


Simulating the aurora

By Gladimir V. G. Baranoski*, Jon G. Rokne, Peter Shirley,
Trond S. Trondsen and Rui Bastos

We present the first computer graphics algorithm designed to simulate the aurora, a natural phenomenon of great visual beauty and considerable scientific interest. The algorithm is based on the current understanding of the physical origin of this natural display. The aurorae are mainly caused by high-energy electrons originating in the sun and entering the earth's atmosphere in narrow regions centered on the magnetic poles. These electrons collide with atmospheric atoms, which are excited to higher energy levels. The excited atoms emit rapidly varying visible light in a curtain-like volume as they return to lower energy levels, thereby creating the aurora. By simulating these light emissions along with the spatial and temporal distribution of the entering electrons, we are able to render the major visual aspects of auroral displays. The applicability of this auroral model for rendering and scientific purposes is illustrated through comparisons of synthetic images with photographs of real auroral displays. Copyright © 2003 John Wiley & Sons, Ltd.

Received: 1 August 2002; Revised 26 August 2002

KEY WORDS: natural phenomena; atmospheric effects; physically based rendering



Introduction

The earth is constantly bombarded by energy from outer space. Sunlight is the most visible and common example. Other parts of this bombardment are invisible, such as cosmic rays, X-rays and atomic particles that stream out from the sun in all directions. These atomic particles, protons and electrons constitute the *solar wind*. A small fraction of these solar wind particles interact with the earth's magnetic field. As a result of these interactions, plasma[†] particles are guided and accelerated to regions around the magnetic north and south poles. These fast energetic particles then collide with high-altitude atmospheric atoms and molecules resulting in light emission. This natural phenomenon is known as the *aurora bore-*

alis, or 'northern lights', and *aurora australis* (its less familiar southern counterpart).

The aurora is considered by many to be one of the most fascinating and mysterious of nature's displays (Figure 1). Lynch and Livingston³ appropriately observe:

Aurorae stand alone for mouth-gaping, awe-inspiring, spell-binding majesty. Their silent play of eerie color is surely one of the Mother Nature's grandest spectacles.

The impressive visual characteristics of auroral displays have fascinated writers, philosophers, poets and scientists over the centuries. Although descriptions and investigations are more prevalent in northern cultures, which live along the region known as the 'auroral oval', we can also find descriptions in the Bible and in the ancient texts of Greek, Roman and Chinese writers and philosophers.

Mankind has been puzzled by the origin of this phenomenon for centuries. The main visual aspects of the aurora, involving the excitation of atmospheric particles indirectly by the solar wind, have only been understood in the twentieth century. Although several theories regarding the aurora have been recently validated by space-born instrumentation,⁴ many questions remain concerning auroral physics. As an example, the mechanism of auroral electron acceleration by

*Correspondence to: G. V. G. Baranoski, School of Computer Science, University of Waterloo, 200 University Avenue West, Waterloo, Ontario, N2L 3G1, Canada.

E-mail: gvgbaran@cgl.uwaterloo.ca

[†]The word *plasma* is used in this context to denote a state of the matter in which the atoms are completely dissociated into nuclei and electrons,¹ forming an ionized stream of charged particles of both signs in equal number.²

Contract/grant sponsor: Natural Sciences and Engineering Council of Canada; contract/grant numbers: 691266; 238337.

Contract/grant sponsor: Canada Foundation for Innovation; contract/grant number: Project 6218.



Figure 1. Simulation of an auroral display.

natural forces constitutes one of the major mysteries confronting space plasma physics today.⁵

Because the aurorae are not often visible far from the poles, the reader might question whether an effort to simulate the aurora is worthwhile. We believe there are several reasons why such simulations should be done. For any educational or entertainment simulation of an environment typical of polar latitudes, the aurora is an obvious visual effect, and in fact it is often bright enough to navigate by or even to cast shadows. The aurora is also a strong visible effect in space, and even occurs on other planets with strong magnetic fields, such as Jupiter and Saturn. Simulating the aurora may also have great scientific value. As in any effort in the physical sciences, models can be tested by how well their simulations predict observed data. Since the aurorae have such an obvious visual manifestation, reproducing their actual visual appearance is one way to evaluate auroral theories and data. The science behind the aurora is furthermore quite crucial because of its links to plasma physics, as stated by Savage⁶:

Down here on Earth, plasmas occur only in human-made devices such as neon lights, mercury-vapor lamps and laboratory apparatus. But elsewhere in the Universe, plasmas are common. As much as 99.9 per cent of the matter in the cosmos is thought to exist

in a plasma state. This includes not only the matter in the Sun, the stars and space, but also in the outer atmosphere, the Earth's magnetic field and the solar atmosphere—all the regions that are involved in producing the aurora. The polar lights offer us a glimpse into the complex workings of the plasma universe and provide us with a natural laboratory in which it can be studied.

The primary goal of our research is to perform realistic simulations of auroral displays, incorporating as many known auroral physics concepts and data as possible, while keeping the complexity of the algorithm compatible with rendering requirements. Initially, we simulate the paths of the falling electron beams to model the dynamic nature of the auroral displays. The spectral and intensity characteristics of these displays are modeled by considering the light emissions along these paths. During the rendering the light emissions are forward mapped to the image plane to account for the view-dependent characteristics of the aurorae. Finally, the temporal variations associated with the spectral distribution of the light emissions are taken into account by applying antialiasing or 'blurring' techniques.

In short, in this paper, which is an extended and upgraded version of a conference presentation,⁷ we present an algorithm that simulates the stochastic processes associated with the auroral emissions in order to represent the major visual features of the auroral displays. It can be used by both designers and researchers to produce auroral images. Designers can use it to include auroral displays in the synthetic reproduction of polar scenes or in simulations of the night sky at different latitudes. Geophysical researchers may invert the model to refine auroral spectral and intensity data as well as to validate concepts regarding the composition of the outer atmosphere.

In the remainder of this paper we initially introduce the scientific background for the auroral phenomena. We then present our auroral modeling approach. Afterwards, we describe the rendering algorithm used to generate the auroral images. Finally, we compare the images produced using our model with photographs of auroral displays, and outline directions for future work in this area.

