



Graphics Gems Revisited

Stéphane Guy, Cyril Soler

► To cite this version:

Stéphane Guy, Cyril Soler. Graphics Gems Revisited. ACM Transactions on Graphics, 2004. inria-00510165


HAL Id: inria-00510165

<https://inria.hal.science/inria-00510165>

Submitted on 13 Oct 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.




Graphics Gems Revisited

Fast and Physically-Based Rendering of Gemstones

Stephane Guy
PRIMA*

Cyril Soler
ARTIS*

GRAVIR/IMAG - INRIA



Abstract

We present an algorithm for rendering faceted colored gemstones in real time, using graphics hardware. Beyond the technical challenge of handling the complex behavior of light in such objects, a real time high quality rendering of gemstones has direct applications in the field of jewelry prototyping, which has now become a standard practice for replacing tedious (and less interactive) wax carving methods. Our solution is based on a number of controlled approximations of the physical phenomena involved when light enters a stone, which permit an implementation based on the most recent – yet commonly available – hardware features such as fragment programs, cube-mapping.

Keywords: Crystal optics, Hardware-based rendering, real time

1 Introduction

Gemstones are fascinating because they display many visually appealing phenomena thanks to their ability to alter light in a number of very specific ways. Examples include brilliance from internal reflections, fire due to dispersion of light, dichroism and doubling due to birefringence, color-shifting because of a camel-shaped absorbance spectrum, and darkening due to polarization. Unfortunately, the complexity of light interaction inside gemstones, due to their particular crystal structure, makes correct renderings of such objects very difficult to obtain.

Furthermore, our investigations among jewelry design packages show that computer aided prototyping has now become a standard [Doyle 2000]. Applications need a fast algorithm for rendering gemstones to allow the user to move and appreciate the variations in color and brightness over several views of the stone. The quality of the rendering is most important during the design of a piece hand-in-hand with a client who may want to see a high quality view of the expected result.

Up to now, the academic solutions which have been proposed for rendering gemstones have all been based on more or less complex ray tracing implementations, and therefore do not offer the possibility of real time display. To our knowledge, no commercial jewelry design software has such a capability. In the specifications of the two software packages *JewelSpace*TM (see www.jewel-space.net) and *JewelCAD*TM (see www.jacadc.com), for instance, one learns that the first

combines radiosity and ray tracing while the second uses OpenGL and ray tracing, which gives nice, but not instant, results.

We show in this paper that the convexity and polyhedral nature of faceted gemstones raises the possibility of an efficient hardware implementation, which we achieve by using the most recent hardware capabilities such as high level fragment programming [Mark et al. 2003], as well as more classical techniques such as cube mapping [Greene 1986]. It is indeed possible to rely on these tools for implementing pixel-based computation of Fresnel terms, directional light sampling and tone reproduction.

Our system introduces new possibilities of observing virtual gemstones: not only can the viewing conditions be changed in real time but also the physical properties (*e.g.*, color, refractive index) and even the geometry of the stone. Indeed, our management of rendering complexity offers the possibility to trade accuracy for speed, so as to maintain interactive rendering speed. Among applications, our algorithm could be a plugin for rendering faceted objects in a more complex scene as our results show. In terms of jewelry design, our contribution should be seen as a potential solution for real time examination of created objects with high rendering quality. For educational purposes it also permits to study interactively the visual properties of gemstone cuts in measured light conditions.

In section 3 we present a physical model for the behavior of light in colored gemstones. The importance of each phenomenon is then examined with a correct order of priority in terms of visual impact in Section 4, in order to derive a model suitable for hardware implementation. This implementation is presented in section 5. We finally give a number of results prior to concluding.

2 Related work

Yokoi *et al.* propose an algorithm [Yokoi et al. 1986] aimed particularly at the reproduction of asterism (*i.e.*, the effect seen on star sapphires) and chatoyancy effects. For this, a model for the dispersion of light rays into a distribution of microfacets is elaborated, which produces the expected result.

Yuan uses an adaptive ray tracing for adequately sampling the light spectrum for various wavelengths [Yuan et al. 1988], while keeping the overall complexity much lower than traditional dispersive ray tracing [Thomas 1986] and obtains fine images of diamonds. Dispersion is further investigated in [Sun et al. 2000b] and [Wilkie et al. 2000] for the rendering of glass-made objects.

Sun [Sun et al. 2000a] obtains images of colored diamonds using Fresnel reflection, volume absorption and light dispersion in a ray tracer. He also proposes a method to handle the non-linearity of absorption and the poor regularity of absorbance spectra using a composite spectral model.

These three papers constitute the few contributions in the field of computer graphics which aiming at rendering gemstones specifically. All these methods share a ray tracing basis with high computation time. None of them implements polarization which however incurs several visually important phenomena in gemstones.

Wolff proposes one of the first attempts to incorporate polarization effects in ray tracing [Wolff and Kurlander 1990], in the context

*Artis and PRIMA are teams within the GRAVIR/IMAG laboratory, a joint research unit of CNRS, INPG, INRIA, and UJF. {stephane.guy|cyril.soler}@imag.fr

of generalizing the Torrance-Sparrow reflectance model. For this, he uses the formalism of coherency matrices, introduced by Wolf in 1959 [Wolf 1959]. Another formalism (Stokes light vectors) was used by Wilkie in 2001 to get rid of complex numbers [Wilkie et al. 2001] and to incorporate polarization along with fluorescence effects in a ray tracer, with convincing results on semi-transparent objects. In 1994, Tannenbaum gave some details about the implementation of the birefringency phenomenon [Tannenbaum et al. 1994] for computing images of highly birefringent media such as calcite. As we explain below, birefringency is responsible for color variations in many stones.

3 Light propagation in gemstones

We start by presenting a physical model for light propagation in gemstones. This model will be used as a starting point for the approximations which lead to our simplified model (in Section 4), which is suitable for hardware implementation. It will also serve in a ray tracer (see Section 6) for validating these approximations.

Although the polarization state of light is, for most rendering applications, an expensive and unnecessary feature, it does play a critical role in the interaction of light with gemstones, and therefore must be incorporated into any rendering model that wants to reproduce these effects. Indeed, while the naked eye is not trained for detecting light polarization, the path of light through a faceted transparent object involves a number of selective changes in the polarization. This succession of changes is responsible for the darkening of some regions and, in the case of anisotropic crystal structures, a color change depending on the direction of propagation (Figure 3, left shows a combination of these two effects).

