREZ, O., VUKADINOVIC, N., AND GREFFET, J.-J. 2007. Light scattering by a random distribution of particles embedded in absorbing media: full-wave Monte Carlo solutions of the extinction coefficient. *Journal of the Optical Society of America* 24, 9, 2953–2962.
- FILIP, J., AND HAINDL, M. 2009. Bidirectional texture function modeling: A state of the art survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence* 31, 11, 1921–1940.
- FOLDY, L. L. 1945. The multiple scattering of waves. I. General theory of isotropic scattering by randomly distributed scatterers. *Physical Review* 67, 107–119.
- HABEL, R., CHRISTENSEN, P. H., AND JAROSZ, W. 2013. Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 32, 4.
- HENYEV, L. G., AND GREENSTEIN, J. L. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93, 70–83.
- JAKOB, W., 2010. Mitsuba renderer. <http://mitsuba-renderer.org>.
- JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., AND HANRAHAN, P. 2001. A practical model for subsurface light transport. *Computer Graphics (Proc. SIGGRAPH)* 35, 511–518.
- KAJIYA, J. T., AND KAY, T. L. 1989. Rendering fur with three dimensional textures. In *Computer Graphics (Proc. SIGGRAPH)*, 271–280.
- KAJIYA, J. T. 1986. The rendering equation. *Computer Graphics (Proc. SIGGRAPH)* 20, 143–150.
- KIMMEL, B. W., AND BARANOSKI, G. V. G. 2007. A novel approach for simulating light interaction with particulate materials: application to the modeling of sand spectral properties. *Optics Express* 15, 15, 9755–9777.
- LEVITZ, P. 1993. Knudsen diffusion and excitation transfer in random porous media. *Journal of Physical Chemistry* 97, 3813–3818.
- LI, H., PELLACINI, F., AND TORRANCE, K. E. 2005. A hybrid Monte Carlo method for accurate and efficient subsurface scattering. In *Proc. Eurographics Symposium on Rendering*, 283–290.
- LUEBKE, D., WATSON, B., COHEN, J. D., REDDY, M., AND VARSHNEY, A. 2002. *Level of Detail for 3D Graphics*. Elsevier Science Inc.
- MATUSIK, W., PFISTER, H., BRAND, M., AND McMILLAN, L. 2003. A data-driven reflectance model. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 22, 3, 759–769.
- MCWHORTER, D. B., AND SUNADA, D. K. 1977. *Ground-water hydrology and hydraulics*. Water Resources Publications, LLC.
- MOON, J. T., AND MARSCHNER, S. R. 2006. Simulating multiple scattering in hair using a photon mapping approach. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 25, 3, 1067–1074.
- MOON, J. T., WALTER, B., AND MARSCHNER, S. R. 2007. Rendering discrete random media using precomputed scattering solutions. In *Proc. Eurographics Symposium on Rendering*, 231–242.
- NEYRET, F. 1998. Modeling, animating, and rendering complex scenes using volumetric textures. *IEEE Trans. on Visualization and Computer Graphics* 4, 1, 55–70.
- OLSON, G., MILLER, D., LARSEN, E., AND MOREL, J. 2006. Chord length distributions in binary stochastic media in two and three dimensions. *Journal of Quantitative Spectroscopy & Radiative Transfer* 101, 269–283.

PEYTAVIE, A., GALIN, E., MERILLOU, S., AND GROSJEAN, J. 2009. Procedural Generation of Rock Piles Using Aperiodic Tiling. *Computer Graphics Forum (Proc. Pacific Graphics)* 28, 7, 1801–1810.

PHARR, M., AND HANRAHAN, P. M. 2000. Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proc. of ACM SIGGRAPH*, 75–84.

PHARR, M., AND HUMPHREYS, G. 2010. *Physically Based Rendering, Second Edition: From Theory to Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc.

RANDRIANALISOA, J., AND BAILLIS, D. 2009. Radiative transfer in dispersed media: Comparison between homogeneous phase and multiphase approaches. *Journal of Heat Transfer* 132, 2, 023405–023405.

RANDRIANALISOA, J., AND BAILLIS, D. 2010. Radiative properties of densely packed spheres in semitransparent media: A new geometric optics approach. *Journal of Quantitative Spectroscopy and Radiative Transfer* 111, 10, 1372–1388.

RUSHMEIER, H. E. 1988. *Realistic Image Synthesis for Scenes with Radiatively Participating Media*. PhD thesis, Cornell University, Ithaca, NY, USA.

SADEGHI, I., MUÑOZ, A., LAVEN, P., JAROSZ, W., SERON, F., GUTIERREZ, D., AND JENSEN, H. W. 2012. Physically-based simulation of rainbows. *ACM Trans. on Graphics* 31, 1, 3:1–3:12.

