# Rendering Realistic Snow

Cameron L. Chrisman*
University of California, San Diego

## Abstract

The vast majority of previous research of snow in computer graphics has been on the construction of relatively high level geometric descriptions. Indeed this is not unreasonable, as it is the most salient visual characteristic of snow, however, a large number of visual cues come from the surface appearance of snow: its density, purity, cues to curvature and thickness. These cues are essential to the realistic representation of snow in computer graphics. We present a brief overview of the physical optics of snow, and examine techniques to accurately model the surface appearance of snow and its small scale geometry.

**Keywords:** snow, ice, subsurface scattering, appearance modeling

## 1 Introduction

A physically accurate technique for rendering snow is important to computer graphics in many ways. Most importantly, snow is a ubiquitous material. During the Northern Hemisphere Winter, snow can cover more than 40% of the world's landmass [Fearing 2000]. Materials that we are visually exposed to the most are in turn those to which we are most visually sensitive. Ubiquitous materials such as skin, hair, plants, animals, food products, water, clouds, ice, and snow, are all require extremely accurate models to be photorealistically convincing. In addition, extreme winter climates present logistic difficulties in film production, making the computer generation of such scenes an attractive alternative.

In addition to being a very commonly utilized material, the appearance of snow falls into the general class of granular or porous materials, such as pumice, sand, salt, flour, and gravel. Most of the current techniques for simulating these kinds of materials are not fully convincing. As a corollary, providing a good solution to modeling the granular/porous surface appearance of snow contributes also to modeling the appearance of a larger class of materials.

In the first section of this paper we briefly survey the progress of research in snow optics from the early 1970s, and present an overview of the physical optics of snow. In the second section, we present a parameterized model for snow as a surface material, and a brute force Monte-Carlo technique implementing this model.

## 2 Previous Work:

There is not a large body of published work in the computer graphics community on the subject of modeling snow, despite the fact that snow is such a common material when rendering natural scenes. Of the work that has been done, the primary focus has been on the large scale geometry of snow. [Nishita et al. 1997] uses meta-balls to describe both the large and small scale geometry of snow. Using a volumetric integration approach, multiple scattering of light within the snow is also taken into account. The results, however, remain photo-realistically unconvincing.

Volumetric integration approaches are inherently costly, especially given the highly scattering nature of snow. In their approach, scattering is only calculated up to the third order. This is a bad approximation. Due to its porous nature, the mean free path of snow is approximately the size of a snow grain, and light undergoes many scattering events after entering before leaving the material again [Kokhanovsky and Zege 2004].

[Fearing 2000] focuses purely on modeling the high level geometry of snow using an iterative accumulation model, which is very effective in providing an automated method to apply a layer of snowfall to a static scene. The shading model used, however, is still the simple lambertian approximation. Providing a convincing surface appearance model for snow can greatly improve realism of scenes generated with Fearing's model, and indeed all future snow geometry simulations.

While the amount of research on the optical characteristics of snow has been all but absent in the computer graphics community, there is a very large body of research on the subject in both the physical optics community and the geophysics and hydrology communities. The most significant research on this subject has been by Bohren, Warren and Wiscombe.

The fundamentals of modern theoretical snow optics were established in [Bohren and Barkstrom 1974]. This is one of the most commonly cited papers in snow optics research, and makes prominent use of radiative transfer concepts. Bohren solidifies the usage of scattering optics in snow modeling, providing estimations for the scattering albedo, asymmetry parameter, and mean free path of snow. This marks a significant change in the type of analysis that had been done previously. Beforehand, the vast majority of research on the optical properties of snow had been empirical, and theoretical models proposed by researchers as corollaries to empirical results, most of which were themselves not universally accepted. This as Bohren claims, is due to the lack of well defined experimental parameters. One researcher might focus on the impact of snowpack density correlated to spectral albedo, whereas another might only consider variation in grain size, both of which have potential contributions to an accurate model. Bohren helps to remedy this situation both by establishing a relatively general parameterized model, verifying observations that have been constant with respect to previous literature, and those that have been erratic, and need either further specification of parameters to verify, or more empirical data to be gathered.