Let $\mathbf{E}$ be the electrical field of a monochromatic planar wave propagating along vector $\mathbf{s}$ at speed v and angular frequency ω . $\mathbf{E}$ is expressed in the plane orthogonal to $\mathbf{s}$ as space and time dependent 2D vector:

$$\mathbf{E}(\mathbf{r}, t) = \begin{bmatrix} E^\perp \cos(\omega(t - \mathbf{r} \cdot \mathbf{s}/v)) \\ E^\parallel \cos(\omega(t - \mathbf{r} \cdot \mathbf{s}/v) + \delta) \end{bmatrix} = \mathcal{R} \left(\begin{bmatrix} E^\perp e^{i(\omega(t - \mathbf{r} \cdot \mathbf{s}/v))} \\ E^\parallel e^{i(\omega(t - \mathbf{r} \cdot \mathbf{s}/v) + \delta)} \end{bmatrix} \right) \quad (1)$$

$E^\perp$ and $E^\parallel$ are the amplitudes of each component and δ is the phase shift between the two components of $\mathbf{E}$. As shown, $\mathbf{E}$ may alternatively be represented as the real part of a complex-valued field. For rendering, we are interested in the intensity of $\mathbf{E}$, which is

$$I = E^{\perp 2} + E^{\parallel 2}$$

3.1 Fresnel laws

Most crystal-structured materials are optically anisotropic (e.g., tourmaline, sapphire, but not diamond). This comes from the geometric asymmetry of the atoms arranged in the structure, which favors different charge transfers – i.e., different wavelength absorptions –, and different propagation speeds, depending on the direction of propagation. In the case of uniaxial crystals the medium is characterized by two indices of refraction n^o and n^e , and its optical axis $\mathbf{a}$. For a given wave propagation vector $\mathbf{s}$, we define the *crystal coordinate system* by its orthonormal basis X_a, Y_a, Z_a and the *principal plane* by the two vectors Y_a, Z_a :

$$\mathbf{X}_a(\mathbf{s}) = \frac{\mathbf{a} \times \mathbf{s}}{\|\mathbf{a} \times \mathbf{s}\|} \quad \mathbf{Y}_a(\mathbf{s}) = \frac{\mathbf{s} \times (\mathbf{a} \times \mathbf{s})}{\|\mathbf{s} \times (\mathbf{a} \times \mathbf{s})\|} \quad \mathbf{Z}_a(\mathbf{s}) = \mathbf{a} \quad (2)$$

Let ϵ be the dielectric tensor of the medium, defined by:

$$\epsilon = {}^t O \begin{pmatrix} n^o & 0 & 0 \\ 0 & n^o & 0 \\ 0 & 0 & n^e \end{pmatrix} O \quad \text{with} \quad O = \begin{pmatrix} \mathbf{X}_a(\mathbf{s}) \\ \mathbf{Y}_a(\mathbf{s}) \\ \mathbf{Z}_a(\mathbf{s}) \end{pmatrix}$$

The dielectric displacement vector D of the wave is related to the electric field E by:

$$\mathbf{D} = \epsilon \mathbf{E}$$

A light ray entering an anisotropic gemstone separates into two sub-rays linearly polarized along orthogonal directions (see Figure 1), called *ordinary ray*, or *o-ray*, and *extraordinary ray*, or *e-ray* [Born and Wolf 1999]. These rays belong to the only two categories of waves which are allowed to propagate in uniaxial crystals, and have their own characteristics:




Figure 1: Formation of an ordinary and extraordinary ray when an incident light ray encounters an anisotropic medium.

The two waves have respective directions of propagation $\mathbf{s}^o$ and $\mathbf{s}^e$. The directions of polarization are such that for the *o-ray*, $\mathbf{D}^o$ vibrates perpendicular to the principal plane (i.e., along $\mathbf{X}_a(\mathbf{s}^o)$), and for the *e-ray*, $\mathbf{D}^e$ vibrates parallel to the principal plane (i.e., along $\mathbf{Y}_a(\mathbf{s}^e)$). The electric field of the *o-ray* is orthogonal to the direction of propagation $\mathbf{s}^o$ of its wave, whereas that of the *e-ray* is not, because $\epsilon \mathbf{E}^e$ is not collinear to $\mathbf{D}^e$. The energy propagation vectors of the two waves (i.e., the directions used for ray tracing) are thus $\mathbf{t}^o = \mathbf{s}^o$ and $\mathbf{t}^e \neq \mathbf{s}^e$.

The speed v of the two rays can be computed from the speed c of light in vacuum, using [Born and Wolf 1999]:

$$\begin{aligned} v^o &= \frac{c}{n^o} \\ v^e &= \left(\left(\frac{c}{n^o} \right)^2 \cos^2 \theta + \left(\frac{c}{n^e} \right)^2 \sin^2 \theta \right)^{\frac{1}{2}} \quad \text{with} \quad \cos \theta = \mathbf{a} \cdot \mathbf{s}^e \end{aligned}$$

While the *o-ray* lies in the plane of incidence and obeys Snell's law, the transmitted *e-ray* does not. Its ray direction $\mathbf{t}^e$ can be computed using the following formula [Beyerle and McDermid 1998] in which $\mathbf{n}$ is the interface normal and n_1 the refractive index of the incident medium:

$$\begin{aligned} \mathbf{t}^e &= \frac{\gamma^2 (\mathbf{n}'^2 \mathbf{s} + (R - \mathbf{n}' \cdot \mathbf{s}) \mathbf{n})}{\|\gamma^2 (\mathbf{n}'^2 \mathbf{s} + (R - \mathbf{n}' \cdot \mathbf{s}) \mathbf{n})\|} \quad \text{where} \quad \gamma = \frac{1}{n^o n^e} \epsilon \quad \text{and} \\ R &= \left((\mathbf{n}' \cdot \mathbf{s})^2 - \mathbf{n}'^2 \mathbf{s}^2 + \mathbf{n}'^2 / n_1^2 \right)^{\frac{1}{2}} \quad \mathbf{n}' = \gamma \mathbf{n} \quad \mathbf{s}' = \gamma \mathbf{s} \end{aligned} \quad (3)$$

Similar to refractions, internal reflections inside an anisotropic crystal splits light waves into ordinary and extraordinary rays. While the direction of the former obeys the classical rule of reflection, the direction of the *e-ray* obtained by internal reflection must be computed using [Beyerle and McDermid 1998]:

$$\mathbf{t}_r^e = \frac{\gamma^2 (\mathbf{n}'^2 \mathbf{s}' + 2(\mathbf{s}' \cdot \mathbf{n}') \mathbf{n}')}{\|\gamma^2 (\mathbf{n}'^2 \mathbf{s}' + 2(\mathbf{s}' \cdot \mathbf{n}') \mathbf{n}')\|} \quad (4)$$

Finally, for a given wave refracting and reflecting at the interface between the air and an anisotropic medium, we need to express the Fresnel coefficients which permit us to compute the refracted and reflected fields. This happens when light enters or exits a gemstone, and depending on the case, reflected or refracted waves may either be the sum of an ordinary and extraordinary rays, or a unpolarized

wave. We treat both cases at once by considering the interface between two anisotropic media. If one medium is isotropic, the waves in this medium are still the sum of two orthogonally polarized wave components and can thus be represented as a 'o-ray' and a 'e-ray' with the same direction of propagation with arbitrary choice of the 'optical axis' of the medium. We take the convention that subscripts **t** and **r** stand for *transmitted* and *reflected* fields, while **o** and **e** stand for *ordinary* and *extraordinary*.

