FEATURE ARTICLE

Rendering the Bluish Appearance of Snow: When Light Transmission Matters

Petri M. Varsa and Gladimir V. G. Baranoski, *University of Waterloo, Waterloo, ON, N2L 3G1, Canada*

Material appearance is largely determined by complex light attenuation processes. The distinct bluish colorations that can be observed when light is transmitted through snow are among the most striking outcomes of these processes. In this article, we present a method for the predictive rendering of this phenomenon taking into account the variability of snow's physical and morphological characteristics. To achieve that, we employ an approach centered on the effective use of spectral transmittance data obtained using a first-principles light transport model for snow. The suitability of the proposed method to rendering applications is illustrated through the synthesis of images depicting the bluish appearance of snow under different illumination conditions.

The rendering of material appearance changes brought about by natural phenomena remains an active area of computer graphics research. The bluish appearance that can be observed when light is transmitted through snow (Figure 1) is arguably among the most fascinating of such changes. In concise terms, it results from the wavelength-dependent attenuation of the transmitted light by the ice crystals that form this granular material.[1] Despite the existing knowledge about its eliciting processes, however, its predictive depiction in computer generated images remains elusive. This open problem has served as the main motivation for the work described in this article.

Snow is often portrayed in feature films, video games, and computer animations. Many scenes in such media make use of illumination configurations, where the light source is obscured by this ubiquitous material. Viewed in this context, the predictive rendering of its bluish appearance elicited by light transmission would enhance the realism of those scenes while mitigating the need for manual tweaks.

Clearly, the predictive reproduction of this distinctive display of a natural phenomenon requires careful simulations of the underlying light attenuation processes. These simulations, in turn, rely on the use of detailed light transport algorithms controlled by the wavelength-dependent optical properties of the target material. We note that academically and commercially available image synthesis pipelines often adopt simplifying assumptions about these properties, notably with respect to fundamental materials like water and ice. These simplifications include, for instance, setting the real component of their index of refraction to a constant value across the visible portion of the light spectrum, and considering its imaginary component (the extinction coefficient) negligible in this spectral domain. Although this strategy can reduce rendering costs, it hinders the predictive rendering of material appearance changes with a markedly noticeable wavelength dependence.

Recently, we developed a first-principles light transport model for snow, known as SPLITSnow, whose predictive capabilities have been extensively evaluated in the visible and infrared spectral domains.[2] This model was designed to support snow-related multidisciplinary investigations and applications in fields as diverse as remote sensing, environmental sciences, and computer graphics, just to name a few. It can be employed to compute high-fidelity radiometric data (e.g., spectral reflectance and transmittance) for snow under different illumination and environmental conditions. As a stochastic model, however, SPLITSnow employs many trials to produce asymptotically convergent simulation results. This aspect makes it unsuitable for use on the fly during rendering applications demanding high interactivity rates.

To overcome the simulation constraints outlined above, we present a practical method for the predictive rendering of the bluish appearance of snow. This method is based on the effective use of snow

Digital Object Identifier 10.1109/MCG.2023.3307517
Date of publication 29 August 2023; date of current version 25 January 2024.

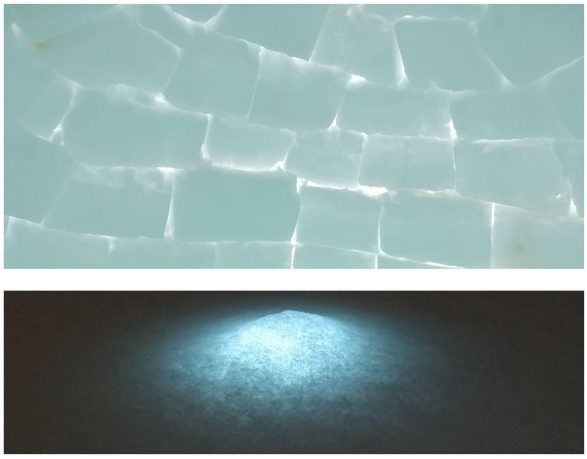

FIGURE 1. Photographs depicting the bluish appearance observed when light is transmitted through snow. Top: daylight transmitted through an igloo wall formed by snow blocks (courtesy of Muntart). Bottom: light transmitted from a garden lamp covered by snow at dusk (courtesy of Gabriel E. C. Baranoski).

transmittance data. It takes into account not only variations in light incidence, but also in the physical and morphological characteristics of snow since all these factors can significantly affect the material's transmittance, particularly within the visible spectral domain.[3]

We note that measured transmittance datasets for snow are scarce in the scientific literature. Furthermore, existing datasets are limited to a small number of snow samples and light incidence geometries.[3] These limitations prompted us to search for a viable alternative to obtain such data. Accordingly, within the framework of the proposed method, we chose to employ the SPLITSnow model to form a simulated spectral transmittance database, henceforth simply referred to as the *transmittance database*. We then used principal component analysis (PCA) (see the "Principal Component Analysis," sidebar) to obtain a compact representation of this database. This representation, which can be extended to account for other snow characterization parameters as needed, further contributes to the low computational costs associated with the incorporation of the proposed method in standard image synthesis pipelines.

The contributions of this work can be summarized as follows. It tackles an elusive topic relevant for two areas that have been focal points of computer graphics research and applications since its early days, namely material appearance modeling and natural phenomena simulation. We remark that despite snow being a ubiquitous material, the predictive depiction of its bluish appearance elicited by light transmission has not been explicitly addressed in the computer graphics literature to date. Thus, by filling this gap, this work is expected to play a part in the achievement of a higher level of realism in the predictive rendering of scenes portraying this material, while keeping the need for manual adjustments low. Moreover, to address this topic, we employed an approach based on the effective use of a spectral database compactly represented using PCA. As such, our work can also be seen as proof of concept for using such an approach in the rendering of other complex phenomena elicited by light transmission. Finally, it has long been recognized that fruitful research requires both theory and data. In line with this guideline, future investigations and applications involving the rendering of natural scenes can benefit from the availability of the spectral database generated for this work.