In response to advances in theoretical models for the optics of snow, several researchers began to fill in the empirical gaps left by Bohren and others. In [Wiscombe and Warren 1980], and [Warren 1982] detailed empirical data is presented, and variant models are proposed, taking into account a widely varying set of parameters including wavelength dependence, grain size and shape, illumination conditions, purity, density, wetness, as well as using Mie theory vs. geometrical optics. In addition, building upon his previous paper,

---

*e-mail: cchrisma@cs.ucsd.edu

[Bohren 1983] presents a more detailed discussion of the absorptive characteristics of snow.

# 3 Snow Optics

## 3.1 Structural Properties of Ice and Snow

Here we present a brief overview of the optics of snow, and some of the visually observable results. To fully understand the optics of snow, one must understand the basic structural composition of snow. Snow is composed of relatively loosely packed ice crystals, making up as little as 5% total spatial volume. The structure of crystalline ice is canonically hexagonal, but the shape of ice crystals in snow is usually quite different. The reason for this is that snow is produced not by the spontaneous formation of ice a surface, but the constant accumulation of snowflakes.

Ice crystals that form in the upper atmosphere are indeed hexagonal, and refract light very much like a prism (which is the cause for many optical effects such as halos), however the structure of ice crystals change dramatically in their descent through the atmosphere. Being hexagonal, ice crystals form facets. These facets are molecularly smooth, and so water molecules tend to accumulate on the facet edges and grow out at symmetrically skew angles to the facet plane. In turn these growths contain new facets. This processes occurs iteratively and produces the fractal pattern commonly associated with snowflakes. When these flakes hit the ground, they are thermodynamically unstable. To reach a state of equilibrium, that is, minimizing the ratio of free surface energy to volume, they converge to spherical granules. This crystalline structure is different from many materials, in that snowflakes are a *single* ice crystal, and not an ordered collection of small stable crystals. As a result, snowflakes reduce to one or sometimes several coupled spherical granules, each of which has no crystalline substructure.

## 3.2 Optical Characteristics of Ice and Snow

First, let us define some notation.

**Scattering Properties:**

- $\sigma_s$ - scattering coefficient
- $\sigma_a$ - absorption coefficient
- $\sigma_t$ - extinction coefficient
- $p(\theta)$ - phase function, outgoing scattering direction given incoming angle of a photon
- $g$ - asymmetry parameter, the cosine weighted mean of the phase function
- $l_0 = \sigma_t^{-1}$ - mean free path, the average distance a photon travels before a scattering event
- $a_s = \frac{\sigma_s}{\sigma_t}$ - single scattering albedo, the probability a photon will scatter rather than absorb while undergoing a scattering event
- $\alpha$ - spectral albedo, the ratio of incoming and outgoing flux at a given point.

**Structural Properties**

- $\rho$ - density
- $\eta = 1.32$ - index of refraction
- $r$ - grain size (radius)
- $G$ - grain shape
- $W$ - wetness, the volume ratio of liquid water to ice
- $P$ - purity coefficient, the volume ratio of foreign material present to pure water or ice

These properties are the necessary components for defining a well developed radiative transfer model for snow. For the purposes of this paper, many of the parameters are simplified compared to presentations in optical literature. The most prominent of these is wavelength dependence. Properties such as the index of refraction, phase function, and the scattering/absorption coefficients are all commonly wavelength dependent, however most of these are close to constant within the visible range, and may be ignored. The exception to this is the absorption coefficient,the difference with respect to wavelength in the visible spectrum is still rather small, but this difference is cumulative with the number of scattering events, which we will see is very large.

These parameters span most of the parameters used in various research models. We will not show the formulas for all of them, but we will discuss some of the more important ones, as well as some common relationships and parameter reductions.