Let $\mathbf{E}_i$ be an incident linearly polarized wave of amplitude E_i , vibrating direction $\mathbf{e}_i$, and speed $v_i = \frac{c}{n_i}$. We want to find the coefficients $\alpha_r^o, \alpha_r^e, \alpha_t^o, \alpha_t^e$ by which to multiply E_i to obtain the amplitudes of the four reflected and refracted fields.

Maxwell theory requires that the tangential components of the total electric field and the magnetic vector $\mathbf{H} = \frac{c}{\mu} \mathbf{s} \times \mathbf{E}$ be continuous across the surface (where μ is the magnetic permissivity, supposed identical in both media):

$$\begin{aligned} \mathbf{n} \times (\mathbf{E}_i + \mathbf{E}_r^o + \mathbf{E}_r^e) &= \mathbf{n} \times (\mathbf{E}_t^o + \mathbf{E}_t^e) \\ \mathbf{n} \times (\mathbf{H}_i + \mathbf{H}_r^o + \mathbf{H}_r^e) &= \mathbf{n} \times (\mathbf{H}_t^o + \mathbf{H}_t^e) \end{aligned}$$

Let $\mathbf{e}_r^o, \mathbf{e}_r^e, \mathbf{e}_t^o, \mathbf{e}_t^e$ and $\mathbf{h}_r^o, \mathbf{h}_r^e, \mathbf{h}_t^o, \mathbf{h}_t^e$ be the respective vibrating directions of the electrical and magnetic fields. By expressing each component $\mathbf{E}_k^l$ (resp. $\mathbf{H}_k^l$) as $\alpha_k^l E_i \mathbf{e}_k^l$ (resp. $\alpha_k^l E_i n_k^l \mathbf{h}_k^l$, with $\mathbf{h}_k^l = \mathbf{s}_k \times \mathbf{e}_k^l / \|\mathbf{s}_k \times \mathbf{e}_k^l\|$), and computing the dot product of these equations with two independent vectors $\mathbf{v}_1$ and $\mathbf{v}_2$ in the interface plane, one obtains [C.McClain et al. 1993]:

$$\begin{bmatrix} -\mathbf{v}_1 \cdot \mathbf{e}_r^o & -\mathbf{v}_1 \cdot \mathbf{e}_r^e & \mathbf{v}_1 \cdot \mathbf{e}_t^o & \mathbf{v}_1 \cdot \mathbf{e}_t^e \\ -\mathbf{v}_2 \cdot \mathbf{e}_r^o & -\mathbf{v}_2 \cdot \mathbf{e}_r^e & \mathbf{v}_2 \cdot \mathbf{e}_t^o & \mathbf{v}_2 \cdot \mathbf{e}_t^e \\ -n_r^o \mathbf{v}_1 \cdot \mathbf{h}_r^o & -n_r^e \mathbf{v}_1 \cdot \mathbf{h}_r^e & n_t^o \mathbf{v}_1 \cdot \mathbf{h}_t^o & n_t^e \mathbf{v}_1 \cdot \mathbf{h}_t^e \\ -n_r^o \mathbf{v}_2 \cdot \mathbf{h}_r^o & -n_r^e \mathbf{v}_2 \cdot \mathbf{h}_r^e & n_t^o \mathbf{v}_2 \cdot \mathbf{h}_t^o & n_t^e \mathbf{v}_2 \cdot \mathbf{h}_t^e \end{bmatrix} \begin{bmatrix} \alpha_r^o \\ \alpha_r^e \\ \alpha_t^o \\ \alpha_t^e \end{bmatrix} = \begin{bmatrix} \mathbf{v}_1 \cdot \mathbf{e}_i \\ \mathbf{v}_2 \cdot \mathbf{e}_i \\ n_i \mathbf{v}_1 \cdot \mathbf{h}_i \\ n_i \mathbf{v}_2 \cdot \mathbf{h}_i \end{bmatrix} \quad (5)$$

This linear system can be solved numerically. In the case of total reflection it should first be simplified into a 2x2 system by suppressing the α_t^o and α_t^e unknowns. Up to now $\mathbf{E}_i$ has been supposed linearly polarized, which applies for *e*-rays and *o*-rays inside an anisotropic medium. If the incident medium is isotropic, the two solutions corresponding to setting $\mathbf{e}_i$ orthogonal and then parallel to the incident plane gives the 8 needed coefficients. If both media are chosen isotropic, and if the 'optical axis' is set to be the normal of the interface, we have checked that the obtained solution corresponds to the well known Fresnel formulas given in Figure 2.

$$\begin{aligned} F_r^\perp &= \alpha_r^{oo} = -\sin(\theta_i - \theta_t) \sin(\theta_i + \theta_t)^{-1} \\ F_r^\parallel &= \alpha_r^{ee} = \tan(\theta_i - \theta_t) \tan(\theta_i + \theta_t)^{-1} \\ F_t^\perp &= \alpha_t^{oo} = 1 + F_r^\perp \\ F_t^\parallel &= \alpha_t^{ee} = (1 - F_r^\parallel) \cos \theta_i \cos \theta_t^{-1} \end{aligned}$$

Figure 2: Geometric configuration and Fresnel coefficients for a light ray refracting and reflecting at the interface between two isotropic media. The crossed coefficients $\alpha_r^{oe}, \alpha_r^{eo}, \alpha_t^{oe}$ and α_t^{eo} are null in this particular case.

3.2 Absorption

As a general rule, a fraction of the light traveling inside a transparent material is absorbed. The absorption along a path (x_0, x) is ruled by the *Bouguer-Lambertian* law [Wyszecki and Stiles 1982], giving the resulting intensity after a distance d :

$$I_\lambda(x) = I_\lambda(x_0) e^{-\kappa(\lambda)d} \quad (6)$$

κ is called the *absorbance* of the medium and depends on the wavelength. Absorption is responsible for the color of gemstones, and the absorbance spectrum acts as a tracer for the chemical nature, geographic source and applied treatments of a stone.