BACKGROUND

Snow consists of ice grains suspended in a medium (pore space) usually occupied by air. These grains accumulate in a range of sizes and can take on a variety of forms.[4] Snow grains can be small, with fine grains being less than 0.2 mm in size, or coarse, with snow grains being 1 mm in size. The grains may be formed of highly faceted crystals or present rounded forms. In the latter case, the surface itself may be covered by microfacets (Figure 2). In this work, we will refer to the quantity of facets on snow grains as *facetness*. Independent of these two snow grain characteristics, the snowpack itself may vary in density[5] (from 100 to 450 kg m^{-3}) and thickness. Spectral curves of light transmitted through snow are affected by all these parameters.[3]

As light penetrates through a snowpack, reflection and refraction events cause the photons (commonly represented by light rays) to be scattered.[1] These events increase the optical path length traversed by the rays. The light and grain interactions along the increased path lengths present several opportunities for

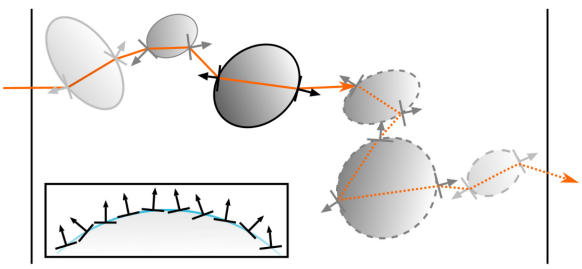

FIGURE 2. Schematic representation of light traversing a snow sample composed of rounded grains. Inset: Surface details represented by microfacets associated with their regular normal vectors.

SIDEBAR: PRINCIPAL COMPONENT ANALYSIS

We applied Principal Component Analysis (PCA)[S1] to reduce the dimensionality of the transmittance database that contains the spectral light transmittance curves obtained for snow. This database consists of a set of 24,750 spectral curves, which quantify the relative spectral power distribution of light transmission through snow for various samples, thicknesses, and angles of incidence. For the purpose of this work, it is important to note that the database can be represented as an $s \times l$ matrix M, where each row s contains the values of a spectral curve in the database, and each column l is associated with a particular wavelength of light λ.

Although several different techniques can be used to implement PCA, we elected to employ singular value decomposition (SVD) due to its numerical stability.[S2] More specifically, we used the version suggested by Pratt.[S1]

SVD factors M into three matrices $M = U\Sigma V^T$, where U is an $s \times s$ orthonormal matrix, Σ is an $s \times l$ matrix whose off-diagonal elements are zeros, and V is an $l \times l$ orthonormal matrix. For $s \geq l$, observe that the values of Σ are the *singular values* of M and they satisfy the condition

$$\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_l \geq 0 \tag{1}$$

where each σ_i is the diagonal value for the ith row of Σ.[S2] We then compute the $s \times l$ matrix $C = MV$. Using this technique, V contains a new set of basis vectors (as column vectors) and C stores the coefficients of each sample, relative to the new basis V.

To reduce the size of the spectral database, we select $l' \ll l$ columns from the matrices C and V. Hence, C' is an $s \times l'$ matrix and V' is an $l \times l'$ matrix. To reconstruct the spectrum for a sample at row i, a $1 \times l'$ matrix, C'_i is created by selecting the ith row from C'. Multiplying $C'_i \times V'^T$ yields the reconstructed spectrum. Note that, although there are many implementation strategies for PCA, the implementation of PCA used in this work aims to minimize the Euclidean norm of the difference between C and C'.[S2]

The results of our application of PCA in this work are exemplified in Figure S1. More precisely, this figure depicts the reconstructed spectral transmittance curves for each of the 225 virtual snow samples employed in this work. These curves were computed considering a snowpack thickness equal to 15 cm and an angle of incidence of 0°. We remark that this number of samples corresponds to the combination of nine different grain sizes, five different degrees of facetness, and five densities.

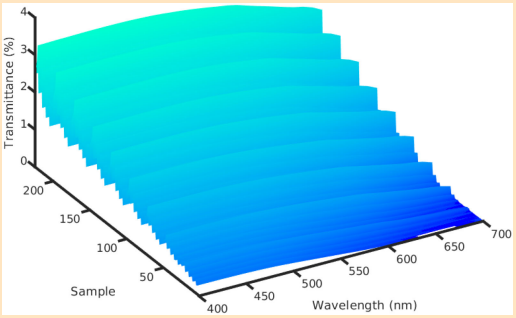

FIGURE S1. Reconstructed spectral transmittance curves for each of the 225 snow samples employed in this work. The plotted spectral curves were computed considering a snowpack thickness equal to 15 cm and an angle of incidence of 0°.

Recall that the chromatic attributes of a given material can be obtained from the convolution of a selected standard illuminant's spectral power distribution, the material's spectral responses, and the spectral responses of the human photoreceptors.[16] Accordingly, we carried out tests to determine the minimum number of principal components (values of l') appropriate to obtain a high fidelity to cost ratio in the reconstruction of the spectral curves, and consequently in the resulting material colors. These tests consisted of the computation of the perceptual difference (Euclidean distance in the CIELAB color space)[20] between the chromatic attributes associated with each reconstructed curve, and those attributes associated with the corresponding reference (including all components) spectral curve. We then selected the smallest number of principal components ($l' = 2$) leading to perceptual differences below 2.3, which is the experimentally determined perceptibility threshold for an average human observer.[20] Incidentally, this number of components also resulted in root mean square errors[16] between reconstructed and reference spectral curves below 0.002.

REFERENCES

S1. W. K. Pratt, *Digital Image Processing*, 4th ed. Hoboken, NJ, USA: Wiley, 2007.

S2. I. T. Jolliffe, *Principal Component Analysis*, 2nd ed. New York, NY, USA: Springer-Verlag, 2002.

absorption events to occur. Since ice is not strongly absorptive, many rays will penetrate to a significant depth within the snowpack. However, as ice preferentially absorbs wavelengths at the red end of the visible spectrum,[6] fewer photons within this spectral region will penetrate to greater depths. This process can ultimately lead to the observed bluish appearance of snow.[1]

RELATED WORK

The modeling, animation, and rendering of snow have long been topics of interest for the computer graphics community. We note that, although substantial efforts have been directed toward snow accumulation[7,8] and rigid body interactions with snow,[9,10] few works have addressed its appearance, the topic with the strongest connection to this work.

An early investigation to directly represent the appearance of snow was conducted by Nishita et al.[11] They calculated the optical path of light through snow, taking into account scattering and absorption events, in order to compute total light intensity. Their work, however, did not explicitly allow for the computation of spectral reflectance and transmittance curves for snow samples. Subsequently, Ohlsson and Seipel[12] implemented a shading technique to compute a snow color function for real-time applications.