[Bohren and Barkstrom 1974] establishes reasonably accurate values for many of the parameters that are fairly constant for snow. Most of these values are based on analytic calculations of dielectric spheres. He takes the phase function to be

$$p(\theta) = \frac{T_2 \eta^2 [n \cos(\Theta/2) - 1][\eta - cos(\Theta/2)}{\bar{T}_2 \cos(\Theta/2)[\eta^2 + 1 - 2\eta \cos(\Theta/2)]}$$

. Where $T_2 = \frac{T_l^2 + T_r^2}{2}$ and $\bar{T_2} = 2 \int_2^{\pi/2} d\theta_i \sin\theta_i \cos\theta_i T_2(\theta_i)$, where $T_l$, $T_r$ the Fresnel transmission coefficients. This is significant, because Bohren disregards reflection, both primary and secondary. Attributing the phase function purely to transmission.

Accordingly, the asymmetry parameter of this phase function is $g = 0.874$. This is still a very forward scattering material, but in comparison to other highly scattering materials, such as skin, it exhibits substantial isotropy.

[Kokhanovsky and Zege 2004] give a more concrete values on the intrinsic scattering coefficients for snow, taking $a_s = 0.995$, and also help confirm the validity of Bohren's assumption that modeling properties on an individual sphere basis is accurate by showing that the mean free path $l_0 \approx r$, as well as examining the effects of the $g = 0.874$ phase function for individual spheres in the context of snowpack.

It is also observed that an effective way to model wetness of snow is to increase the grain size. This has the effect of increasing the length of the mean free path, and therefore making the likelihood that light travels through the snow and into the underlying material more probable given a certain level of thickness. In reality wet snow achieves this by having most of the air gaps between grains filled with water, minimizing the effects of the medium change.

# 4 Our Method

## 4.1 Parameterized Model

Now let us present a parameterized model for practical use. This model takes a similar approach to that of [Bohren and Barkstrom

1974], and models ice as a collection of dielectric spheres. While this technique is not without limitations, it provides an adequate first step in modeling the optical characteristics of snow for use in computer graphics. This is especially true when none of the models commonly used (lambertian or simple multiple scattering) come close to producing accurate representations. In addition to this, the Bohren model founded the basis for most of the radiative transfer work to follow, and is still often used today, although a large reason for this is the simplicity of Mie scattering calculations for spheres. Even models that take into account the effects of non-spherical geometry often reduce it down to a collection of spheres in specific configurations. [Kokhanovsky and Zege 2004].

The parameters to the model are outlined below:

- $n$ - number of total snow particles
- $\rho$ - density
- $\eta$ - index of refraction
- $r$ - grain size (radius)
- $r_\sigma$ - grain variation

We take a brute-force Monte-Carlo simulation approach to rendering snow. Snow is modeled as a collection of packed dielectric spheres. Each of these spheres represents a single grain of snow, with a radius of $r \approx .1mm - 10mm$, although exaggerating this parameter can lead to some interesting results. The grain variation $r_\sigma$ represents the uniformity of the grains, and can range from $0 \cdots r$.

Images are rendered using Monte-Carlo raytracing from the eye. Conceptually, the way light reaches our eyes from snow is that light from the sun, or another powerful source hits the snow surface, and scatters (possibly a great number of times), and either gets absorbed, or shoots back out, possibly to your eye. We model this process in reverse. Rays are shot from the eye at jittered pixel locations, when they hit a snow sphere, we use probabilistic sampling to determine if light reflects, or refracts at the interface. This is done by calculating the Fresnel reflection and transmission coefficients $K_t$ and $K_r$ for the given angle. Since $K_t + K_r = 1$ we can treat them as an interval. We choose a uniformly sampled random number $R \in [0..1]$ and transmit if $R \in [0 \cdots K_t)$ and reflect if $R \in [K_t \cdots 1]$. This has the major advantage that there is no combinatorial explosion in the number of rays, and is a commonly used technique in the general simulation if dielectrics.

The ray follows a complex path through the packed spheres, and in our model undergoes four final events to produce the desired shading. One event is that the ray hits a diffuse object which the snow might be covering, in which the shading of the diffuse object is returned. The other possibility is that the ray leaves leaves the snow material and does not a diffuse object along its path. Here the ray can either hit a light source or not. In our model, the light source is modeled as a diffuse hemi-sphere at infinite distance. This can be simulated effectively by simply taking the sign of the y-component of the ray after it has left the material. If the ray hits the light, it gets a contribution equal to the wattage of the light source, and if not, it gets no contribution. These two components make up most of the shading, although underlying materials can be seen when the snow cover is light, and color bleeding effects are often visible.