Auroral Science

Auroral physics involves several complex energy reaction and transfer mechanisms. In this section we present

an overview of auroral phenomena, focusing on physical aspects directly related to its morphology, spectrum and intensity. The reader interested in a more detailed description of auroral physics is referred to the various comprehensive texts on these phenomena.^{2,4,5,8}

Phenomena Overview

During solar flares and coronal mass ejections,⁹ plasma particles, namely protons and electrons, are emitted from the sun primarily from highly magnetized areas in the solar photosphere known as sunspots, which have a lower temperature than the solar corona. As the sun rotates, these particles are thrown out in spiraling streams, and after a few days the 'solar wind' hits the earth's magnetosphere. A shock front is then formed facing the sun (day side), while a plasma sheet is created in the opposite direction (night side) (Figure 2). In the core of this plasma sheet there is a magnetic field reversal where a small fraction of the solar wind particles are trapped. As a result, due to mechanisms yet unknown, plasma particles are accelerated towards the regions around the earth's magnetic poles. These particles create a global instability known as a magnetic substorm.

These magnetic substorms are characterized by large variations in the magnetic field and can occur over a short period of time. They may affect systems of energy transmission and the operation of satellites used in communication and global positioning systems.¹⁰ The bright and colorful displays of the aurorae are manifestations of such phenomena. In other words, the upper atmosphere is like a 'TV screen' that conveniently allows us to view and study phenomena in the distant




Figure 2. The solar wind forms a shock front when it hits the earth's magnetosphere. Some electrons interact with a neutral plasma sheet and travel back toward the earth, eventually descending in the regions around its magnetic poles.

tail of the plasma sheet (Figure 2). More 'explosive' substorms, and consequently more impressive auroral displays, are more likely to occur when the sun goes through periods of high sunspot activity at approximately 11-year intervals; the last peak of solar activity occurred in the 2000–2001 period.^{4,9}

Light Emission

The protons are not as efficient generators of auroral light as the electrons, so they will not be discussed further. We will focus our discussion on the electron aurorae. The directed velocity of a precipitating auroral electron can be separated into two components parallel and perpendicular to the magnetic field lines represented by \vec{B} . The parallel component is unaffected by the magnetic field, and the electron, therefore, moves along \vec{B} with a speed equal to its parallel velocity. The electron does not cross the magnetic field line. Instead, it travels around on a circle at a speed equal to the perpendicular component of its velocity. The combination of both components leads to a spiraling motion for the electron with an overall macroscopic progression parallel to the field lines (Figure 3).¹¹

As the electrons travel down along the earth's magnetic field lines they suffer many random deflections, which are caused by collisions with atoms of the atmosphere (Figure 3). These deflections spread the electrons out horizontally. When the electrons interact with atoms of the atmosphere, the atoms are excited and after a period of time they may emit a photon. Statistically, several collisions must occur before a photon is emitted. The wavelength of this photon, a *spectral emission line*,




Figure 3. A single electron undergoes several collisions on its travel toward earth. Each collision can potentially change the electron's energy and position.



Figure 4. Collisions (c) between auroral electrons and atoms or molecules in the upper atmosphere may generate photon emissions (e) at different wavelengths. The emission of a red photon is preceded by a long delay (around 110 s), while the photon emissions at the other wavelengths are characterized by short delays (less than 1 s).

depends primarily on the type of atmospheric constituent hit by the electron and the stability of the atom's excited state. In practice the wavelength will correlate with height (Figure 4). Almost all of the visible aurorae occur at altitudes between 100 km and 300 km above the earth's surface and thus appear to be part of the distant night sky. Atomic oxygen and molecular and atomic nitrogen are the principal constituents of the upper atmosphere available for and involved in the production of auroral emissions at these altitudes.

Besides the spectral emission lines generated by collisions between electrons and atoms, the auroral spectrum is also composed of *spectral emission bands* generated by collisions between electrons and molecules. The most common and the brightest visible feature of the aurorae, the atomic oxygen 'green line' at 557.7 nm, is dominant in the lower parts of auroral displays, around 100 km. It is mainly due to this emission line that most aurorae appear yellow-green (Figure 5). Because the peak of human light sensitivity is about 555.0 nm, these aurorae are particularly bright-looking. The red in the upper parts is caused mainly by another commonly observed line, the atomic oxygen 'red line' at 630.0 nm. The bluish color, seen sometimes on the lower border of auroral displays, comes mainly from the ionized nitrogen 'blue band' at 427.8 nm. These



Figure 5. Photograph of a real auroral display (courtesy of Jan Curtis). Two auroral curtains produced by two sheets of precipitating electrons causing emissions primarily at a wavelength of 557.7 nm.

emissions may be considered to form a triad with a highly saturated blue, yellowish-green, and red, which are almost the ideal primaries for an RGB display system. The spectral variety of auroral displays is further contributed to by weaker light emissions at other wavelengths across the visible spectrum. The mixtures in various ratios of all of these components may result in a wide variety of hues and colors.

Even though green and red, which are the strongest lines in the auroral spectrum, both originate from excited atomic oxygen, they behave quite differently. The transition state that produces the atomic oxygen green line only exists for up to 0.7 s, and the excited atom cannot move far before its photon is emitted. As a result the green line is often visible in structured forms. The transition state that produces the atomic oxygen red line can exist for 110 s, and the atom can travel a longer distance from the point at which it was excited. As a result, the red emissions are spread over a wider area. The ionized nitrogen blue band has a spatial distribution similar to the green line, since the transition state that gives rise to this band can exist for less than 0.001 s.⁴ The longer the life of an excited atom, the greater the chance it has of colliding with other atmospheric particles and losing its capacity to emit light, a process known as *quenching*.⁸ This explains why the red oxygen line is weaker than the green line at lower altitudes (higher atmosphere density) within an auroral display, and why some auroral features are red at high altitude and green lower down.

Auroral Shape

The basic shape of an aurora is determined by the energy and density of the electrons entering the atmosphere, as well as the local variations in the electric field of the earth. The most basic aurora can be thought of as a 'curtain' of light emissions from 'sheets' of falling electrons. This curtain will be colored, brightly yellow-green at the bottom, perhaps red at the top, and a yellow/orange transition may also be present. It can bend and fold almost like a real curtain, thereby generating bright vertical streaks (Figure 6). Such features occur at many scales in real aurora. In fact, the thin sheets of precipitating electrons are often subject to quasi-periodic, rotational distortions:^{12,13} spirals, curls and folds. The spirals are large-scale distortions (≈ 50 km apart), the curls are small scale (≈ 2 – 10 km apart) and the folds are intermediate scale (≈ 20 km apart). Folds and curls are more common in auroral displays observed from the ground¹³ (Figure 7).

In addition to different spatial scales, these distortions have also different time-scales. Folds can exist for more




Figure 6. Because the emission is in all directions within an auroral curtain and the curtain is no more opaque than the regular atmosphere, the apparent surface brightness is proportional to the thickness of the curtain as seen by the viewer. (a) A cross-section of a curtain and (b) a graph of apparent thickness.