Figure 3: *left*: Photo of pleochroism in a green tourmaline. The optical axis is left-to-right oriented. Rays entering on the left side perpendicular to the page have most of their path parallel to the optical axis before they come out to the eye from the right side. They are thus mainly attenuated by the $E_\parallel$ (mainly green) spectrum. *Photo: Wimon Manorkul/Pala International (www.palagems.com); used with permission.* *Right*: absorbance spectra of tourmaline for waves components polarized along/perpendicular to the optical axis.

Here again the optical anisotropy of gemstones plays an important role: because the *o*-rays and *e*-rays have different polarization directions, they are absorbed differently by the stone. The following formula [Born and Wolf 1999] gives the absorbance for the two rays:

$$\begin{aligned} \kappa^o &= K^o \\ \kappa^e &= K^o \cos^2 \theta + K^e \left(\frac{n^o}{n^e}\right)^2 \sin^2 \theta \end{aligned} \quad (7)$$

In this formula, K^o and K^e are characteristic constants of the medium, and $\cos \theta = \mathbf{s}^e \cdot \mathbf{a}$ is the cosine of the angle between the extraordinary wave propagation direction and the optical axis of the stone.

Depending on the angle with which a ray enters a birefringent gemstone, the refracted *e*-ray will have variable direction; its absorption and color contribution to the outgoing light will consequently vary. This phenomena, called *pleochroism*, is illustrated on Figure 3 for a tourmaline gemstone.

Many gemstones display this behavior (e.g., sapphires, rubies, tourmalines [Hughes 1997; Nassau 2001]). Other gemstones (like andalusite) display three distinct colors. This comes from their crystal structure, which has two optical axes instead of one. The computation of absorption in this context is beyond the scope of this paper. Some of these stones, however, behave visually as if they were uniaxial (e.g., peridot and andalusite) because two of the three absorbance spectra are nearly identical.

The model presented above represents the state of the art in the understanding of light transport in gemstones, except for two deliberate approximations: First, the model treats wavelengths independently, and therefore cannot represent fluorescence in gemstones. Although there is fluorescence in some gems (e.g., some sapphires, synthetic diamonds), it is weak and mainly affects UV light outside the visible spectrum, so we ignore it. Second, in absorption, the polarizations of the *e*-ray and *o*-rays are in fact very slightly elliptical rather than strictly linear [Born and Wolf 1999]; we nonetheless treat them as linear.

4 Adopted model

At this point of the paper, we discuss the importance and priorities that should be given to the phenomena previously described, keeping in mind a possible hardware implementation. In section 4.1 we justify our choices for representing color as three separate wavelengths, and examine in section 4.2 how light will be represented for each wavelength.

4.1 Representation of color

Choosing a good representation of color is usually a compromise between visual accuracy of the result and a bulky set of coefficients.

Because the constraints of the hardware, we have chosen to work with 3 color components. Richer spectral rendering [Peercy 1993] could still be achieved by adding more passes and a final reconstruction, at the expense of rendering time [Sun et al. 1999].

We thus limit ourselves to spectral sampling, although the non linear relationship between absorbance and transmittance tends to saturate colors when darkening them [Sun et al. 1999]. Such an approximation on the facility to approach the attenuation by its linear counterpart is justified below.

One other important aspect of spectral sampling is to correctly choose the absorbance coefficients κ_r, κ_g and κ_b for the R, G and B channels. Directly reading spectral absorption coefficients from the spectral absorbance curve at the exact wavelengths of red (700.0nm), green (546.1nm) and blue (435.8nm) introduces significant errors due to peaks in the absorbance curve [Sun et al. 1999]. We thus need a way to extract significant enough absorbance values from the stone's spectral absorbance curves. Inspired by existing approximations for reflectance [Borges 1991], we propose the following:

The R, G, B color components perceived by human eye for a given spectrum S are computed using the color matching functions $\bar{r}, \bar{g}$ and $\bar{b}$ of λ [Wyszecki and Stiles 1982] by:

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \int_{\lambda} \begin{bmatrix} \bar{r}(\lambda) \\ \bar{g}(\lambda) \\ \bar{b}(\lambda) \end{bmatrix} S(\lambda) d\lambda$$

Following Equation 6, a spectrum S_0 corresponding to white light will transform, after a path of length x , into

$$S(x, \lambda) = S_0(\lambda) e^{-\kappa(\lambda)x} \quad \text{e.g.,} \quad R(x) = \int_{\lambda} \bar{r}(\lambda) S_0(\lambda) e^{-\kappa(\lambda)x} d\lambda \quad (8)$$

We are looking for an absorbance value κ_r such that $R(x) = R(0) e^{-\kappa_r x}$ approximates equation 8 for small values of x . We thus take:

$$\kappa_r = -\frac{1}{x} \ln \frac{R(x)}{R(0)} = -\frac{1}{x} \ln \left[\frac{1}{R(0)} \int_{\lambda} S_0(\lambda) \bar{r}(\lambda) e^{-\kappa(\lambda)x} d\lambda \right]$$

which, for small distance values x , is approximated by:

$$\kappa_r = \frac{1}{R_0} \int_{\lambda} \kappa(\lambda) \bar{r}(\lambda) S_0(\lambda) d\lambda \quad \text{using} \quad R_0 = \int_{\lambda} \bar{r}(\lambda) S_0(\lambda) d\lambda \quad (9)$$

Proceeding identically for the green and blue components we obtain suitable absorbance coefficients from the absorbance spectra, while avoiding artifacts of direct sampling of absorbance functions with peaks.

The error of the above approximation for $\kappa_r, \kappa_g, \kappa_b$ depends on the extent to which absorption differs from its linear approximation in Equation 9. The error is thus small because the diameter L of the stone, times the absorbance κ is small (the absorbance is computed piecewise between successive internal reflections of light). For typical values of $L = 1\text{cm}$ and $\kappa = 0.4$ for instance, we get an error of $\frac{e^{-\kappa L} - (1 - \kappa L)}{e^{-\kappa L}} = 0.00125\%$

4.2 Representation of monochromatic light

We adopt the formalism of *coherency matrices* [Glassner 1995; Wolf 1959] for representing the intensity and the polarization state of the electric field along a ray of light. The coherency matrix of a field $\mathbf{E} = (e^{\perp}(t), e^{\parallel}(t))$ is defined as

$$J = \begin{bmatrix} \langle e^{\perp} e^{\perp*} \rangle & \langle e^{\perp} e^{\parallel*} \rangle \\ \langle e^{\parallel} e^{\perp*} \rangle & \langle e^{\parallel} e^{\parallel*} \rangle \end{bmatrix} = \langle E E^* \rangle = \begin{bmatrix} J_{xx} & J_{xy} \\ J_{yx} & J_{yy} \end{bmatrix}$$

where E is the complex representation of $\mathbf{E}$, E^* is the conjugate transpose of E , and $\langle u \rangle$ denotes the mean value of u over time. The intensity of the field is given by $I = J_{xx} + J_{yy}$. From this, the

coherency matrix of an incoherent (non polarized) light ray of intensity I_0 is [Wolf 1959]:

$$J_{\text{incoherent}} = \frac{1}{2} I_0 \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