Later, Frisvad et al.[13] implemented an analytic model to compute the scattering properties of various materials. These include ice with briny water and air inclusions that act in suspension of a pure ice host material. Their adaptation of Lorenz–Mie theory represents nonspherical particles with the quantity of equivalent spheres that preserves the volume-to-area ratio of the nonspherical particles, and employs a constant for the real part of the index of refraction. However, they are able to attain a bluish hue by incorporating the complex part of the index of refraction for water and ice into their calculations. They parameterize their solution by using brine and air fractions and suggest that algae and minerals could also be included. Similarly, Jarabo et al.[14] made use of the optical properties discussed by Frisvad et al. in their own work. In addition to these two works, Meng et al.[15] implemented a multiscale level of detail technique to render granular materials. Among other materials, they also render images of snow at various scales.

We note that there are a number of computer graphics works, such as those mentioned earlier,[13,14,15] that propose models aimed at the rendering of a relatively broad range of materials (e.g., milk, ice, sugar, sand, and snow) designated as either particulate or granular. Due to their intended generality, however, the parameter space and formulation of their proposed models lack the specificity required to guide the predictive reproduction of spectral responses of snow samples with distinct characteristics. In fact, to the best of our knowledge, the output of snow spectral reflectance and transmittance quantities, let alone their direct comparison with measured data, has not been reported in the computer graphics literature to date.

We remark that our goal is to enable not only believable, but also predictable depictions of snow appearance variations elicited by light transmission. Accordingly, we elected to employ a first-principles spectral model for snow[2] specifically developed to predictively output radiometric data for this material. This aspect is further addressed in the next section.

TRANSMITTANCE DATABASE

In this work, we considered virtual snow samples with nine different grain sizes, five different degrees of facetness, and five densities, yielding 225 unique snow exemplars. In addition to these three sample characteristics, light transmission through the material is also affected by other factors, notably the thickness of the snowpack and the angle of light incidence. To account for these two factors, we also considered eleven different snowpack thicknesses and ten distinct angles of incidence. Hence, these two factors plus the three sample characteristics formed the basis of five dimensions of snow variability totalling 24,750 spectral curves. Finally, for each transmittance curve, we considered a spectral resolution of 10 nm, yielding 31 wavelengths in the visible spectrum (400–700 nm). Table 1 lists the specific values considered for each parameter.

We remark that the transmittance curves in the transmittance spectral database were computed using the SPLITSnow model, which makes use of geometric (ray) optics to simulate light interactions with snow grains. For the computation of each curve in the transmittance database, 10^5 rays (per sample wavelength) were employed in order to attain asymptotically convergent radiometric quantities.[16]

We then applied PCA to reduce the dimensionality of the transmittance database that contains the spectral light transmittance curves obtained for snow. For the results presented in this work, only two principal components were necessary to reproduce the original spectral curves without decreasing the fidelity of the resulting images. For details regarding our implementation of PCA, please refer to the "Principal Component Analysis," sidebar.

It is worth mentioning that the proposed method can employ transmittance data provided by different models. The use of the SPLITSnow model for this

TABLE 1. Parameter values used to create the transmittance database of spectral curves.

Parameter	Values
Mean grain size (μm)	100, 150, 200, 250, 300, 350, 400, 450, 500
Facetness (unitless)	0.1, 0.3, 0.5, 0.7, 0.9
Density (kg m^{-3})	200, 250, 300, 350, 400
Thickness (cm)	1, 1.5, 2, 3, 5, 7, 10, 15, 25, 35, 50
Angle of incidence (degrees)	0, 10, 20, 30, 40, 50, 60, 70, 80, 89

purpose was determined not only by its specificity and availability, but also by its predictive capabilities.[2] These provide the empirically supported foundation for this work. For more information regarding these capabilities, please refer to the "Snow Spectral Responses," sidebar.

SPECTRAL RESOLUTION

In order to render images, we employed the completely reconstructed spectra (31 spectral samples), rather than tristimulus values (three spectral samples). This choice was based on the importance of spectral information for the rendering of real world materials.[18]

To illustrate this aspect, we generated swatches for a given snow sample through the convolution of selected illuminants' spectral power distributions [see Figure 3 (top)], the sample's transmittance curve [see Figure 3 (middle)] and the broad spectral responses of the human photoreceptors.[19] This last step was performed using a standard CIEXYZ to sRGB color system conversion procedure[16] and considering the CIE standard illuminants E (equal energy), D55 (5,500 K daylight), D65 (6,500 K natural daylight), LED-B5 (phosphor-type LED #5), and FL-15 (fluorescent type 15).[17] As a final step, we applied a grayscale texture to present the obtained colors on the swatches [see Figure 3 (bottom)].

The swatches shown in Figure 3 (bottom) were generated using the convolution procedure outlined earlier. Those in the top row were obtained using only three sample wavelengths (455, 551, and 608 nm), which correspond to the chromaticity coordinates of standard Society of Motion Picture and Television Engineers monitors.[16] The swatches in the bottom row depict the same snow samples generated using 31 sample wavelengths. It can be observed that there is an noticeable difference between the swatches obtained using tristimulus values and those obtained using 31 samples taken across the visible spectrum.

By visually inspecting the swatches shown in Figure 3 (bottom), it can also be noted that the choice of illuminant can have a significant impact on the depicted colorations. Nonetheless, examining the swatches generated using a flat, equal energy light source (illuminant E), one can verify that the colorations are still largely dependent on the transmittance values.