Figure 7. Cross-sections of an auroral curtain as viewed from above, illustrating two common rotational distortions: (a) fold and (b) curl.

than a second, while curls have a lifetime in the preferred range of 0.25–0.75 s.¹⁴ Both types of rotational distortions are responsible for distinct visual features present in auroral displays. The light emanating from convoluted folds in auroral arc curtains oftentimes creates the impression of vertical, or near-vertical, 'stripes'.⁵ Despite their highly transient nature, curls are largely responsible for another important auroral feature, namely electron beams evolving into thin field-aligned filaments or 'rays'. An auroral ray typically has a diameter of 1 km or less and a vertical dimension up to several hundred kilometers.

Auroral Morphology

The auroral displays present a variety of forms when observed from the ground. These forms have given rise to a terminology discussed extensively in the International Auroral Atlas.¹⁵ For our purposes we divide the auroral forms into two groups: those with and those without a rayed structure. Among the auroral forms without a ray structure we may find:

- Arcs and bands: homogeneous structures that can extend over 1000 km, whose width (thickness) may vary from several hundred meters to more than 10 km, while their vertical dimension is around 20–30 km (Figure 8a and c). Sometimes a band may be twisted into horseshoe bends.
- Diffuse patches: these forms have a cloud-like appearance and may cover several hundred square kilometers. Because they are weak and lack definite structures they can be difficult to see with the naked eye.

Among the auroral forms with a ray structure we may find:

- Rayed arc: a homogeneous arc broken up into vertical striations (Figure 8b).
- Rayed band: a band made up of numerous vertical striations (Figure 8d).
- Rays: ray-like structures, appearing singly or in bundles separated from other forms.
- Corona: a rayed aurora seen near the magnetic zenith, which due to a perspective effect gives the appearance of a crown with rays converging on one point (Figure 8e).
- Drapery: a band made up of long rays, around 70–100 km, giving the appearance of a curtain, which may be folded (Figure 8f).

Rayed bands and draperies are visually very similar. The difference between these auroral forms is mostly




Figure 8. Artist's conception of some typical auroral forms: (a) homogeneous arc, (b) rayed arc, (c) homogeneous band, (d) rayed band, (e) corona and (f) drapery (courtesy of Adriane Elena Baranoski).

associated with the length of the rays, which are shorter in rayed bands, and the overall distribution of bends, which tend to appear more often in draperies.

Groups of arcs or bands frequently occur simultaneously, exhibiting similar geometrical behavior and forming an arc or band system (Figure 5). The spacing between arcs and bands in a system is typically 10–100 km.¹³ Two or more systems may be visible during some phases of a large auroral display.

Temporal Variations

The aurorae exhibit time variations with respect to spatial movement and intensity. Some of these variations may be of the order of many minutes, while others can occur in less than a second. For example, arcs can drift with typical velocities of 100–200 m/s¹⁶ in the north–south direction (although speeds up to 1 km/s can also be observed,¹⁷) and 10 km/s in the westward (high-latitude) and eastward (low-latitude) directions.¹⁸ Smaller-scale arc distortions such as folds and curls have much higher speeds. Folds have apparent horizontal velocities in the range of 0–5 km/s.¹³ The apparent horizontal velocities of curls lie in the range of 0–90 km/s, with preferred speeds in the range of 0–8 km/s.¹⁴ To have a better idea of the magnitude of these speeds, the reader has to consider that an auroral arc may extend over the entire field of view of an observer on the ground. For this observer the auroral rays will 'travel' back and forth between two extreme points in the horizon in few seconds, giving the visual impression of 'dancing lights'. The intensity variations of the

aurorae are usually a function of where in the substorm cycle they occur. This cycle consists of a quiet phase, a growth phase (up to 2 h), an active expansion phase (≈ 20 min) and a recovery phase (30 min to 2 h). The highest intensity variations occur during the expansion phase.¹⁹

Aurora Modeling

Despite their inherent complexity most auroral displays present a set of features that make them readily recognizable by viewers: arcs and band shapes with and without rayed structures, characteristic spectral variation, vertical distribution of intensity dependent on the type of auroral display and apparent surface brightness dependent on the direction of observation. Our modeling approach focuses on the realistic simulation of these characteristics.

In this section we describe how we simulate the underlying physical processes affecting the auroral visual features using the available data. The world coordinate system considered in our simulations as well as the elevation angle θ used in the auroral observations are sketched in Figure 9. Table 1 summarizes the parameters used in our auroral modeling approach.

Auroral Shapes and Internal Structures

Our model of auroral shapes is based on the particle nature of the aurorae. However, instead of following the trajectories of individual electrons, we simulate the paths taken by beams of electrons. These beams represent the auroral rays or curls. This modeling approach is divided into three parts: the determination of the starting points for the beams, the simulation of their paths along the geomagnetic field lines and the computation of the light emissions along their paths.




Figure 9. Sketch showing the world coordinate system considered in our simulations and the elevation angle θ .

Symbol	Definition
w	Electron sheet thickness
Φ	Phase shift of electron sheet boundaries
λ	Electron sheet wavelength
n	Number of electron beams
b	Number of sheet internal layers
P_i, P_f	World coordinates of the sheet extremities
s	Sheet longitudinal parametric coordinate
Φ_r	Sheet angular parametric coordinate
d_s	Longitudinal distance between beams
ξ_1	Random increment to d_s
d_r	Angular increment on Φ_r
\vec{B}	Magnetic field
\vec{B}_{100}	Magnetic field at 100 km
\vec{u}_a	$P_f - P_i$
D	$ \vec{u}_a $
\vec{u}_b	$ \vec{B} \times \vec{u}_a$
d_f	Quasi-periodic interval between folds
α_D	Loss cone angle
h	Altitude of the loss cone
t	Vertical parametric coordinate
d_t	Vertical distance between collision points
ξ_2	Random perturbation on d_t
\vec{v}	Velocity vector
\vec{v}_p	Precipitation vector
\vec{v}_t	Temporary precipitation vector
α	Polar (pitch) angle between \vec{v} and \vec{v}_t
$\Delta\alpha$	Adjustment of α with respect to the altitude
β	Azimuthal angle between \vec{v} and \vec{v}_t
ξ_3	Random perturbation on β
L	Electron beam path length
ΔL	Variation on the path length
ξ_4	Random perturbation on ΔL

Table I. Table of symbols used in the aurora modeling




Figure 10. Top view of cross-sections of two electron sheets with boundaries represented by sine waves with a given phase shift Φ : (a) without folds and (b) with folds.