By necessity the model also contains a parameter for an accumulation step, the total number of particles $n$. Snow without some a high level geometric context is very hard difficult to discern, since most of the light entering snow leaves it, the result is usually a very uniform intensity, provided a sufficient number of samples are taken. A uniform cube of packed spheres is included in our results for completeness, and is a useful tool for gauging the convergence of the Monte-Carlo simulation. Empirically, however, a simple accumulation model is needed to provide a high level geometric context to the surface appearance, and indeed as it turns out, produces a more realistic packing of spheres.

### 4.2 Accumulation Model

The accumulation model used is not a physically based simulation of snow evolution as seen in [Fearing 2000], and was chosen solely to produce a packing of spheres with the desired properties, and to provide some plausible geometric context for snow-covered objects. The approach taken for developing the accumulation model focuses on speed and ease of implementation, as it is only to showcase advanced surface appearance modeling. The results of this naive accumulation model, however, are surprisingly good. Since snow is modeled as spheres optically, we model it the same way physically.

First a bounding box of the scene to be covered is calculated. Ray locations are randomly generated on the top plane of this box with a fully downward orientation. The ray is then cast, and the hit location cataloged. A sphere is added to the scene with a random radius $r'$, within the grain size and variation parameters, with a center offset $r'$ from hit location along its normal. This produces a reasonably good packing, with a relatively small number of overlapping between spheres. In addition, if the hit normal is close enough to perpendicular (i.e. a steep cliff), the particle is discarded.

Some mention should also be made of the need for efficiency in this process. The acceleration structure that was use was a 2-level hierarchical grid, which proved to be very fast for this type of downward raycasting. The disadvantage, however, is that for the acceleration structure is not adaptive. That is, in order for it to be fast and accurate, it must be rebuilt after every object added. Several strategies were tested to help reduce this overhead. The one that was eventually used was a simple “buffer” for accumulated spheres. It consisted of a naive $n^2$ intersection test list, which new spheres were added to. When the size of this structure reached a certain optimal threshold (for most tests it was $\approx 4000$), its contents were removed and added to the grid, which was rebuilt. This proved to work reasonably well, but investing in an adaptive data acceleration structure could conceivably reduce the running time significantly.

## 5 Results

The results of this model look quite a bit like real snow. The brute force nature of our simulation technique captures both the small scale geometry of fallen snow as well as accurately simulating the light transport at or slightly beneath the surface. We show the results of a $120^3$ volume element of snow, as well as a red diffuse sphere and teapot that have undergone snow accumulation. Not all of the parameterizations tested will be shown (and indeed a full parameter sweep was not conducted due to computing constraints), but altering the grain size, index of refraction and number of particles all yielded intuitive results consistent with the appearance of different types of snow.

Figure 3 illustrates a brute force Monte Carlo simulation of snow using packed dielectric spheres at various sampling rates. In this configuration each sphere is of uniform size and are as tightly packed as possible. One can see that the brightest areas occur along closest top edge, where the medium is most optically thick with respect to the viewing angle.

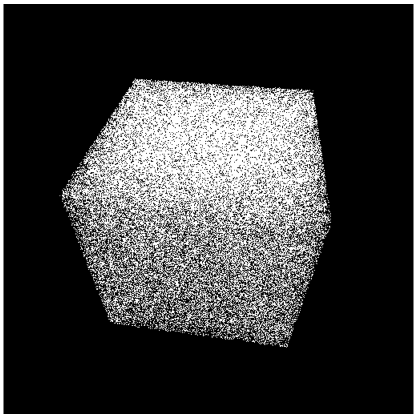

Figure 1: A Monte-Carlo simulation using a large number, $\approx 120^3$ of packed dielectric spheres using 1, sample per pixel