For any linear transformation M (also called a *modifier matrix*) applied to $\mathbf{E}$, the corresponding coherency matrix becomes, from the definition of J , $J' = M J M^*$. It is thus possible to compute the matrix J along a ray by applying successive modifier matrices corresponding to refraction and internal reflections on the faces of the gemstone, and rotations to account for the change in coordinate systems between two successive media interfaces. The matrices involved for refraction, reflection and rotation with an angle of θ , are respectively:

$$M^t = \begin{bmatrix} \alpha_{tr}^{oo} & \alpha_{tr}^{oe} \\ \alpha_{tr}^{eo} & \alpha_{tr}^{ee} \end{bmatrix} \quad M^r = \begin{bmatrix} \alpha_{rr}^{oo} & \alpha_{rr}^{oe} \\ \alpha_{rr}^{eo} & \alpha_{rr}^{ee} \end{bmatrix} \quad R_{\theta} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$$

In birefringent gemstones, the direction of the ordinary ray is easily computed using Snell's law, but the intervention of a different direction for the extraordinary refracted and reflected rays make the computation cost grow exponentially with the depth of paths inside the stone. We computed the maximum angle between the o -ray and e -ray for all incidence angles in $[0, \pi/2]$ and any orientation of the optical axis, for common gemstones (zircon being an extreme case):

Material	Zircon	Tourmaline	Sapphire	Emerald
Angle (deg.)	2.328	0.767	0.45	0.324
Error	4.06 %	1.34 %	0.785 %	0.565 %

The "error" field gives the distance between the two images of an edge on the far side of the stone, as a fraction of the total size of the stone. For a 1cm Zircon, for instance, the two images would be 0.5mm apart; on a 1000-pixel image, this would 50 pixels, but this is the extreme case (by far). The validity of this approximation is confirmed by the photos of the tourmaline gemstones on Figures 3 and 11: no doubling of the edges seen by refraction is perceptible whereas in the orientation of the optical axis corresponds to a case of maximum deviation in Figure 3.

Computing Fresnel coefficients using Equation 5 can not conveniently be implemented in graphics hardware. Contrarywise, Fresnel coefficients at the interface of two isotropic media (given in Figure 2) can be tabulated very efficiently. Because o -rays and e -rays propagate along close directions and have orthogonal polarization, the generalized Fresnel coefficients are very close to the isotropic ones when the difference $n_e - n_o$ is small, if expressed into the same coordinate systems. Let $R_{\theta_i}, R_{\theta_r}, R_{\theta_t}$ be the rotation matrices which align the coordinate systems for the isotropic Fresnel coefficients of Figure 2 to the corresponding implicit coordinate systems of the general Fresnel coefficients of equation 5 in the coordinate system on the principal planes (Equation 2). We have:

$$M^t \approx R_{\theta_i} \begin{bmatrix} F_{\parallel}^t & 0 \\ 0 & F_{\perp}^t \end{bmatrix} R_{\theta_t} \quad \text{and} \quad M^r \approx R_{\theta_i} \begin{bmatrix} F_{\parallel}^r & 0 \\ 0 & F_{\perp}^r \end{bmatrix} R_{\theta_r} \quad (10)$$

Equality holds when $n^o = n^e$. For instance let's consider the case of a ray entering an anisotropic medium of optical axis $\mathbf{a}$ from an isotropic medium. $\mathbf{n}$ is the interface normal and $\mathbf{s}_t$ the propagation vector of the transmitted wave. Generalized Fresnel coefficients already relate the reflected field to the incident field in the same coordinate system than in the isotropic case, but the transmitted field is expressed in the coordinate system with the orthogonal vibrating direction given by $\mathbf{a} \times \mathbf{s}$, so one should take:

$$\theta_i = 0, \quad \theta_r = 0, \quad \cos \theta_t = \mathbf{n} \times \mathbf{s}_t \cdot \mathbf{a} \times \mathbf{s}_t, \quad \sin \theta_t = \mathbf{n} \times \mathbf{s}_t \cdot \mathbf{s}_t \times (\mathbf{a} \times \mathbf{s}_t)$$

The proposed approximation works very well, even on zircon, as illustrated on figure 4.




Figure 4: The colored curves present the 8 Fresnel coefficients $\alpha_{t|r}^{[o|e][o|e]}$ at the interface between the air and a birefringent medium (zircon). The black-dotted curves represent their approximations obtained through equations 10. The interface is the plane $z = 0$, the optical axis is $\mathbf{a} = (\cos(\pi/6), 0, \sin(\pi/6))$ and the incidence plane is chosen so as to correspond to the case of highest deviation between the o -rays and the e -rays.

Contrarywise, the change in absorption due to the polarization of rays with respect to the optical axis of the stone dramatically affects the rendering color and should not be neglected (see Figure 3).

Because the waves along the o -ray and the e -ray are polarized orthogonally to each other, we can represent them using a single coherency matrix in the coordinate system based on the ordinary and extraordinary vibration directions and their – supposed shared – direction of propagation $\mathbf{s}$. In this coordinate system, the matrix of a wave propagating inside the crystal is therefore diagonal, and is attenuated by a diagonal modifier matrix depending on the direction of propagation, using the attenuation coefficients of Equation 7:

$$A(s) = \begin{bmatrix} e^{-\kappa^o(s)l} & 0 \\ 0 & e^{-\kappa^e(s)l} \end{bmatrix} \quad (11)$$

In the case of isotropic crystals (e.g., diamonds, garnets) the matrix $A(s)$ becomes identity times the attenuation given by Equation 6.

5 Rendering algorithm

Figure 5 shows the path of the light obtained by tracing a ray from the eye to a point on the stone. To compute the resulting intensity J_0 along such a path, we need to add the contributions of light at each interface P_k between the stone and the air, and accounting for attenuation $A_{k \rightarrow k+1}$ along segments $[P_k, P_{k+1}]$ inside the stone. Denoting by J_k the coherency matrix of the light from point P_k in the path, we have:

$$\begin{aligned} J_0 &= M_0^r J_0^i M_0^{r*} + M_0^t A_{0 \rightarrow 1} J_1 A_{0 \rightarrow 1}^* M_0^{t*} \\ J_k &= M_k^r J_k^i M_k^{r*} + M_k^t A_{k \rightarrow k+1} J_{k+1} A_{k \rightarrow k+1}^* M_k^{t*} \end{aligned} \quad (12)$$

For rendering a gemstone using classical ray tracing, one would collect these contributions from back to front (i.e., at P_n, P_{n-1} and finally P_0), transforming rays at each successive interface to account for refraction, external reflection, or internal reflections encountered along the path.

