State of the Art in Perceptually Driven Radiosity

Jan Přikryl and Werner Purgathofer

Institute of Computer Graphics, Vienna University of Technology
Karlsplatz 13/186/2, A-1040 Wien, Austria
e-mail: {prikryl,wp}@cg.tuwien.ac.at

Abstract

Despite its popularity among researchers the radiosity method still suffers of some disadvantages over other global illumination methods. Besides the fact that the original method allows only for solving the global illumination of environments consisting of purely diffuse surfaces, the method is rather computationally demanding. In the search for possible speed-up techniques one of the possibilities is to take also the characteristic features of the human visual system. Being aware of how the human visual perception works, one may compute the radiosity solution to lower accuracy in terms of physically based error metrics, but being sure that the physically correct solution won't bring any improvements in the image for the human observer.

In the following report we briefly summarize achievements in the radiosity research in the past years and present the state of the art in perceptual approaches used in computer graphics nowadays. We will give an overview of known tone-mapping and perceptually-based image comparison techniques that can be used in the scope of the radiosity method to further speed up the computational process. In the second part of the report we concentrate on known radiosity methods that already use these perceptual approaches to predict different visible errors of the result of the radiosity computation. We will not speak about importance-driven radiosity solutions, as those methods are based on using geometric visibility rather than on using human perception-aware techniques.

1. Introduction

Radiosity is a global illumination method that is able to compute physically correct lighting for closed scenes consisting of usually diffuse surfaces. As the first physically based method for computing global illumination, it has been an all-time favorite of all computer graphics events over the past decade. Despite its ability to produce high-quality images of virtual environments, the method puts heavy demands on computational power and memory space needed for the computation.

The radiosity method uses criteria based on radiometric values to drive the computation — to decide about sufficient mesh quality or to estimate the error of the simulation process. This is absolutely correct for the case of radiometric simulation, when the user is interested in actual values of radiometric quantities. On the other hand, the radiosity method is very often just used to generate pictures for the human observer and those pictures do not require to be correct physical simulations, they just have to look the same for the observer.

Therefore it seems beneficiary to take the behaviour of the human visual system into account when producing radiosity images and using this knowledge combined with the knowledge about the display transformation to develop criteria in a perceptually mapped space in order to decide what actually causes a visible error in the radiosity output and what can be safely ignored. Exploiting the fact that current display devices can not by large reproduce the real world range of luminances, or knowing which difference in contrast or colour can be noticed by a human under given viewing conditions, we hope to reasonably decrease the time needed for a radiosity simulation.

In the following report we briefly summarize achievements in the radiosity research in the past years and present the state of the art in perceptual approaches to image synthesis. In the second part we concentrate on radiosity methods that use perceptual approaches to further speed up the radiosity computation.

2. Radiosity methods

The radiosity method (see[8, 16, 39] for an introduction) belongs to the broad family of finite element methods. In general the method works as follows: The light propagation in the scene is described using a radiometric quantity called radiosity. The scene is given as a mesh of planar patches, which makes it possible to approximate the original integral equation describing the propagation of radiosity in the space by a set of linear equations. The influence of radiosity some patch leaves on others is described by a set of form factors. The radiosity function representing the light distribution in the scene is then computed by solving the system of linear equations.

First methods used Gauss-Seidel or Southwell relaxation to solve the equation system. Early approaches to compute form factors used Z-buffer-like approaches and suffered form aliasing problems. This section gives a short overview of radiosity methods that are in use nowadays.

Using a simple mesh structure, one often computes radiosity transfers that have a very small influence on the final error of the radiosity system solution. Moreover, if the scene is large, the number of form factors that have to be computed grows rapidly, slowing down the computation to the very edge of usability. A first step towards improving this situation was a two level hierarchy proposed together with adaptive mesh refinement techniques. From this algorithm, more sophisticated algorithms of rapid hierarchical radiosity have been developed in the early 90s by Hanrahan et al.[18]. Followups to this work include adopting the concepts to wavelet radiosity, parallelization of the method and using clustering strategies to further speed up the computation (Gortler et al.[17], Smits et al.[40], Bohn and Garman[5], Stuttard et al.[42], Bekaert and Willems[4]).