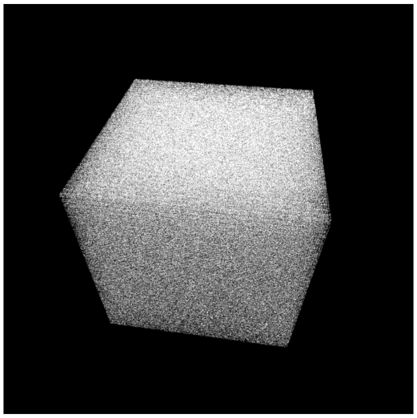

Figure 2: A Monte-Carlo simulation using a large number, $\approx 120^3$ of packed dielectric spheres using 10 samples per pixel

Figure 4 shows a diffuse red sphere covered with thin and thick layers of snow respectively, Both use the same number of particles, and the thickness varies due to grain size.

Figure 5 shows a diffuse red teapot under the same conditions as Figure 4.

## 6 Conclusions & Future Work

While this model represents a promising first step to physically accurate rendering of snow, and indeed can be of practical use, it has several limitations. The most obvious limitation is both processing and memory overhead. Even for small scenes, memory overhead can be on the order of 100-500MB. Processing power requirements are somewhat less of a concern, due both to the fact that efficient acceleration structures can do a very good job of culling unnecessary intersections, and that the model is reasonably tolerant to noise, unlike other scattering phenomenon simulated with Monte-Carlo techniques. Another major limitation is that the lighting model is implicitly fixed to be a completely diffuse hemi-spherical skylight. While this is actually a reasonable approximation for outdoor light under cloud cover, a very common scene in which one would model snow, the appeal of this technique is that it is capable of modeling highly detailed scenes, or “close-ups”, in which local lighting conditions often play a crucial role.

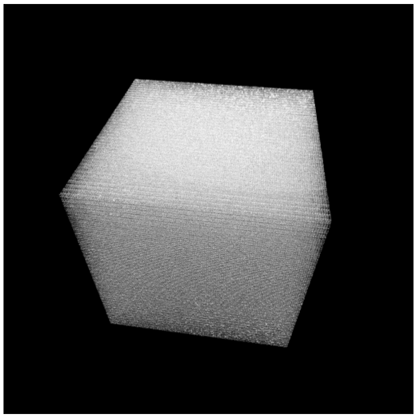

Figure 3: A Monte-Carlo simulation using a large number, $\approx 120^3$ of packed dielectric spheres using 100 samples per pixel

As was discussed above, this model is intended as a starting point, and many of the limitations give rise to a large potential body of future work. The first and foremost task is to produce a snow appearance model that is efficient, but does not lose the effect of small scale detail, or light transport accuracy provided by the present model. A second goal for future models, as mentioned above, is to handle arbitrary lighting conditions.

A third, and an unfortunately neglected characteristic in appearance models in general, is scalability. Most materials by nature look substantially different when viewed up close, and a good appearance model should reflect this. Snow is a quintessential example of this phenomenon, and indeed, this paper is meant to address the gap at one end. However, a bridge that spans both levels of detail is necessary for producing an accurate model with tractable efficiency. General advances on scalable appearance models would be a great boon to both the research community and industry.

Another component for future work is to decouple the high level accumulation model with the small scale surface detail. Since it was discovered that an accumulation model greatly improved the realism not only by providing plausible geometry, but also proved to be a good method for producing desired properties for sphere packing. Of course, our method would be infeasible for large scenes, but it is conceivable that a hybrid method could be use techniques like those in [Fearing 2000] in conjunction with the techniques presented on a level-of-detail base.

One final component of future work that should be addressed is the special needs for tone mapping in snow. While the results given in this paper are plausible, they are somewhat “plain” since they do not have the glints and highlights often associated with snow. Although these highlights are often a result of refreezing and other artifacts of snow metamorphosis, and are rarely visible on fresh pure snow, they are a still a distinguishing characteristic. The problem in simulation is that these glints must be brighter than snow, which is already one of the brightest naturally occurring substances. The glints are in fact brighter because they are usually a direct reflection of sunlight, and thus would not occur under diffuse illumination. It is likely to be necessary then, for special needs to be taken to the tone mapping of snow with glints in order to produce an image consistent with the response of our visual system.