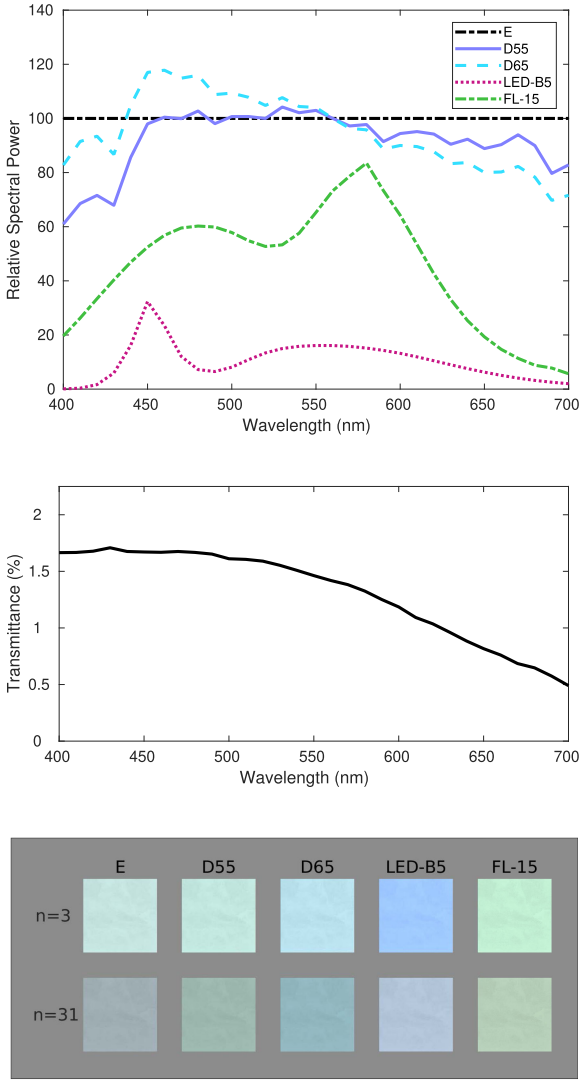

FIGURE 3. Components of a convolution process employed to generate swatches for a backlit snow sample characterized by a mean grain size, facetness, thickness, and density equal to 400 μm, 0.3, 15 cm, and 300 kg m^{-3}, respectively. Top: plots of the relative spectral power distributions of selected CIE illuminants (E, D55, D65, LED-B5, and FL-15).[17] Middle: transmittance curve computed for the snow sample considering an angle of incidence of 0°. Bottom: snow swatches generated using different number (n) of wavelengths, namely 3 (455, 551, and 608 nm) and 31 (equally spaced in the 400–700 nm interval).

SIDEBAR: SNOW SPECTRAL RESPONSES

In order to establish a physical basis supported by empirical evidence, we elected to employ in this work a specific light transport model, SPLITSnow, whose predictive capabilities have been comprehensively verified through comparisons with ground truth data.[2] To further illustrate such capabilities, we present here additional examples (Figure S2) of quantitative and qualitative agreement between its predictions and *in situ* collected radiometric information for snow.

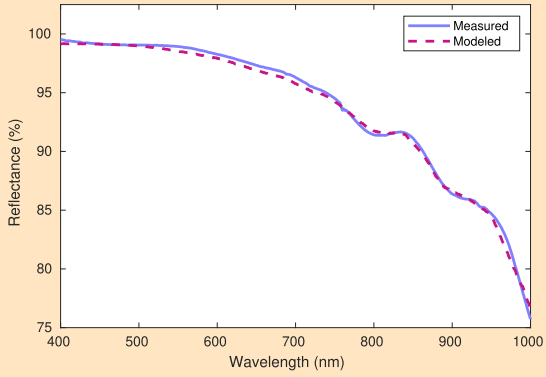

(a)

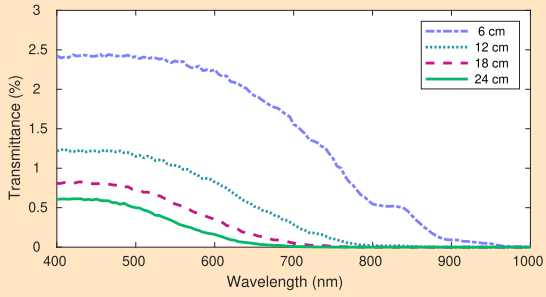

(b)

FIGURE S2. Comparison of modeled snow spectral responses with measured data and empirical observations reported in the literature. (a) Modeled reflectance data compared with measured data provided by Salvatori et al.[S3] (b) Modeled transmittance data depicting nonlinear trends reported by Robson and Aphalo.[3] The modeled curves were obtained using SPLITSnow[2] and the parameter values presented in Table S1.

Without loss of generality, we chose to use as a quantitative reference for the provided examples a measured reflectance dataset (SISpec spectral file #6) provided by Salvatori et al.[S1] This dataset was collected at the Kulmodden sampling site on the Svalbard island in the Arctic at an elevation of 35 m above sea level.

TABLE S1. Parameter values employed in the characterization of the snow sample considered in the computation of the modeled radiometric curves presented in Figure S2.

Parameter	Value
Grain size range (μm)	200–300
Temperature (°C)	−5
Thickness (cm)	18
Water saturation (%)	0
Density (kg m^{-3})	360
Facetness range	0.05–0.25
Facetness mean	0.15
Facetness standard deviation	0.075
Sphericity range	0.7–0.9
Sphericity mean	0.8
Sphericity standard deviation	0.1

The selection of this measured dataset also took into account the fact that the nivological characterization of the target sample was recorded along with its spectral reflectance. This supporting information was then employed to guide the assignment of appropriate values (presented in Table S1) for the model parameters. This use of well-grounded parameter values, in turn, underscores the close agreement observed between the modeled and measured reflectance curves depicted in Figure S2(a).

In Figure S2(b), we present transmittance curves obtained by varying the thickness of the snow sample considered in computation of the modeled reflectance curve depicted in Figure S2(a). More precisely, except for the thickness, all other parameter values considered in the computation of the transmittance curves correspond to the values provided in Table S1. As it can be observed in these curves, the spectral transmission of blue (≈400–500 nm) and green (≈500–570 nm) light are greater than that of red or far-red light. Furthermore,

the increase in the sample's thickness led to a nonlinear reduction in its transmittance across the spectral region of interest. These predictable behaviors have been reported by Robson and Aphalo in their field investigation of light transmission across a seasonal snowpack.[3]

For consistency, we adopted a spectral resolution of 5 nm in all modeled radiometric curves presented in Figure S2, and considered the same angle of incidence of 0° employed in the *in situ* measurements conducted by Salvatori et al.[S3] Lastly, to enable the straightforward reproduction of the presented modeled curves, we have made the SPLITSnow model available for online use.[S4]

REFERENCES

S3. R. Salvatori, R. Salzano, M. Valt, R. Cerrato, and S. Ghergo, "The collection of hyperspectral measurements on snow and ice covers in polar regions (SISpec 2.0)," *Remote Sens.*, vol. 14, no. 9, pp. 2213:1–14, 2022.

S4. Natural Phenomena Simulation Group, *Run SPLITSnow Online (Spectrometric Mode)*, 2020. [Online]. Available: http://www.npsg.uwaterloo.ca/models/splitsnow.php.