Sheet Model. Studies of auroral physics^{12,20} show that the stream of precipitating electrons that cause auroral displays can accurately be represented by sheets with boundaries given by sine waves. These waves have equal amplitudes, but they are not in phase. The phase shift, Φ , may be as high as 0.35π radians.¹² Each sheet can be formed by a few internal layers, and its shape defined by two parameters: the thickness, w , and the wavelength, λ (Figure 10a). The thickness of an electron sheet may vary from about 1 to 10 km, and the ratio represented by $\lambda/(2w)$ has a preferred range of 5.2–31.4.¹² By varying these parameters one can simulate auroral disturbances that appear as a series of asymmetric bulges in the sheet.

The boundaries as well as the internal layers of these sheets are modeled using sine curves. Hence, the starting points for the electron beams used in our simulations are obtained through the discretization of these curves. The world coordinates of the extremities of the electron sheets, P_i and P_f , the number of points in each sine wave, n , the number of internal layers, b , and the parameters Φ , w and λ defined above are the input parameters used to define these sheets.

Intuitively, a starting point P can be seen as a point of a sine curve described parametrically by $s \in [0..1]$ and $\Phi_r \in [\Phi_i.. \Phi_f]$ (Figure 11). Due to the phase shift, each sine curve representing a sheet layer starts with a different initial angle Φ_r^i , and the difference between the initial angles of the boundary layers corresponds to the overall phase shift Φ .




Figure 11. Geometry for generating the parametric sine waves.




Figure 12. Stages of the simulation of an electron beam path: (a) precipitation, (b) computation of deflection points and (c) displacement of the electron beam along the magnetic field lines.

Computationally, a starting point P is obtained using the equation below, which takes into account the geometry shown in Figure 11:

$$P(s, \Phi_r) = \vec{u}_a(s) + w \sin(\Phi_r) \frac{\vec{u}_b}{|\vec{u}_b|} \quad (1)$$

where \vec{u}_a corresponds to the vector from P_i to P_f and \vec{u}_b corresponds to the vector resulting from the cross-product of \vec{u}_a and the geomagnetic field vector \vec{B} .

The angular displacement given by Φ_r is computed incrementally such that $\Phi_r^{new} = \Phi_r^{old} + d_r$, where the increment d_r is given by

$$d_r = \frac{D 2\pi}{n \lambda} \quad (2)$$

where D represents the distance between P_i to P_f , which corresponds to $|\vec{u}_a|$, and the constant 2π indicates the sine functions period.

The parametric interval represented by s is obtained using a parametric longitudinal threshold d_s , i.e., $s^{new} = s^{old} + (d_s \pm \xi_1)$, where ξ is a uniformly distributed random number in the interval $[0..1]$. If $\xi_1 > 0.5$, then we add to d_s , otherwise we subtract from d_s . The use of a random displacement as opposed to a regular one is consistent with the stochastic nature of the phenomena.

In our simulation strategy, folds are modeled by replacing portions of the sine waves by Bézier curves. The number of starting points between folds is also an input parameter, n_p , which is randomly displaced, such that we can have quasi-periodic intervals, d_f , between these folds (Figure 10b). The starting points placed on the folds are also randomly displaced, and by varying the control points of the Bézier curves we can simulate folds in different rotational stages. The quasi-random parametrization of the sine waves and Bézier folds results in approximations for electron sheets

(Figure 10), which are consistent with the current auroral theories.^{12,20}

Precipitation of Electron Beams

The electrons are randomly deflected after colliding with the atoms of the atmosphere. These deflections play an important role in the dynamic and stochastic nature of the auroral displays, and hence they are taken into account in our simulations. We consider the deflection points as emission points, and they are used to determine the spectral and intensity variations of the modeled auroral displays.

The tracking of each electron beam starts with the computation of the starting points described in the previous section. The electron beam's velocity vector, \vec{v} , is defined as the overall direction of progression of the particle during its spiraling descending motion (Figure 3). The angle between the electron's velocity vector and the geomagnetic field vector \vec{B} is called the pitch angle, α . A 'loss cone' of pitch angles is bounded at an altitude h by an angle α_D (Figure 12a). This angle is given by an adiabatic invariant,[†] which takes into account the ratio between the strength of \vec{B} at h and at an altitude of 100 km (the auroral lower border).²²

Electrons with $\alpha \leq \alpha_D$ are in the loss cone and are precipitated ('lost') into the atmosphere. The boundaries of this loss cone are somewhat diffuse ($\alpha_D \approx 2 - 3^\circ$), since the value of α_D decreases with altitude.²² The length of the path is given by a parameter L which is associated with the height chosen for the modeled auroral display. As mention earlier, it assumes typical values around 20–30 km for arcs and bands, and around 70–100 km for draperies.

[†] $\alpha_D = \arcsin \sqrt{\frac{B}{B_{100}}}$ ^{21,22}

Each path is simulated incrementally, through the vertical displacement t such that $t^{new} = t^{old} + (d_t \xi_2)$, where ξ_2 is a uniformly distributed random number in the interval $[0..1]$. The use of this random displacement is consistent with the spatial inhomogeneity of auroral electrons.²³ The vertical threshold d_t is an input parameter which depends on the initial energy of the electrons. For instance, on average an electron with 10 keV (60,000 km/s) can collide 300 times before being brought to a halt[‡] at an altitude of about 100 km above the ground.⁴ In this case, since we assume that $t \in [0..1]$, we could use $d_t = \frac{1}{300}$.

The deflections are simulated using the following sequence of steps, which is based on the diagram sketched in Figure 12(b). An intermediate point $P^{new'}$ is computed using

$$P^{new'} = P^{old} + (L d_t \xi_2) \frac{\vec{B}}{|\vec{B}|} \quad (3)$$

A temporary precipitation vector $\vec{v}_t = P^{new'} - P^{old}$ is then computed. The perturbation of this vector using a polar angle α and an azimuthal angle β provides the direction for the precipitating vector \vec{v}_p . The new point is computed using \vec{v}_p as a directional displacement such that $P^{new} = P^{old} + \vec{v}_p$.

The polar angle α corresponds to a pitch angle adjusted to follow the reduction of α_D with altitude. This adjustment is performed using

$$\alpha = \alpha_D - (t \Delta\alpha) \quad (4)$$

where $\Delta\alpha$ is an input angular parameter usually of the order of 0.015 radians.²²

The azimuthal angle β is obtained randomly, i.e., $\beta = 2\pi \xi_3$, where ξ_3 is a uniformly distributed random number in the interval $[0..1]$. This choice for β also follows the stochastic characteristics of the space distribution of auroral electrons.²³ The net result of this displacement is the electron beam being spread out onto other field lines (Figure 12c).