Our hardware-based algorithm relies on the fact that, for a given depth k , the set of points P_k that contribute to the image through the same succession of transformations, can be rendered at the same time using a fragment program. Such a set is called a *facet*. We regroup facets in a tree structure called the *facet tree*. Each node of the facet tree at depth k contains a product of k transformations and a region of points included in one face of the gemstone.




Figure 5: Different light intensities $J_{0,1,2,3,4}^i$ contribute to the final intensity at point P_0 along incident directions associated to external reflections (at P_0) and internal refraction (at P_1, P_3, P_4). The ray is followed until the attenuation lowers significantly the chances of missing a hotspot in the incoming light.

Rendering a gemstone from a given viewpoint thus requires for each frame to (1) build the facet tree corresponding to current viewpoint and (2) accumulate contributions of all facets of the facet tree, from back to front, using a fragment program. The lighting environment used for rendering gemstones is stored in a cube-map [NVIDIA Corporation 2000]. As illustrated in the result section, rendering gemstones needs to deal with high dynamic range lighting and rendering. All computations are thus performed with floating point values and a final tone reproduction pass (3) turns the result of pass (2) into a 8-bits rgb image. Pass (1), (2) and (3) are detailed in Sections 5.1, 5.2 and 5.3.

When changing the viewpoint the facet tree changes and must therefore be updated. This means that not only the viewpoint, but also the geometry of the gemstone as well as its physical parameters (refractive indices, orientation of optical axis, attenuation) can arbitrarily be changed during rendering, at the same cost.

5.1 Pass 1: construction of the facet tree

Because refraction is not in general a linear transform, the facets as defined above are not polygons. However the nonlinearity does not noticeably affect the refracted images of short segments such as the edges of a gemstone, as shown on figure 6, and we adopt, for representing refraction through a front face of the gemstone, the linear approximation proposed by Heckbert for beam tracing [Heckbert and Hanrahan 1984]. To each facet is thus associated a fictive viewpoint. Note however, that when we render the facet tree, the refraction direction of incoming light I_k at point $P_{k \neq 0}$ will be computed exactly by the fragment shader. Approximations of refraction only affect the point P_0 .



Figure 6: *Left*: image computed using exact refraction with our ray tracer. *Center*: image computed with graphics hardware algorithm and linearized refraction. *Right*: difference image.

Thanks to this approximation and to the linearity of internal reflections on the gemstone faces, each facet is an actual polygon and a subset of a single face of the gemstone. We compute facets with the OpenGL feedback buffer using the following algorithm:

At level 0 of the tree the facets are the faces of the gemstone polyhedron directly seen from the viewpoint. At level 1 the child facets of level 0 facet f_i^0 correspond to the intersection of a beam traced from the viewpoint through f_i^0 with the gemstone transformed by refraction. At subsequent levels the gemstone is further transformed

by reflection through the support face of the parent facet and clipped with this facet.

The window coordinates of each facet are computed by rendering the initial geometry of the gemstone (stored in a *display list*) into the *OpenGL* feedback buffer [Neider et al. 1993], using the appropriate viewpoint and transformation, while inserting clipping planes corresponding to the edges of the parent facet as shown on Figure 7.




Figure 7: Construction of the *facet tree*, using a fictive viewpoint for linearized refraction and successive symmetries for internal reflections. For level k the gemstone model is transformed through reflection/refraction around the support plane of the parent facet, then drawn in the *OpenGL* feedback buffer, and clipped by the edges of the parent facet. The resulting mesh is displayed on the right, using the same color for facets of the same level.

To minimize costly access to the feedback buffer, the facet tree is constructed breadth-first. The depth of the tree is limited by three conditions: (1) the total path length to a facet attenuates light strongly enough to reduce the resulting intensity under a fixed value; (2) the area of the facet is smaller than a user defined threshold; (3) the computation time down to the current level exceeds the frame-rate limit.

Because refraction indices depend on wavelengths, the facet tree should not be the same for the red, green and blue components. Depending on the intensity of dispersion, we compute either one single facet tree corresponding to an average refractive index, or three distinct facet trees, at the expense of frame rate. However, the fragment shader which computes the refraction of light at points $P_{1,...,n}$ still uses the correct indices of refraction. We thus still achieve in the worst case an approximation of dispersion effects, as shown on Figure 8.



Figure 8: Image of a highly dispersive stone, computed using exact refraction with our ray tracer (top left) and with our hardware algorithm with one facet tree (top right). Left: a difference image shows that some—but not all, as explained in the text—rainbow effects are missed by our algorithm in this case.

For a given maximum depth, fully building the facet tree is not in fact necessary: during this operation, we estimate the cumulative effect of attenuation and Fresnel coefficients for a single point on each facet down to the next level. We locally prune the construction if the contribution of next facet will be less than a threshold in percentage to the accumulated energy before the current facet. Section 6.2 shows a practical example of this.

5.2 Pass 2: rendering the facet tree

We implemented and tested the rendering algorithm of the facet tree on a *Nvidia GeForceFX 5900* using *Cg* programming [Lindholm et al. 2001].

Figure 9 summarizes our implementation: the facet tree is traversed breadth-first, by decreasing order of level. At a given level $k > 0$, a P -buffer is filled with the result of the *internal reflection* fragment program, which computes the exact refraction of light entering the gemstone from the cube-map, as well as the internal reflection of the contribution of level $k + 1$ stored in the accumulation buffer for level $k - 1$. At level 0 the *external reflection* fragment program is used instead, for adding the result of the previous calls seen by refraction, with the light reflecting on the faces directly visible. Both fragment programs compute the attenuation using Equations 7 and 11. The path length is obtained by subtracting the distance from the fictive viewpoint to the current pixel (c.f. Figure 7) to that of the pixels in the previous result. The Fresnel coefficients are computed using equation 10, from the formulas of Figure 2, tabulated in 1D float textures.




Figure 9: Implementation using *Cg* programming. See text.