SCHRÖDER, K., KLEIN, R., AND ZINKE, A. 2011. A volumetric approach to predictive rendering of fabrics. *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering)* 30, 4, 1277–1286.

SINGH, B., AND KAVIANY, M. 1992. Modelling radiative heat transfer in packed beds. *International Journal of Heat and Mass Transfer* 35, 6, 1397–1405.

SKOGE, M., DONEV, A., STILLINGER, F. H., AND TORQUATO, S. 2006. Packing hyperspheres in high-dimensional Euclidean spaces. *Physical Review E* 74, 4, 041127.

SONG, C., WANG, P., AND MAKSE, H. A. 2008. A phase diagram for jammed matter. *Nature*, 7195, 629–632.

STAM, J. 1995. Multiple scattering as a diffusion process. *Proc. Eurographics Workshop on Rendering*, 41–50.

TONG, X., WANG, J., LIN, S., GUO, B., AND SHUM, H.-Y. 2005. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 24, 3, 1054–1061.

TORQUATO, S., AND LU, B. 1993. Chord-length distribution function for two-phase random media. *Physical Review E* 47, 2950–2953.

TORQUATO, S. 2001. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Interdisciplinary Applied Mathematics. Springer.

TORRANCE, K. E., AND SPARROW, E. M. 1967. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* 57, 9, 1105–1112.

WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proc. SIGGRAPH)*, 255–264.

WU, H., DORSEY, J., AND RUSHMEIER, H. 2013. Inverse bi-scale material design. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 32.

ZHAO, S., HAŠAN, M., RAMAMOORTHY, R., AND BALA, K. 2013. Modular flux transfer: efficient rendering of high-resolution volumes with repeated structures. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 32, 4, 131:1–131:12.

ZINKE, A., AND WEBER, A. 2006. Global illumination for fiber based geometries. In *Proc. Ibero-American Symposium in Computer Graphics (SIACG)*.

ZINKE, A., AND WEBER, A. 2007. Light scattering from filaments. *IEEE Transactions on Visualization and Computer Graphics* 13, 2, 342–356.

ZINKE, A., YUKSEL, C., WEBER, A., AND KEYSER, J. 2008. Dual scattering approximation for fast multiple scattering in hair. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 27, 3, 32:1–32:10.

A Diffusion Details

Given a source location \mathbf{x}_i^k —generated by VPT—we generate the incident location \mathbf{x}_i on the boundary by sampling a uniform random direction $\vec{\omega}_i$ and intersecting the aggregate mesh [Christensen et al. 2012]. We sample the direct illumination at \mathbf{x}_i and set up our dipoles, where the positive receiver is the source location \mathbf{x}_i^k , and the negative receiver is a reflection of \mathbf{x}_i^k about the extrapolated boundary a distance $z_b(0)$ above the plane defined by the surface intersection \mathbf{x}_i and its normal \mathbf{N}_i . One of our differences to Li et al. [2005] is that they define the reflection plane using \mathbf{x}_i and the direction $\vec{\omega}_i$ (instead of the normal \mathbf{N}_i). Another difference is that they assume $r = 0$ by construction, while we compute r as the projected distance of \mathbf{x}_i^k on the plane defined by $(\mathbf{x}_i, \mathbf{N}_i)$, thereby computing the correct distances for scenes adhering to the planar slab assumption.

Additionally, instead of assuming a semi-infinite medium, we estimate a thickness $d = z_r + z_a$ for a finite slab, where z_r is the scalar projection of $(\mathbf{x}_i - \mathbf{x}_i^k)$ onto \mathbf{N}_i , and z_a the scalar projection of $(\mathbf{x}_a - \mathbf{x}_i^k)$ onto \mathbf{N}_a , obtained with an additional ray shot in direction $-\vec{\omega}_i$ from \mathbf{x}_i^k . We define $d_r = \sqrt{z_r^2 + r^2}$ and $d_v = \sqrt{z_v^2 + r^2}$, with $z_v = -z_r - 2z_b(0)$. We then derive the distances to the negative receiver and consecutive dipoles using the standard multi-pole expansion [Donner and Jensen 2005] but using the improved diffusion model [d’Eon and Irving 2011; Habel et al. 2013]. The definitions of the extrapolated boundaries at $z_b(0)$ and $z_b(d)$ can be found in [Donner and Jensen 2005]. We assume that the boundary of the granular medium is index matched (i.e. $\eta = 1$), since our grains are suspended in vacuum. Finally, we resume VPT from \mathbf{x}_i . The auxiliary point and normal $(\mathbf{x}_a, \mathbf{N}_a)$ are only used for computing the thickness of the approximate parallel slab and subsequent multi-pole mirroring.