We have presented a general overview of the optical properties of snow, from which we proposed general parameterized model for the appearance model of snow, coupled with a simple accumula-

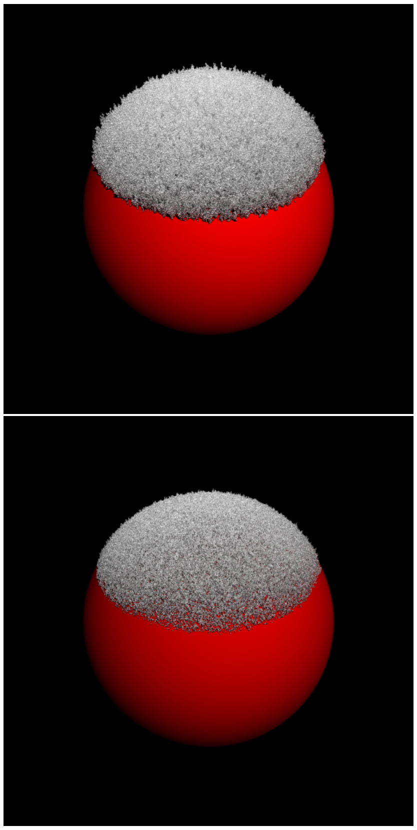

Figure 4: A Monte-Carlo simulation of a snowy sphere under thin and thick snow cover

tion model. We have presented a Monte-Carlo based technique for implementing the parameterized model, which is tractable (if not fast) for small scenes on modern hardware, and produces images of snow at a level of realism surpassing previous methods.

## References

Bohren, C. F., and Barkstrom, B. R. 1974. Theory of the optical properties of snow. *Journal of Geophysical Research 79*, 30 (October), 4527–4534.

Bohren, C. F. 1983. Colors of snow, frozen warterfalls, and icebergs. *Journal of the Optical Society of America 73*, 12 (December), 1646–1652.

Fearing, P. 2000. Computer modelling of fallen snow. In *Siggraph 2000, Computer Graphics Proceedings*, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, K. Akeley, Ed., 37–46.

Jensen, H. W., and Buhler, J. 2002. A rapid hierarchical rendering technique for translucent materials. In *Proceedings of SIGGRAPH 2001*.

Jensen, H. W., Marschner, S. R., Levoy, M., and Hanrahan, P. 2001. A practical model for subsurface light transport. In *Proceedings of SIGGRAPH 2001*, 511–518.

Kokhanovsky, A. A., and Zege, E. P. 2004. Scattering optics of snow. *Applied Optics 43*, 7 (March), 1589–1602.

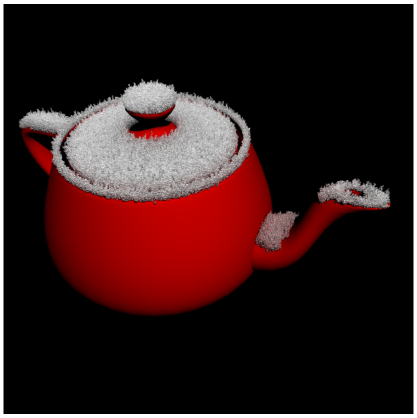

Figure 5: A Monte-Carlo simulation of a snowy teapot under thick snow cover

Meirold-Mautner, I., 2004. A physical snow-radiation model: Measurements, model development, and applications to the snow ecosystem.

Nishita, T., Iwasaki, H., Dobashi, Y., and Nakamae, E. 1997. A modeling and rendering method for snow by using metaballs. *Computer Graphics Forum 16*, 3, C357–C364.

Sean A. Twomey, C. F. B., and Mergenthaler, J. L. 1986. Reflectance and albedo differences between wet and dry surfaces. *Applied Optics 25*, 3 (February), 431–437.

Warren, S. G. 1982. Optical properties of snow. *Rev. Geophys. Space Phys. 20*, 1 (February), 67–82.

Wiscombe, W. J., and Warren, S. G. 1980. A model for the spectral albedo of snow i: Pure snow. *Journal of the Atmospheric Sciences 37* (August), 2712–2733.