In the meantime, a faster Monte Carlo approach to solving integral equations has been applied to radiosity (Pattanaik and Mudur[36], Neumann et al.[28, 33]). Further improvements include better sampling techniques (Keller[21], Neumann et al.[32]) and hierarchical methods (Tobler et al.[44], Bekaert et al.[1]).

As the radiosity scene grows in complexity, it is becoming less desirable to compute a radiosity solution for the whole complex scene, since the observer is often interested in the image of his/her immediate environment only. For such cases importance-driven radiosity solutions have been proposed, that use a patch property dual to radiosity to express importance of a patch for the observer. Extensions to this method include the combination of progressive refinement and importance radiosity, hierarchical importance-driven radiosity and the combination of stochastic and importance-driven radiosity (Smits et al.[41], Bekaert and Willems[2, 3], Neumann et al.[27]).

Nowadays, higher order elements and Galerkin approaches to solving radiosity are often used (Zatz[48], Feda[10]). Form factors in hierarchical radiosity are usually computed using stochastic ray-casting (Wallace et al.[46], Keller[20]), the stochastic radiosity approach does not even require form factors to be explicitly computed. A fast hybrid algorithm combining a stochastic solution with hardware-accelerated rendering has also been presented (Keller[22]).

3. Perceptual approaches in image synthesis

In other areas of computer graphics, as image processing and coding or image reproduction, human perception-aware approaches have already been used to drive the comparison of the images, to predict the distortion of compressed images or to transform a high range of image luminances onto a narrow contrast range of display material. As those methods form a good basis for algorithms that can be used in perceptually-driven radiosity approaches, we give an overview of them in this section.

In the area of image processing, perceptual error metrics are used for image comparison and image coding that enable to better predict differences between two images as opposed to the perceptually inappropriate and widely used mean-squared error metrics. All the metrics are based on the fact that if some image feature is physically important it does not imply that this feature will also disturb the human observer. As the metrics are usually based on some approximate model of first stages of the human vision, the comparisons in perceptual space yield principally better results than classical comparisons of radiometric values stored in the image (Daly[9], Teo and Heeger[43], Boker[6], Rushmeier et al.[37], Gaddipatti et al.[13]).

Tone reproduction operators known from image synthesis make it possible to map a bright scale of image luminances onto a narrow scale of CRT luminances in such a way that the perceived CRT image can be thought of as producing the same mental image as the original image (Tumblin and Rushmeier[45], Chiu et al.[7], Schlick[38], Ward [47], Larson et al.[23], Neumann et al.[29, 30]).

Modelling features of the human visual system allow us to simulate the effects of visual masking and light adaptation dependent on current lighting conditions, and use the simulation results to determine the acceptable error that may be introduced by our computation, but stays unnoticed by the observer (Ferwerda, Pattanaik et al.[11, 12, 35]).

3.1. Terminology

The following terms will be used often in the further discussion:

Luminance L is a photometric equivalent of radiance, *i.e.* it describes visual sensation produced by the electromagnetic spectrum that can be physically quantified with given radiance. The luminance unit is called *nit*, but the equivalent unit of cd/m^2 is also frequently used.

Since our visual system does not perform equally under all lighting conditions, the *adaptation level* of the human visual system, usually denoted L_a, makes it possible to describe its the current working conditions. The adaptation level is given by the luminance to which the human visual system is adapted to.

As the adaption luminance changes, the minimum discernible differences in luminance change as well. *Just noticeable difference* tells us which minimal difference in luminances can be detected at a given adaptation level.

3.2. Tone mapping

The physical accuracy in rendering itself does not yet guarantee that the displayed images will have a realistic visual appearance[11]. Nowadays, typical CRTs can only display luminances in the range 1 to 100 cd/m^2 and typical print reproduction materials (film slides, paper) even less, whereas the range of luminances found in the real world can vary between 10^{-5} to 10^5 cd/m^2 [14,45]. Due to human visual system adaptation we can achieve that the subjective levels of brightness and contrast that occur in a real environment can be reproduced on a display device, even if the luminances themselves cannot. This adaptation process is simulated by *tone reproduction operators*.