It is worth noting that, depending on the spectral characteristics of a given illuminant, it may be possible to employ less than 31 equally spaced spectral samples and still maintain the visual fidelity of the generated swatches. However, in order to determine the ideal number and location of spectral samples for a given illuminant, one may need to employ a perceptually based approach similar to that proposed by Kravchenko et al.[20] We intend to explore this possibility in our future investigations in this area.

METHOD

The spectral curves in the transmittance database are incorporated into a typical rendering pipeline in order to compute the color exhibited when light is transmitted through snow in winter scenes. To this end, scene geometry that is identified as a snow material is considered to be translucent in nature. Thus, snowy objects that are hit by a primary ray are treated as (potentially) emissive, with light permeating through the material. The emissive color is retrieved from the compressed version of the transmittance database, with the specific curve being determined by the characterization of the snow sample and the local geometry of the object in relation to a traditional light source.

We remark that the focal point of this work is on appearance changes elicited by light transmitted through snow. Thus, although the proposed method can be extended to also employ snow reflectance data, for simplicity and conciseness, the light reflected by snowy objects was accounted for using a standard local illumination model for diffuse materials.

The implementation of the proposed method provides a new material definition class for a typical path tracing renderer. In the remainder of this article, we will refer to this class as *PCASnow*.

When an object associated with the PCASnow material class is hit by a ray, the distance d travelled inside the material and the angle of incidence θ with respect to the light source (Figure 4) are used in combination with the three snow sample properties

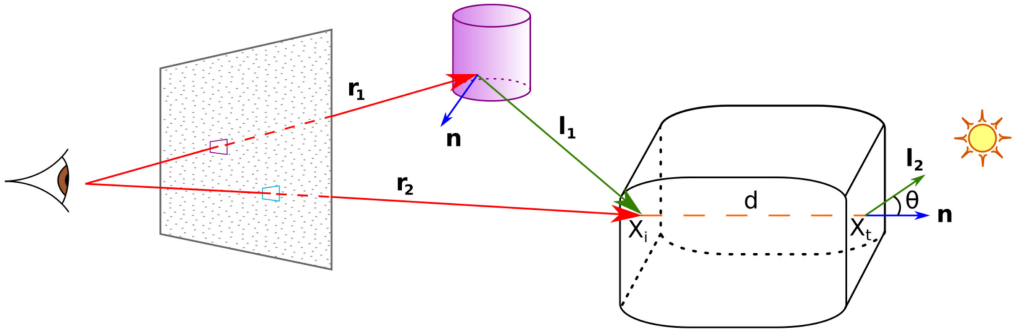

FIGURE 4. Sketch illustrating key parameters employed in the incorporation of the PCASnow material class into a path tracing algorithm. When a ray hits an object, such as a snow block, associated with the PCASnow material class, the spectral distribution of the light transmitted through this object is obtained from the transmittance database employing the distance d travelled inside the material (which corresponds to the thickness t of the reference snow sample) and the angle of incidence θ. Note that $\mathbf{r}_1$ and $\mathbf{r}_2$ represent primary rays, while $\mathbf{l}_1$ and $\mathbf{l}_2$ represent secondary rays.

(facetness, grain size, and density) in order to determine the spectral light distribution that has been transmitted through the snow. To index the compressed database, the three snow sample properties, as well as the sample thickness (represented by the distance d) at the location of interest, and the angle of incidence (represented by θ) must be known. Again, since our goal is to reproduce the bluish appearance observed when light is transmitted through snow, we concentrate our attention on the spectral distribution of light transmitted through snow samples, and account for its spatial distribution implicitly.

Algorithm 1. Pseudocode of the algorithm used to determine the spectral contribution of light transmitted through snow. In Figure 4, this spectral value may be used by the primary ray $\mathbf{r}_2$, or the secondary ray $\mathbf{l}_1$.

```
▷ Initialization of the Material ◁
▷ facetness (f), grain size (gs), and density (den) ◁
M ← new PCASnow(f, gs, den)
G ← new SceneObject(M)        ▷G decorated with M
▷ Light source, L, is a traditional light source ◁
L ← new Light(position)

▷ Indexing of Transmittance Database ◁
▷ Ray-object intersection at X_i ◁
d ← computeDistanceThroughObject(X_i, G)
X_t ← determineExitPoint(X_i, d)
θ ← computeAngleToPrimaryLightSource(L, X_t)
▷ Use f, gs, den, d, and θ as index values ◁

▷ Interpolation of Spectral the Responses ◁
▷ ⌊floor⌋ and ⌈ceiling⌉ functions denote lower and
  upper index values ◁
spectra_ll ← M[f, gs, den, ⌊d⌋, ⌊θ⌋]
spectra_hl ← M[f, gs, den, ⌈d⌉, ⌊θ⌋]
spectra_l ← cubicSpline(spectra_ll, spectra_hl)
spectra_lh ← M[f, gs, den, ⌊d⌋, ⌈θ⌉]
spectra_hh ← M[f, gs, den, ⌈d⌉, ⌈θ⌉]
spectra_h ← cubicSpline(spectra_lh, spectra_hh)
spectra ← linear(spectra_l, spectra_h)
return spectra
```

Referring to Figure 4 to schematically illustrate the aforementioned light and material interactions, a primary ray $\mathbf{r}_1$ hits an object in the scene. To compute the object's reflectance, secondary ray $\mathbf{l}_1$ is cast toward a block of snow, which indirectly serves as a light source. Similarly, a primary ray $\mathbf{r}_2$ strikes the snow block directly. For $\mathbf{r}_2$, both the reflectance and the emissive components contribute to the shading.

To compute the spectral appearance attributes for geometric objects whose colors are affected by light transmitted through snow, there are three aspects of the PCASnow material class implementation to be considered: the snow characterization specification, the retrieval of the required spectral curves, and the interpolation between the retrieved curves representing the discrete samples. (See Algorithm 1.)

To address the first aspect, the *initialization of the material*, the sample characterization values must be specified when initializing an instance of a PCASnow material class. More precisely, the facetness, grain size, and density must be provided as parameters to the PCASnow material class instantiation.