The incoming light is represented using a HDR cube-map composed of three traditional 4×8 -bits cube-maps, one per color channel, each representing 32-bits fixed point values. The computation by the fragment programs is performed using floating point numbers. We therefore can not simply blend the contributions of facets into the (8-bits RGBA) frame buffer. This justifies the use of a 32-bits floating point RGBA P -buffer. Fresnel relationships being always expressed in coordinate systems where the matrices are diagonal, we only need to store two components per channel (E^o , E^e for anisotropic media, resp. $E^\perp$, $E^\parallel$ for isotropic ones). Each component being represented as a 16-bits *half float* [Bogart et al. 2003], one can fit in the *RGB* channels of the P -buffer the electric fields for red, green and blue. This encoding requires to *un-pack* the values before using them. The *A* channel contains the distance of the pixel to the fictive viewpoint at the last treated level in the facet tree.

5.3 Pass 3: tone reproduction

Tone reproduction is achieved entirely in hardware using a publicly available shader of Industrial Light+Magic [Bogart et al. 2003], to display floating point images stored in the *OpenEXR* format. Such an approach makes comparisons between ray-traced images, pictures of gemstones, and images rendered using the hardware method very convenient, since the first two are produced in *OpenEXR* format and displayed using the same shader as the one used in the hardware implementation. We also tested hardware based glare effect [Spencer et al. 1995; Debevec 1998] on our HDR images with success, as in the teaser on page 1, and the diamond on Figure 11. Glare effect was not used anywhere else in the paper, to prevent masking important artifacts and details of the method.

6 Results

6.1 Modus operandi

The work presented may seem hard to reproduce without a clear road map, which we give here.

Using the reference ray tracing algorithm requires working with coherency matrices. Directions of the ordinary reflected and transmitted o -rays can be computed using Snell's law, while those of the e -rays are given by Equations 3 and 4. Absorption is computed by Equations 7 and 11. General Fresnel coefficients at each interface are obtained by solving the 4×4 linear system in Equation 5.

Our hardware based algorithm differs from a standard ray tracer in that ray directions always follow Snell's law; Fresnel coefficients are obtained using Equation 10, and refraction is linearized as explained in Section 5.1.


Both methods use equation 12 for composing coherency matrices at media interfaces, and use the same set of parameters, as described in the following three paragraphs:

System parameters We detail in the table below the parameters we used for our experiments. One needs the nature of the medium ('I'=isotropic, 'U'=birefringent uniaxial, 'B'=biaxial) and the refractive indices n^o and n^e when applicable. Values are indicated for the green channel whereas red and blue are obtained after adding or subtracting the value after ' $\pm$ ' (half the *dispersion* value found in usual tables). Tricolored absorbance values are computed using Equation 9 from the absorbance spectra of each stone¹.

Gemstone	Type	$(K_r, K_g, K_b)_{ole}$	Colors	n_o, n_e
Garnet	I	(0.136, 0.153, 0.175)	orange-red	$1.730 \pm .014$
Tourmaline (Dravite)	U	(0.033, 0.034, 0.082) _o (0.010, 0.076, 0.015) _e	yellow green blue green	$1.642 \pm .011$ $1.619 \pm .011$
Peridot	B	(0.023, 0.015, 0.051) _o (0.011, 0.003, 0.028) _e	green yellow green	$1.680 \pm .010$ $1.640 \pm .010$
Diamond	I	(0.001, 0.001, 0.001)	white	$2.410 \pm .022$
Sapphire	U	(0.332, 0.270, 0.156) _e (0.165, 0.147, 0.185) _o	light blue violetish blue	$1.768 \pm .009$ $1.760 \pm .009$
Andalusite	B	(0.0056, .006, .0183) _o (0.170, 0.175, 0.257) _e	greenish red yellowish green	$1.635 \pm .005$ $1.644 \pm .005$

Geometric models Models of standard gemstone cuts are readily available on the internet². However, problems arise when a particular gemstone needs to be simulated, as in our comparisons with photographs. While expensive gemstones (e.g., very clear and fine quality stones) tend to be cut using very precise standards, more affordable pieces often display an *ad-hoc* cut so as to respect constraints such as avoiding inclusions. Laser-scanning gemstones is not applicable due to both their size and specularity. We also tried X-ray tomography with very poor results.

The solution we used is based on manually designing the mesh and applying an iterative relaxation algorithm on vertex positions so as to fit symmetries, distance and planarity constraints measured on the stones. This works well provided that the hand-designed mesh is not too far from the real object. The model *at right* was designed that way and corresponds to the tourmaline of Figure 11.



Acquisition of light We used traditional methods for acquiring cube-maps. Photographs of a mirrored ball were taken along two orthogonal directions and using a set of 10 different exposures. Each set is first converted into a HDR image in Radiance format. The two HDR images are combined using HDRShopTM

1 - e.g., at <http://minerals.gps.caltech.edu/FILES/Visible/>
2 - e.g., at <http://www.3dlapidary.com/>

(www.debevec.com/HDRShop), and the result saved as a cube-map. We proceed identically to obtain HDR images of sample gemstones in the same lighting environment.

6.2 Additional validation tests





Figure 10: (Bars charts, scale at right:) Computation times in milliseconds for the construction of the facet tree only and for the full algorithm. (Lines, scale at left:) number of facets in the facet tree as a function of the number of internal reflections, (1) in the case where the facet tree is full down to the requested reflection level, and (2) when the facet tree is optimized for an error bound of 4%. Two different models were used: model 1 (130 vertices, 115 faces) and 2 (120 vertices, 93 faces). The picture row shows, for model 1, the raw and optimized facet trees as well as the corresponding images. Parameters correspond to violet/blue sapphire.

Computation times In Figure 10 we present the computation times (in *ms*) and the numbers of facets in the resulting facet trees for two different models. This experiment shows that cleverly pruning the facet tree makes the computation more efficient. Typical framerates for less complex models (40 to 80 facets), in a 800×600 window, range from 12 to 30 *fps* at depths 2 and 3.


Comparison to real images On the top-right of Figure 11 one can see a comparison between a simulated tourmaline and a picture of a real one. The simulation was performed using the absorbance values and the geometric model displayed in Section 6.1, and captured lighting conditions.

Because it was impossible to precisely duplicate the geometry of the gems and the lighting conditions, and because the appearance of gems is very sensitive to both, one cannot expect a pixel-by-pixel correspondence between the images; one can, on the other hand, evaluate the phenomenological similarities and differences quite well. Both images display a change in color along the optical axis (oriented at approximately 60° in the image plane) hence the bluish tint on the bottom of each image (this effect, also seen on stones (e), (f) and (g) proves the need to account for birefringency in the simulation). Color ranges are quite similar, as well as flash effects (from a qualitative point of view) thanks to the HDR rendering.

However one can complain that the luster of the stone (responsible for the iridescent colors on the top faces in the photography)




Grenat for different depths of the facet tree. (a) no internal reflection; (b) single internal reflection; (c) and (d): 2 and 3 reflections; Note that after 2 reflections, very little is added to the image.