Because the electron precipitation is governed by statistical processes, the actual penetration depths are not identical even for two electrons with identical initial conditions.^{11,24} In our simulations we account for this aspect through two kinds of perturbations on the electron beams. The first consists of changing the interval for the parametric variable t , which becomes

$[0..(1.0 - \xi_4 \Delta L)]$, where ξ_4 corresponds to a uniformly distributed random number in the interval $[0..1]$, and ΔL corresponds to an input parameter representing a variation on the path length. The second consists of perturbing the z coordinate of the starting point of an electron beam (Figure 12a). This perturbation is performed by applying a 3D correlated noise function,²⁵ and it also causes a variation in the beam path length. However, in the second case the variation will be related to perturbations performed in the path length of the neighbor beams, and it will be more noticeable in a plane perpendicular to the auroral display.

Auroral Spectrum

Since the aurora is characterized by its storm-like behavior, the variations of spectral ratios (Figure 13a) and intensities (Figure 13b) according to the auroral heights are given in the literature as average values. According to Romick and Belon,²⁶ these vertical spectral and intensity profiles are a good approximation of the vertical emission profile. For these reasons our simulation of the auroral spectrum consists of sampling the spectral curves presented in Figure 13(a) using the heights of the deflection points as the sampling parameter.

As mentioned earlier, the spectral variety of auroral displays is further contributed to by several weaker




Figure 13. (a) Spectral emission curves (redrawn from Brekke and Egeland.⁴) (b) Auroral intensity profiles along various forms (redrawn from Vegard and Krogness).²⁷

[‡]These electrons are not destroyed, but in the course of their passage through the atmosphere they eventually become indistinguishable from the ambient electron population.²³

light emissions at other wavelengths across the visible spectrum. For the purpose of simulating the aurora, however, a viable approach is to focus on the bright emissions (at 630.0 nm, 557.7 nm, 427.8 nm), which are most obvious to the casual observer. The spectral variety seen in the aurora can then be simulated by mixing the components of this triad in various ratios. As the data for other wavelengths become available they can be easily incorporated in our simulations and accounted for during the rendering process.

Auroral Intensities

The light intensity from an aurora is proportional to the deposition of energy into the atmosphere by the precipitating electrons. As a result, the height and the intensity distribution of each auroral form are related to the average energy distribution of the particles. Vegard and Krogness²⁷ performed a series of ground-based measurements that showed not only characteristic differences of the distribution of intensity of various auroral forms, but also very different intensity distributions even for the same auroral form (Figure 13b). In some of the auroral forms the intensity is concentrated into a band of only 10–20 km vertical distance, and the lower border can be quite sharp. The intensity of a discrete arc typically falls to 10% within a few kilometers below the maximum, and to 1% a kilometer or two below that.²⁸ The findings of Vegard and Krogness have been recently corroborated by rocket-based measurements.^{4,24,26} These vertical variations of intensity of auroral forms are also accounted for in our simulations by sampling the intensity curves shown in Figure 13(b), also using the height of the deflection points as the sampling parameter. The incorporation of these intensity values in the computation of the volume emissions is described in the next section.

As mentioned previously, the observed intensity of a particular auroral display depends on the direction of observation (Figure 6). In this context the apparent surface brightness is used to define the intensity of an aurora, and it is given in rayleighs (R).²⁹ The faintest auroral lights that can be detected with the naked eye are of the order of 1 kR in the green line. One rayleigh is equal to an integrated emission rate of 10^6 photons per square centimeter per column per second.* The auroral

*In this context 'column' refers to the unknown height of the column above the apparent source, and it is included to show that this is a volume emission.²⁴

IBC	Intensity comparable to	κR
1	Milky Way	1
2	Thin moonlit cirrus clouds	2
3	1/10 of the intensity of full moonlight	3
4	1/3 of the intensity of full moonlight [†]	4

Table 2. Intensity classification of aurora intensities (redrawn from Eather² and Brekke and Egeland⁴)

intensities are commonly classified in terms of the international brightness coefficient (IBC), as shown in Table 2.

In short, the apparent surface brightness of an aurora is proportional to the integrated emission per unit volume along the line of sight.^{14,24} A thin auroral layer covering a large part of the sky is, therefore, most intense when viewed at low elevation angles (Figure 9) due to the high density of emissions observed in the lower border of several auroral forms (Figure 13b). This dependence on the direction of observation of auroral displays is taken into account through the use of a view-dependent rendering approach, which is described in the next section.

Rendering

The rendering of auroral displays is performed in three stages. In the first stage the light emissions are mapped[‡] to a screen plane and stored in an image raster array. Afterwards these values are converted to RGB values in the second stage. In the last stage the image raster array is convolved to simulate auroral temporal and spatial variations.

Forward Mapping of Emissions

Each electron beam is tracked from its starting point until it reaches the low border of the electron sheet. The deflection points along the path of an electron beam can be regarded as emission points. The world coordinates of each emission point E are used to compute an emitted ray.

[†]The rate of photons hitting a detector faceplate aimed at this type of aurora and subtending a solid angle ω is three times smaller than the rate of photons hitting the same detector faceplate aimed at the full moon and subtending the same solid angle ω .

[‡]This mapping is similar to splatting techniques used in volume-rendering applications.³⁰




Figure 14. Geometry for the forward mapping of auroral emissions to the screen, or image, plane of a virtual camera.

The direction of this ray corresponds to the vector \vec{e} given by $Q - E$, where Q corresponds to the projection center of a pinhole camera (Figure 14). This emitted ray represents a stream of photons emitted at E in the direction given by \vec{e} .

In the next step we determine if this emitted ray intersects the screen plane. If there is no intersection, we move to the next deflection point in the path. Otherwise, we determine the pixel on the screen that corresponds to the intersection (Figure 14). The screen coordinates of this pixel correspond to the indices used to access the corresponding element of the image raster array. This array is used to accumulate the weight of the emitted rays that hit a certain pixel. This weight has three components that correspond to the RGB channels. These three components are obtained from a look-up table, which is derived from spectral emission curves presented in Figure 13(a) and accessed using the height of the emission point given by its component E_y . These weight components are also multiplied by an intensity factor, which accounts for the different vertical variations of intensity of the auroral forms. Its value is obtained from the look-up table which is derived from intensity curves presented in Figure 13(b) and also accessed using the height of the deflection point given by E_y .

After performing the forward mapping of emissions for all deflection points along a given path, the algorithm proceeds to tracking another electron beam. If there are no more beams to be tracked then the weight components stored in the image raster array are converted to RGB values as described next.