As for the second implementation aspect, *indexing of the transmittance database*, it requires correct indexing into the matrix containing the spectral curves. Our implementation computed the row index using an elementary array indexing technique. Since there are five dimensions to consider, index retrieval is akin to a 5-D array index computation. The indexing order is arbitrary, so long as the indexing of retrieval is consistent with the indexing used when creating and compressing the transmittance database with PCA. We chose to index first by the grain properties (facetness and grain size), next by the snow sample's density, and then by the geometric properties required for rendering (the angle of incidence and sample thickness). This strategy ensures that all spectral transmittance curves for each sample are contiguous in memory. This organization, in turn, allows for the set of transmittance spectra associated with a given snow sample to be cached if desired.

Subsequently, to address the third implementation aspect, the *interpolation of the spectral responses*, we selected ten angles of incidence and eleven sample thicknesses for each snow sample. Due to the relative coarseness of these choices, interpolations between angles of incidence and sample thicknesses were required. We chose to first use cubic spline interpolation between thicknesses, interpolating between the thicknesses for both the relevant angles of incidence. We then interpolate linearly between these two curves, which represent the low and high angles of incidence for the cubically interpolated thicknesses. The former (cubic) interpolation scheme was selected due to the reported nonlinear attenuation of light traversing snow samples.[3] Note that the interpolated spectra are convolved with the relative spectral power distribution of the light transmitted through the snow block before it is used in local reflectance calculations, or added as an emissive contribution.

The boundaries adopted for the minimum and maximum angles of incidence were 0° and 89°. The

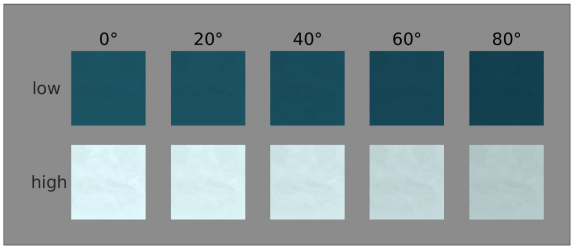

FIGURE 5. Swatches generated for a selected snow sample considering different values for its grain size and distinct angles of incidence. Top row: low grain size value (100 μm). Bottom row: high grain size value (500 μm).

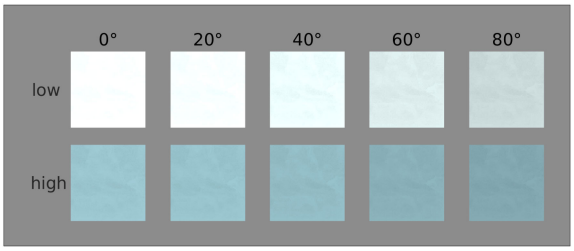

FIGURE 7. Swatches generated for a selected snow sample considering different values for its density and distinct angles of incidence. Top row: low density value (200 kg m^{-3}). Bottom row: high density value (400 kg m^{-3}).

spectra for any angle of incidence outside these bounds were assumed to exhibit no color. Note that the upper limit of the angle of incidence was set to 89° to avoid numeric computation issues with incident light that is parallel to a surface. Similarly, the boundaries for the snow thicknesses were chosen to be 0 and 100 cm. Transmittance spectra through a 0 cm snow thickness were assumed to have the maximum magnitude and assigned the color white. On the other hand, transmittance spectra through snow thicknesses beyond 100 cm were assumed to be negligible and assigned the color black.

RESULTS

Snow Characterization Variability

To illustrate the effects that different material characterizations have on the bluish appearance of snow, we initially generated swatches (Figures 5 to 8) for a selected sample considering variations in its parameters and distinct angles of light incidence. For this sample, except for the parameter under test, we set its grain size, facetness, density, and thickness to 550 μm, 0.3, 275 kg m^{-3}, and 12 cm, respectively. The swatches were obtained using the D65 illuminant.[17]

By visually inspecting the obtained swatches presented in Figures 5 to 8, we observe that high values for the grain size led to markedly lighter hues, while high values for the facetness, density, and thickness led to darker hues. Moreover, we note that the variation in the grain size parameter resulted in more prominent appearance differences. Also, in all instances, the impact of the snow characterization parameters is modulated by the light incidence geometry, with an increase in the angle of incidence leading to darker hues. This aspect can be explained by the fact that greater angles of incidence result in more light being reflected on a snow sample surface, and less light being transmitted through its volume.[2] It is important to mention that these behaviors are consistent with qualitative trends identified by *in situ* investigations on the transmission of light through snowpacks.[3] For more information regarding related observations reported in the literature, please refer to the "Snow Spectral Responses" sidebar.

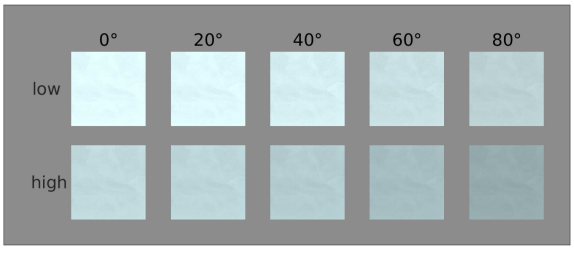

FIGURE 6. Swatches generated for a selected snow sample considering different values for its facetness and distinct angles of incidence. Top row: low facetness value (0.1). Bottom row: high facetness value (0.9).

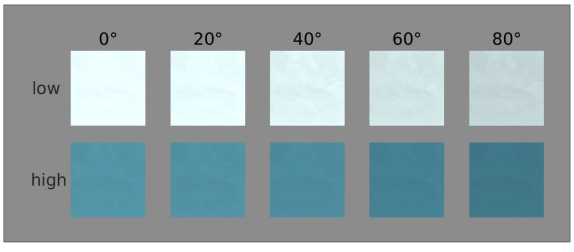

FIGURE 8. Swatches generated for a selected snow sample considering different values for its thickness and distinct angles of incidence. Top row: low thickness value (10 cm). Bottom row: high thickness value (25 cm).