Andalusite. Optical axis (e) toward the eye, (f) updown and (g) left-to-right.



Ring with a peridot



Left: real tourmaline, right: simulated stone with same model and similar lighting condition and viewpoint. See Section 6.2.



Example of application in Jewellery prototyping (diamond).



Figure 11: Various examples of results from our hardware rendering algorithm.

is not simulated, which would be a challenging problem to solve using graphics hardware.

7 Conclusion

We have presented a solution for rendering faceted colored gemstones based on the analysis of optical phenomena, their simplification and their adaptation to a hardware-based implementation. We have shown – we think for the first time – that it is possible to obtain some very visually pleasant and mostly physically correct images of faceted gemstones in real time by using a number of wise approximations of the real phenomena.

Our implementation benefits from the high level programming of today's graphic cards [Mark et al. 2003], which allows a comprehensive and portable implementation. Moreover, the rendering speed of our system offers new possibility for observing gemstones: not only the viewpoint can be changed in real time, but also the stone physical properties and geometry.

Our model can easily be extended to biaxial media. Other important features of gemstones, such as inclusions, color zoning and luster would be interesting to incorporate as well.

Acknowledgments

Many thanks to John Hughes for his careful reading of the paper, to Laurence Boissieux for the jewellery models and to the reviewers for helpful suggestions.

References

- BEYERLE, G., AND McDERMID, I. S. 1998. Ray tracing formulas for refraction and internal reflection in uniaxial crystals. *Applied Optics* 37, 34 (Dec.), 7947–7953.
- BOGART, R., KAINZ, F., AND HESS, D. 2003. Openexr image file format. *Siggraph Technical Sketches and Applications* (July).
- BORGES, C. F. 1991. Trichromatic approximation for computer graphics illumination models. In *Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, ACM Press, Eurographics, 101–104.
- BORN, M., AND WOLF, E. 1999. *Principles of Optics, Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press.
- C. McCLAIN, S., HILLMAN, L. W., AND CHIPMAN, R. A. 1993. Polarization ray tracing in anisotropic optically active media. ii. theory and physics. *Journal of Optical Society of America* 10, 11 (Nov.).
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, ACM Press, 189–198.
- DOYLE, A. 2000. Screen gems: Cad/cam technologies are replacing traditional methods of jewelry making. *Computer Graphics World* (July).
- GLASSNER, A. S. 1995. *Principles of Digital Image Synthesis*. Morgan Kaufmann.

- GREENE, N. 1986. Environment mapping and other applications of world projections. *IEEE Comput. Graph. Appl.* 6, 11, 21–29.
- HECKBERT, P. S., AND HANRAHAN, P. 1984. Beam tracing polygonal objects. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, H. Christiansen, Ed., vol. 18.
- HUGHES, R. W. 1997. *Ruby & Sapphire*. RWH Publishing.
- LINDHOLM, E., KILGARD, M. J., AND MORETON, H. 2001. A user-programmable vertex engine. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ACM Press, 149–158.
- MARK, W. R., GLANVILLE, R. S., AKELEY, K., AND KILGARD, M. J. 2003. Cg: a system for programming graphics hardware in a c-like language. *ACM Trans. Graph.* 22, 3, 896–907.
- NASSAU, K. 2001. *The Physics and Chemistry of Colour*. John Wiley & Sons.
- NEIDER, J., DAVIS, T., AND WOO, M. 1993. *The OpenGL Programming Guide – OpenGL Version 1.2*. Addison-Wesley. Third Edition.
- NVIDIA CORPORATION. 2000. Perfect reflections and specular lighting effects with cube environment mapping. Tech. rep. <http://developer.nvidia.com/>.
- PEERCY, M. S. 1993. Linear color representations for full spectral rendering. In *Comp. Graphics (SIGGRAPH '93 Proceedings)*, J. T. Kajiya, Ed., vol. 27, 191–198.
- SPENCER, G., SHIRLEY, P., ZIMMERMAN, K., AND GREENBERG, D. P. 1995. Physically-based glare effects for digital images. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, ACM Press.
- SUN, Y., FRACCHIA, F. D., AND DREW, M. S. 1999. Rendering the phenomena of volume absorption in homogeneous transparent materials. In *2nd Annual IASTED International Conference on Computer Graphics and Imaging (CGIM'99)*, 283–288. <http://fas.sfu.ca/pub/cs/mark/Cgim99/volumeAbs.ps.gz>.
- SUN, Y., FRACCHIA, F. D., AND DREW, M. S. 2000. Rendering diamonds. In *Proceedings of the 11th Western Computer Graphics Symposium (WCGS)*, 9–15.
- SUN, Y., FRACCHIA, F. D., AND DREW, M. S. 2000. Rendering light dispersion with a composite spectral model. In *International Conference on Color in Graphics and Image Processing - CGIP'2000*.
- TANNENBAUM, D. C., TANNENBAUM, P., AND WOZNY, M. J. 1994. Polarization and birefringency considerations in rendering. In *Comp. Graphics (SIGGRAPH '94 Proceedings)*, ACM Press, 221–222 (Extended version available on CD-ROM).
- THOMAS, S. 1986. Dispersive refraction in ray tracing. *Visual Computer* 2, 3–8.
- WILKIE, A., TOBLER, R. F., AND PURGATHOFER, W. 2000. Raytracing of dispersion effects in transparent materials. In *WSCG Conference Proceedings*. <http://citeseer.nj.nec.com/wilkie00raytracing.html>.
- WILKIE, A., TOBLER, R. F., AND PURGATHOFER, W. 2001. Combined rendering of polarization and fluorescence effects. Tech. Rep. 186-2-01-11. Available at www.cg.tuwien.ac.at/research/TR/01/.
- WOLF, E. 1959. Coherence properties of partially polarized electromagnetic radiation. *Il Nuovo Cimento* 8, 6 (september), 1165–1181.
- WOLFF, L., AND KURLANDER, D. J. 1990. Ray tracing with polarization parameters. *IEEE Computer Graphics and Applications* 10, 6 (november/december), 44–55.
- WYSZECKI, G., AND STILES, W. 1982. *Color science: Concepts and Methods. Quantitative Data and Formulas*. Wiley.
- YOKOI, S., KURASHIGE, K., AND ICHIRO TORIWAKI, J. 1986. Rendering gems with asterism and chatoyancy. *The Visual Computer* 2, 5 (Sept.), 307–312.
- YUAN, Y., KUNII, T. L., INAMATO, N., AND SUN, L. 1988. Gemstone fire: Adaptive dispersive ray tracing of polyhedrons. *The Visual Computer* 4, 5 (Nov.), 259–270.