Conversion to RGB Values

The conversion of the auroral spectral emissions at wavelengths of 630.0 nm, 557.7 nm and 427.8 nm to

RGB values associated with SMPTE* chromaticity coordinates follows two steps. In the first step CIE† spectral tristimulus values for the wavelengths of the auroral emissions, $\tilde{x}(\lambda)$, $\tilde{y}(\lambda)$ and $\tilde{z}(\lambda)$,³¹ are converted to RGB tristimulus values, $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$, using

$$\begin{bmatrix} r(\lambda) \\ g(\lambda) \\ b(\lambda) \end{bmatrix} = M \begin{bmatrix} \tilde{x}(\lambda) \\ \tilde{y}(\lambda) \\ \tilde{z}(\lambda) \end{bmatrix} \quad (5)$$

where M corresponds to a transformation matrix set according to SMPTE chromaticity values.³²

In the second step the RGB color is quantified by sampling the spectral curve $f(\lambda)$ of the emitting source:

$$R = \int f(\lambda)r(\lambda)d\lambda \quad (6)$$

$$G = \int f(\lambda)g(\lambda)d\lambda \quad (7)$$

$$B = \int f(\lambda)b(\lambda)d\lambda \quad (8)$$

In practice, these integrations are performed by summations in which the spectral energy distribution represented by $f(\lambda)$ is discretized.³³ Since in our case we are considering auroral emissions in only three wavelengths, these summations become

$$R = \sum_{i=1}^3 f'_i(\lambda)r_i(\lambda) \quad (9)$$

$$G = \sum_{i=1}^3 f'_i(\lambda)g_i(\lambda) \quad (10)$$

$$B = \sum_{i=1}^3 f'_i(\lambda)b_i(\lambda) \quad (11)$$

where $f'_i(\lambda)$ corresponds to the sum of weights stored in a pixel for each wavelength considered.

This discretization results in a loss of accuracy for light sources with a continuous spectral energy

*Society of Motion Picture and Television Engineers.

†Commission Internationale de L'Eclairage.

distribution. However, since the three auroral spectral emissions considered in our simulations are dominant in the auroral spectrum, this loss of accuracy is likely to be small. Moreover, as data for auroral emissions at other wavelengths become available, they can be easily incorporated in this conversion scheme.

Antialiasing

After being hit by an auroral electron, an atmospheric excited atom may move away from the collision point before emitting a photon. Hence, the auroral emissions from a given emission point may contribute to multiple pixels. Statistically, the intensity contribution spreads radially around a central direction of motion and follows a Gaussian distribution of intensity along that dimension.³⁴ The different lifetimes of transition states in auroral emissions cause distinct spread areas around the principal direction of propagation of an electron—110 s for the red line, 0.7 s for the green line, and 0.001 s for the blue line.

For simulating this distribution of the auroral emissions, we use a Gaussian kernel to blur the image obtained with strict forward mapping. That is, we spread the contribution at each pixel to a surrounding pixel neighborhood (Figure 15), according to a normalized Gaussian distribution.³⁵ The Gaussian kernels used in our experiments are given by

$$G(i, j) = \frac{1}{2\pi\sigma^2} e^{-\frac{(i^2 + j^2)}{2\sigma^2}} \quad (12)$$

where i and j are the indices of a Gaussian kernel whose size in each dimension is in the range $[1..m]$, where m corresponds to the mean and σ to the standard variation




Figure 15. (a) Convoluting the image raster array using a Gaussian kernel. (b) Example of a Gaussian kernel with $m = 31$ and $\sigma = 5.25$.

of the Gaussian kernel. The larger the value of σ for a given kernel size, the greater the blurring of the intensity variations.

We apply different kernel supports to each color channel, for approximating the distinct transition states at different spectrum lines. The pixel values of the red channel receive the widest spreading factor in image space, whereas the blue channel receives the narrowest factor. For an image with 218×383 pixels and a field of view of 35 degrees vertical, we determined the Gaussian standard deviations and the kernel supports per color channel empirically. For the red channel we used a standard deviation of 4.75 and a kernel support of 21×21 pixels, whereas for the green channel we used a standard deviation of 2.0 and a support of 5×5 pixels. We also determined no blurring for the blue channel empirically. Scaling of the image resolution or of the field of view implies proportional scaling of the kernel parameters. Mathematically, this blurring process is the convolution of the image with a color-dependent Gaussian low-pass filter.³⁰

In addition to the local motion that each individual excited atom undergoes, an aurora also exhibits global temporal variations that affect the entire phenomenon visually. These variations are usually captured in photographs as blurred auroral displays, due to finite exposure times. For simulating this global blurring effect, we sample the temporal variations of an aurora along time, and compute the accumulation of all these independent contributions on a single final image. This corresponds to generating an image of an aurora for each sample in time, and then computing the arithmetic average of all the generated images. Mathematically, this corresponds to another convolution, but using a temporal low-pass box filter.³⁵ The longer the sampled window is in time, the blurrier the results, similarly to the effects captured in real photographs with longer exposure times.

Results

For illustrating the applicability of our model, we simulated several auroral displays. To facilitate the comparison with real photographs of aurorae, we also added a background—night sky—and a foreground—mainly vegetation—with respect to the simulated aurora display (Figure 16). The image synthesis process starts by initializing the color buffer with a background scene, then superimposing the auroral display, and finally filtering out regions of the image that would be covered by foreground objects.



Figure 16. Simulation of two auroral curtains in different settings.

The dynamic nature of the aurorae precludes a quantitative analysis of their visual simulations. The alternative available is to qualitatively analyze the simulations in comparison with photographs of real aurorae. However, whereas photographs of an aurora by necessity contain a certain amount of blur (Figure 17) due to low light conditions and significant exposure times,³⁶ instantaneous simulated images contain more defined structures (Figure 18), which is closer to a naked eye observation of the phenomenon. Figure 19 also illustrates this aspect.




Figure 17. Early twentieth-century black and white photographs of auroral bands.



Figure 18. Simulations of auroral bands.




Figure 19. Left: Photograph of two ray-filled curtains (courtesy of Jan Curtis, 5 s of exposure time). Right: Simulation of auroral curtain showing the ray structure.

Figure 20 illustrates the application of local antialiasing techniques, which account for small temporal variations of auroral emissions. The net effect of these variations, i.e., the subtle color change due to the local 'blur', can be captured in photographs of auroral displays depending on the exposure time used. Therefore, besides accounting for the physical phenomena, the application of these techniques contributes to making the appearance of the synthetic images closer to photographs of real auroral displays (Figure 21). Although in our experiments we have obtained reasonable results using $m = 5$ and $\sigma = 2.0$ for the green channel and $m = 21$ and $\sigma = 4.75$ for the red channel, these parameters can be fine tuned in order to obtain a closer match. However, as mentioned earlier, our main goal is the realistic simulation of the auroral phenomena,



Figure 20. Auroral simulations. Left: Without local blur. Right: With local blur.