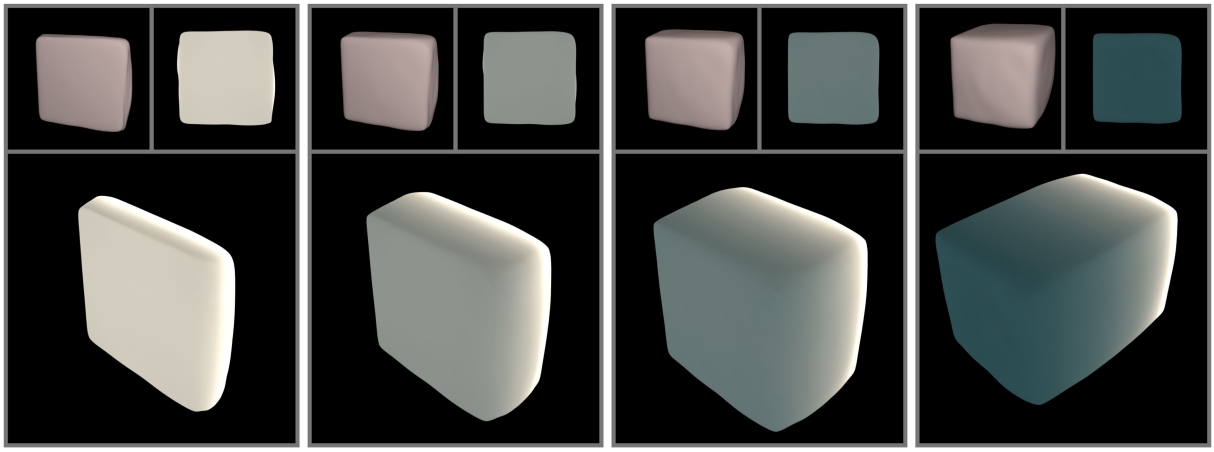

FIGURE 9. Sequence of rendered images depicting light interactions with four different blocks of snow. From left to right the snow blocks are approximately 5, 10, 15, and 30 cm thick. Each sequence component contains three images. Upper-left image: the block geometry is rendered using a standard local illumination model and a point light source situated near the eyepoint. Upper-right image: the geometry is rendered using the proposed method, with the point light source positioned behind the snow block and the viewing direction perpendicular to it. Bottom image: the geometry is also rendered with the proposed method, but with the viewing direction oblique to the snow block.

Image Synthesis Applications

To showcase the use of the proposed method in image synthesis applications, we present in Figure 9 a sequence of images rendered using the described modified path tracer. Each group of images depicts appearance changes elicited by light transmission through snow blocks of distinct thicknesses. To render these images, the blocks are backlit by an unseen point light source with the relative spectral power distribution of the D65 illuminant.[17] In order to clearly depict the coloration changes associated with volume thickness, no textures were employed in the generation of these images. The PCASnow material class was initiated using a snow sample with a mean grain size of 400 μm, a density of 300 kg m^{-3}, and a moderately low facetness of 0.3.

The oblique-angled image in Figure 9 (right), depicting the 30 cm-thick snow block, took less than 1 min to render on an AMD Ryzen Threadripper 3990X Processor with 64 computing cores. This image has a resolution of 1080×1080 pixels. It was rendered using 100 rays per pixel, and it involved approximately 90 million iterations of the proposed method's core algorithms. For comparison purposes, the time required to conduct a full light transport simulation to obtain a *single* spectral transmittance curve for a 30 cm-thick snow sample with the same physical and morphological characteristics and utilizing the same CPU is approximately 7.5 min.

To further illustrate the applicability of the proposed method, we generated two additional demonstration scenes with the same path tracer. The first scene, presented in Figure 10, was rendered considering the virtual camera positioned inside an igloo formed by snow blocks whose thicknesses are approximately 15 cm, and a light source located outside of it. In Figure 10 inset, we provide an alternative visualization of the geometries employed to represent the snow blocks. The inset image was rendered using a standard local illumination model employed for opaque objects and considering a point light source placed inside of the igloo. To extend our scope of

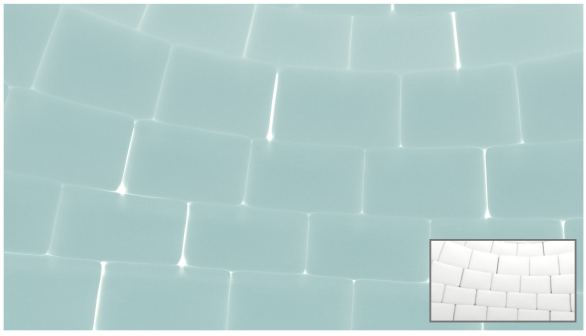

FIGURE 10. Rendered images depicting the inside of an igloo wall formed by snow blocks. The main image was rendered using spectra that were produced by employing the techniques discussed in this work. Inset: the same scene rendered using an interior light source and considering the blocks to be opaque.

FIGURE 11. Rendered images depicting the Stanford Dragon represented by three virtual snow samples characterized by distinct grain sizes, namely 250 (top), 100 (bottom left), and 500 μm (bottom right). The diameter of the dragon's body is approximately 30 cm. All three samples are characterized by a facetness of 0.3 and a density of 300 kg m^{-3}.

observations, the main image was generated using a directionally uniform light source with the relative power distribution of the D55 (5,500 K daylight) illuminant.[17] All light rays inside the igloo have been transmitted through the igloo wall, and handled by the proposed method, in order to account for indirect illumination effects. In addition to the bluish coloration exhibited by the light transmission, one can observe that exterior light also *bleeds* through the snow at the seams between the blocks, where the geometry is comparatively thin, thus brightening the surface.

The second demonstration scene, presented in Figure 11, depicts the Stanford Dragon rendered considering three different virtual snow samples. The geometry has been scaled so that the diameter of the body's main section is approximately 30 cm in diameter, with the thickest part of the geometry (including the front legs) being approximately 80 cm in thickness. Note that the variations of the thickness of the geometry contribute to the tonal gradations of color including white, various shades of blue, and dark tones. Each rendering is backlit using a point light source of the same intensity and utilizes the D65 standard illuminant.[17] The main (top) image illustrates more diversely the snow appearance changes associated with thickness variations affecting light transmission. The virtual snow sample considered in the rendering of this image is characterized by a facetness of 0.3, grain size of 250 μm, and a density of 300 kg m^{-3}.

According to Robson and Aphalo,[3] grain size has a large effect on the transmission of light through snow. To showcase the effects of concomitant grain size and thickness variations, we also rendered the Stanford Dragon scene using low and high values for the samples' grain size, namely 100 μm and 500 μm. The remaining parameter values were kept the same as those used to render the main (top) image. The resulting images are presented in the bottom row of Figure 11. Observe that the virtual snow sample with the larger grain size (right) allows more light to be transmitted and exhibits a more vivid blue hue than the main (top) image. This is in contrast to the virtual snow sample with the smaller grain size (left) that displays the opposite effect when compared to the main image (top).

We remark that the presented demonstration scenes (Figures 10 and 11) are intended to serve as illustrative examples of potential uses of the proposed method in realistic image synthesis applications. These applications may involve objects with distinct geometrical complexities and subjected to different illumination conditions. Furthermore, the depicted visual aspects can be regarded as a small subset of the myriad of snow appearance changes that can be predictively prompted by systematically varying the values of the material characterization parameters considered in this work.

CONCLUSION

In this article, we presented a method for the predictive rendering of the bluish appearance that can be observed when light is transmitted through snow. This phenomenon results from complex light and snow interactions. The involved light attenuation processes are implicitly accounted for by effectively employing a compact transmittance database obtained using a first-principles light transport model for snow. This approach allows for the generation of images depicting variations in the appearance of different snow samples traversed by light originating from illuminants with distinct relative spectral power distributions, while maintaining a high fidelity to cost ratio in their rendering.

There are a number of avenues for future work associated with possible extensions to the employed transmittance database. For instance, it would be worth examining the impact that the presence of water in the snow's pore space can have in its appearance when light

is transmitted through it. In the case of a significant impact, another dimension associated with the snow samples' water content could be added to the transmittance database. Furthermore, to enable the predictive rendering of other subtleties associated with the appearance of snow, the database could be extended to include reflectance data for samples with different characterizations. Finally, for artistic and scientific visualization applications, it would also be worth exploring the inclusion of snow radiometric data obtained for the infrared domain.

ACKNOWLEDGMENTS AND DATA AVAILABILITY

This work was supported by the Natural Sciences and Engineering Research Council of Canada under Grant 238337. The transmittance database can be obtained from the contact author upon request.

REFERENCES

1. C. F. Bohren, "Colors of snow, frozen waterfalls, and icebergs," *J. Opt. Soc. Amer.*, vol. 73, no. 12, pp. 1646–1652, 1983.
2. P. M. Varsa, G. V. G. Baranoski, and B. W. Kimmel, "SPLITSnow: A spectral light transport model for snow," *Remote Sens. Environ.*, vol. 2552021, Art. no. 112272.
3. T. Robson and P. Aphalo, "Transmission of ultraviolet, visible and near-infrared radiation to plants within a seasonal snowpack," *Photochem. Photobiol. Sci.*, vol. 18, no. 1, pp. 1963–1971, 2019.
4. C. Fierz et al., "The international classification for seasonal snow on the ground," UNESCO-IHP, Paris, France, Tech. Rep. 83, 2009.
5. C. F. Bohren and R. L. Beschta, "Snowpack albedo and snow density," *Cold Reg. Sci. Technol.*, vol. 1, no. 1, pp. 47–50, 1979.
6. S. G. Warren and R. E. Brandt, "Optical constants of ice from the ultraviolet to the microwave: A revised compilation," *J. Geoph. Res.-Atmos.*, vol. 113, no. D14, pp. D14220:1–10, 2008.
7. P. Fearing, "Computer modelling of fallen snow," in *Proc. 27th Annu. Conf. Comput. Graphics Interactive Techn.*, 2000, pp. 37–46.
8. N. v. Festenberg and S. Gumhold, "Diffusion-based snow cover generation," *Comput. Graphics Forum*, vol. 30, no. 6, pp. 1837–1849, 2011.
9. R. Sumner, J. F. O'Brien, and J. K. Hodgins, "Animating sand, mud, and snow," *Comput. Graphics Forum*, vol. 18, no. 1, pp. 17–26, Mar. 1999.
10. A. Stomakhin, C. Schroeder, L. Chai, J. Teran, and A. Selle, "A material point method for snow simulation," *ACM Trans. Graphics*, vol. 32, no. 4, pp. 1–9, 2013.
11. T. Nishita, H. Iwasaki, Y. Dobashi, and E. Nakamae, "A modeling and rendering method for snow by using metaballs," *Comput. Graphics Forum*, vol. 16, no. 3, pp. 357–364, 1997.
12. P. Ohlsson and S. Seipel, "Real-time rendering of accumulated snow," in *Proc. SIGRAD Conf.*, 2004, pp. 25–31.
13. J. R. Frisvad, N. J. Christensen, and H. W. Jensen, "Computing the scattering properties of participating media using Lorenz–Mie theory," *ACM Trans. Graphics*, vol. 26, no. 3, pp. 60:1–10, Jul. 2007.
14. A. Jarabo, C. Aliaga, and D. Gutierrez, "A radiative transfer framework for spatially-correlated materials," *ACM Trans. Graphics*, vol. 37, no. 4, pp. 1–13, 2018.
15. J. Meng et al., "Multi-scale modeling and rendering of granular materials," *ACM Trans. Graphics*, vol. 34, no. 4, 2015, Art. no. 49-1.
16. G. V. G. Baranoski and A. Krishnaswamy, *Light & Skin Interactions: Simulations for Computer Graphics Applications*. Burlington, MA, USA: Morgan Kaufmann/Elsevier, 2010.
17. E. Carter et al., "Colorimetry," Int. Commiss. Illumination (CIE), Tech. Rep. CIE 015:2018, 2018.
18. G. Johnson and M. Fairchild, "Full-spectral color calculations in realistic image synthesis," *IEEE Comput. Graphics*, vol. 19, no. 4, pp. 47–53, Jul./Aug. 1999.
19. R. Hunt, *Measuring Colour*, 2nd ed. Chichester, U.K.: Ellis Horwood Limited, 1991.
20. B. Kravchenko, G. V. G. Baranoski, T. Chen, E. Miranda, and S. Van Leeuwen, "High-fidelity iridal light transport simulations at interactive rates," *Comput. Anim. Virtual Worlds*, vol. 28, no. 3/4, 2017, Art. no. e1755.

PETRI M. VARSA is currently working toward the Ph.D. degree with the School of Computer Science, University of Waterloo, Waterloo, ON, N2L 3G1, Canada. His research interests include the simulation of natural phenomena with applications in remote sensing and computer graphics. He is a member of Natural Phenomena Simulation Group and the corresponding author of this article. Contact him at pmvarsa@uwaterloo.ca.

GLADIMIR V. G. BARANOSKI is a professor with the University of Waterloo, Waterloo, ON, N2L 3G1, Canada, where he has established the Natural Phenomena Simulation Group. His research interests include primarily the simulation of light and matter interactions aiming at applications in computer graphics, remote sensing and biomedical optics. Baranoski received his Ph.D. degree in computer science from the University of Calgary, Canada. Contact him at gvgbaran@uwaterloo.ca.