Figure 21. Left: Photograph of a broad drape (courtesy of Jan Curtis, 8 s of exposure time). Right: Simulation of a broad drape with local blurring (using $m=3$ and $\sigma=1.5$ for the green channel and $m=31$ and $\sigma=5.25$ for the red channel).

instead of the simple reproduction of their photographic appearance. As mentioned earlier, these photographs are used just as references in our qualitative comparisons.

However, the local antialiasing is not enough to entirely mimic the blur due to finite exposure times of the auroral photographs. As shown in Figure 22, this effect is better approximated with a more global antialiasing technique based on the temporal sampling of auroral displays. In order to further illustrate this sampling approach, we produced a few animation sequences. We simulated movements often noticed in auroral displays, namely the drift of arcs and the horizontal motion of curls. The simulation of the drift of arcs was performed by changing the coordinates of the ending points of an electron sheet, according to a specified time interval determined by a given drift velocity. For simulating the horizontal swift motion of curls along arcs, we varied the number of electron beams along each electron sheet and shifted the position of the folds. Figure 23 shows frames of an animation sequence (60 frames per second) in which arc drifting (2.5 km/s) and a horizontal motion (left to right) of curls along the arc (8.5 km/s) are illustrated. It should be noted that these simulations do not account for complex electromagnetic plasma instabilities, which is an area of research in itself.³⁷

Besides visually comparing our synthetic images with static photographs of real aurorae, we requested feedback from several experts in polar science at various universities. One of the most detailed feedbacks came from Jan Curtis, a climatologist who has often experienced and photographed auroral displays. When shown



Figure 22. Simulation of finite exposure time using the average of (top left) 2, (top right) 4, bottom (left) 8, and (bottom right) 16 images uniformly sampled in time (exposure time ranging from 1/30 to 4/15 s). Longer exposure times imply blurrier images or fewer fine details.

our synthetic images, he indicated that the shape is accurately represented and that there is a reduced gradation from the brightest light to the faintest pale in our images, which can be observed in real aurorae (J. Curtis, personal communication, 2000). This observation indicates that undersampling followed by blurring in our simulations has not made our images unrealistic. Indirectly, this observation also addressed the lower dynamic range in synthetic image representation and display (24 bits split into three color channels) than in photographic film and the human eye.

Thomas Hallinan, one of the leading scientists in the auroral physics field, also provided us valuable feedback. He stated that 'the auroral drape simulations are quite good' (T. J. Hallinan, personal communication, 2001) and remarked on 'a curious (and accurate) feature of the simulations' (T. J. Hallinan, personal communication, 2001), namely the sharp upper border of some of our computer-generated rayed arcs. In his opinion this accords well with photographs of the aurora, but not



Figure 23. Frames from an animation sequence in which an arc drift and a horizontal motion (left to right) of curls are illustrated: (top left) frame #1, (top right) frame #143 (bottom left) frame #320 and (bottom right) frame #1200. The yellow arrows indicate the relative movement of auroral rays.

quite with some models of the vertical luminosity profile. He attributes this feature to the fact that the spectral data used in our simulations came from densitometry of auroral photographs.

The simulations presented in this paper were implemented on a Compaq Alpha cluster of 30 machines, consisting of 500 MHz Alpha EV56 processors.³⁸ The approximate time taken to generate each image using five processors in parallel was 22 s.

Conclusion and Future Work

We presented a physically based approach for simulating auroral phenomena. The major auroral visual features apparent to naked eye can be promptly identified on our images. The results presented in the previous section demonstrate that the proposed approach qualitatively represent real auroral displays shown in photographic images commonly published in the literature.

Thus, it can be used for the generation of realistic images of these natural phenomena aiming at educational and scientific applications. Its modular design also allows its use on the investigation of auroral theories and evaluation of auroral data.

Our current efforts are focused on the auroral dynamics,³⁷ more specifically on the simulation of plasma instabilities responsible for the spectacular auroral motions and shape changes. The most fascinating consequences of these processes are the formation of rotational distortions having different spatiotemporal scale, notably the majestic auroral spirals or surges. We also believe that the rendering of auroral displays can be done in real time. As future work we intend to examine techniques for real time rendering of aurorae by exploiting graphics hardware.

Finally, the complexity and large variety of shapes, colors and movements of the aurora and its direct connection with plasma physics present several avenues for future research, with artistic, educational, and scientific applications. However, as suggested by Eather,² perhaps the main motivation for continued study of the aurora is the aurora itself, with its beauty and mystery.

ACKNOWLEDGEMENTS

The authors would like to thank Jan Curtis, Thomas J. Hallinan and Ove Harang for their valuable feedback, and the anonymous reviewers for their useful insights. This research was supported by the Natural Sciences and Engineering Council of Canada (grants 691266 and 238337) and the Canada Foundation for Innovation (Project 6218).

References

1. Musset P, Lloret A. *Concise Encyclopedia of the Atom*. Collins: Glasgow, 1968.
2. Eather RH. *Majestic Lights*. American Geophysical Union: Washington, DC, 1980.
3. Lynch DK, Livingston W. *Color and Light in Nature*. Cambridge University Press: Cambridge, UK, 1995.
4. Brekke A, Egeland A. *The Northern Lights: Their Heritage and Science*. Grøndahl og Dreyers Forlag: Oslo, 1994.
5. Bryant DA. *Electron Acceleration in the Aurora and Beyond*. Institute of Physics: Bristol, 1999.
6. Savage C. *Aurora: The Mysterious Northern Lights*. Sierra Club Books: San Francisco, 1994.
7. Baranoski GVG, Rokne JG, Shirley P, Trondsen T, Bastos R. Simulating the aurora borealis. In *8th Pacific Conference on Computer Graphics and Applications*, IEEE Computer Society: Los Alamitos, CA, October 2000; 2–14.

8. Jones AV. *Aurora*. D. Reidel: Dordrecht, 1974.
9. Burtnyk K. Anatomy of an aurora. *Sky and Telescope* 2000; **99**(3): 35–40.
10. Odenwald S. Solar storms: the silent menace. *Sky and Telescope* 2000; **99**(3): 41–56.
11. Borovsky JE, Suszcynsky DM, Buchwald MI, DeHaven HV. Measuring the thickness of auroral curtains. *Arctic* 1991; **44**(3): 231–238.
12. Hallinan TJ. Auroral spirals. 2. Theory. *Journal of Geophysical Research* 1976; **81**(22): 3959–3965.
13. Hallinan TJ, Davis TN. Small-scale auroral distortions. *Planetary Space Science* 1970; **18**: 1735–1744.
14. Trondsen TS. High spatial and temporal resolution auroral imaging. PhD thesis, Department of Physics, Physics, Faculty of Science, University of Tromsø, Norway, November 1998.
15. Edinburgh University Press. *International Auroral Atlas*, 1963.
16. Evans S. Horizontal movements of visual auroral features. *Journal of Atmospheric and Terrestrial Physics* 1959; **16**: 191–193.
17. Omholt A. *The Optical Aurora*. Springer: New York, 1971.
18. Oguti T. Rotational deformations and related drift motions of auroral arcs. *Journal of Geophysical Research* 1974; **79**(25): 3861–3865.
19. Brekke A. *Physics of the Upper Polar Atmosphere*. Wiley in association with Praxis: Chichester, 1997.
20. Chmyrev VM, Marchenko VA, Pokhotelov OA, Shukla PK, Stenflo L, Streltsov AV. The development of discrete active auroral forms. *IEEE Transactions on Plasma Science* 1992; **20**(6): 764–769.
21. Chamberlain JW. *Physics of the Aurora and Airglow*. Academic Press: New York, 1961.
22. Haymes RC. *Introduction to Space Science*. Wiley: New York, 1971.
23. Rees MH. *Physics and Chemistry of the Upper Atmosphere*. Cambridge University Press: Cambridge, UK, 1989.
24. Kivelson MG, Russell CT. *Introduction to Space Physics*. Cambridge University Press: Cambridge, UK, 1995.
25. Watt A, Watt M. *Advanced Animation and Rendering Techniques*. Addison-Wesley: New York, 1992.
26. Romick GJ, Belon AE. The spatial variation of auroral luminosity: I. Determination of volume emission rate profiles. *Planetary Space Science* 1967; **15**: 1695–1716.
27. Vegard L, Krogness Q. The variation of light intensity along auroral ray-streamers. *Geofysiske Publikationer* 1920; **1**(1): 149–170.
28. Hargreaves JK. *The Solar–Terrestrial Environment*. Cambridge University Press: Cambridge, UK, 1992.
29. Hunten DM, Roach FE, Chamberlain JW. A photometric unit for airglow and aurora. *Journal of Atmospheric and Terrestrial Physics* 1956; **8**: 345–346.
30. Westover LA. Splatting: a parallel, feed-forward volume rendering algorithm. PhD thesis, Department of Computer Science, University of North Carolina at Chapel Hill, NC, November 1991.
31. CIE. *Colorimetry Official Recommendations of the International Commission on Illumination*. Commission Internationale de L'Eclairage (CIE): Paris, 1970.
32. Lilley C, Lin F, Hewitt WT, Howard TLJ. *Colour in Computer Graphics*. ITTI Computer graphics and Visualisation, Manchester Computing Centre, University of Manchester, UK, December 1993.
33. Hall R. Comparing spectral color computation methods. *IEEE Computer Graphics and Applications* 1999; **July/August**: 36–45.
34. Borovsky JE, Suszcynsky DM. Optical measurements of the fine structure of auroral arcs. In *Auroral Plasma Dynamics*, Lysak R (ed.). American Geophysical Union: Washington, DC, 1993.
35. Castleman K. *Digital Image Processing*. Prentice-Hall: New York, 1996.
36. Eather RH. An aurora watcher's guide. *Sky and Telescope* 2000; **99**(3): 42–48.
37. Baranoski GVG, Wan J, Rokne JG, Bell I. Simulating the dancing lights. Technical Report CS-2002-16, School of Computer Science, University of Waterloo, loo, Ontario, April 2002.
38. Baranoski GVG, Rokne JG. Using a HPC system for the simulation of the trajectories of solar wind particles in the ionosphere. In *High Performance Computing Systems and Applications*, Dimopoulos N, Li KF (eds). Kluwer: Norwell, MA, 2002; 317–329.

Authors' biographies:

Gladimir V. G. Baranoski is an assistant professor in the School of Computer Science at the University of Waterloo, Canada. He has a BSc in Mechanical Engineering and an MSc in Computer Science from the Federal University of Rio Grande do Sul (Brazil), an MSc in Computer Science from Indiana University, and holds a PhD in Computer Science from the University of Calgary (1998). The main topic of his doctorate research was the rendering of natural scenes, namely the development of reflectance and scattering models for plants. As a postdoctoral fellow at the University of Utah he worked on the modeling and rendering of auroral phenomena. His research interests include biologically and physically based rendering, and simulation of natural processes involving radiative transfer of energy. His personal interests involve football (the Brazilian version) and athletics.

Jon G. Rokne was born in Bergen, Norway, in 1941. He is currently a professor and former Head of the Computer Science Department at the University of Calgary, Canada, where he has been a faculty member since 1970. He has been a visiting professor at the University of Canterbury, New Zealand (1978), University of Freiburg, Germany (1980), University of Grenoble, France (1983–84) and University of Karlsruhe, Germany (1984). He holds a PhD in Mathematics from the University of Calgary (1969). His research interests span interval analysis, computer graphics and global optimization.

Peter Shirley is an associate professor in the School of Computing at the University of Utah. He has a BA in Physics from Reed College and a PhD in Computer Science from the University of Illinois at Urbana–Champaign. He spent four years as an assistant professor at Indiana University and two years as a visiting assistant professor at the Cornell Program of Computer Graphics

before moving to Utah. His professional interests include computer graphics, statistical computing, visualization, and immersive environments. In his spare time enjoys golf, hiking, backgammon, and beer.

Trond S. Trondsen received his doctorate degree in Cosmic Geophysics from the University of Tromso, Norway, in 1998. His dissertation was on the development of a high temporal and spatial resolution portable camera suitable for imaging the Northern Lights at video rates. He currently works at the University of Calgary's Institute for Space Research, where he is

involved in the development of both ground-and space-based auroral imaging instrumentation, and the analysis and publication of the resulting image data.

Rui Bastos is a 3D graphics architect at nVIDIA Corporation. He holds a PhD and an MSc in Computer Science from the University of North Carolina at Chapel Hill, and an MSc in Computer Science and a BSc in Physics from the Federal University of Rio Grande do Sul in Brazil. His research interests include graphics architectures and hardware, and real-time simulation and photo-realistic rendering of natural phenomena.