The aim of tone reproduction operators is to compute intensities for display on a device such as a CRT, so that when those intensities are viewed by an observer, the mental image they form is as close as possible to that of a real world scene. In the next sections we will briefly present the most widely known tone reproduction operators.

3.2.1. Tumblin & Rushmeier's mapping

Tumblin and Rushmeier's model[45] uses results obtained by Stevens and Stevens regarding the brightness associated with a luminance at a particular adaptation level. The tone reproduction is considered for gray-scale images only, since in this environment the spectral radiance is uniform for all wavelengths and the luminance is therefore just a constant times the uniform spectral radiance. Another reason for this limitation is that for each gray input value the display device gives a unique output luminance, which is not true for colour systems where many different colours may have the same luminance.

The whole operator consists of three transformations: First, the real world luminance is transformed to the true real world brightness. Then, using inverse display observer and inverse display device transformations the value of display input that generates the desired real world response is obtained.

The original mapping of a real world luminance value L_w to the corresponding display value n from interval $\langle 0,1 \rangle$ is described as

$$n = \left[\left(\frac{L_w^{\left(\frac{\alpha_w}{\alpha_d}\right)}}{L_{d,max}} \right) \cdot 10^{\left(\frac{\beta_w - \beta_d}{\alpha_d}\right)} - \left(\frac{1}{C_{max}} \right) \right]^{\left(\frac{1}{\gamma_d}\right)} \quad (1)$$

where α_w and β_w are constants expressing the real world observer adaptation, α_d and β_d are constants expressing the display observer adaptation, $L_{d,max}$ is the maximum displayable luminance of the display device, C_{max} is the maximum achievable contrast of the display device and γ_d is the gamma correction value.

For a given adaptation level L_a the adaptation coefficients α_a and β_a can be written as

$$\alpha_a = 0.4 \log_{10}(L_a) + 2.92 \quad (2)$$

$$\beta_a = -0.4 \left[\log_{10}(L_a)\right]^2 - 2.58 \log_{10}(L_a) + 2.02 \quad (3)$$

The displayed and the real world luminances are bound together by the formula of

$$L_d = \zeta(L_w) = L_w^{\frac{\alpha_w}{\alpha_d}} \cdot 10^{\frac{\beta_w - \beta_d}{\alpha_d}}, \quad (4)$$

where $\zeta()$ is the tone mapping function.

Gibson and Hubbold report poor performance of the Tumblin and Rushmeir model in mesotopic and scotopic conditions — they observations show that the images appear much lighter than they should[14].

3.2.2. Non-uniform extension by Chiu et al.

Chiu et al.[7] proposed the mapping function to be spatially non-uniform in order to enlarge the dynamic range of the displayed image. Their mapping exploits a fact that if we preserve the qualities of the image that the visual system is particularly sensitive to, such as local contrast, we can change image characteristics that the visual system is particularly insensitive to, such as overall luminance.

The authors state that a slow spatial variation in luminance is to some extent ignored by the human eye. This implies that we can display pictures with a wider range than the dynamic range of the display and that as long as the scaling function has a low magnitude gradient, we should not see the scaled image as "looking" very different than the unscaled image.

3.2.3. Schlick's mapping

Schlick[38] tested several widely used mapping approaches (gamma-corrected linear mapping, gamma-corrected clamping, logarithmic and exponential mapping) in the search for a simple and fast mapping operator for high dynamic range images. He proposes a rational mapping function as the operator. The advantage of such a function is that for a given input there is only one parameter that influences its behaviour. This parameter corresponds to the number of different intensity values visible on the display and to the overall dynamic range of the image.

Schlick also proposes a non-uniform mapping extension of the technique. In general, he modifies the actual mapping function according to an average difference in the given area. Surprisingly, the best results have been achieved with the area size set to one, *i.e.* taking into account the given pixel only. This produces a uniform rational mapping again, this time with a modified parameter selection function.

The advantage of Schlick's approach is straightforward: since he uses a rational function as the mapping operator his mapping approach is fast. Further on, the picture does not have to be computed in the absolute units, as no knowledge about real world adaptation is required. However, the parameter selection mechanism fails in some cases (*e.g.* small bright area in a low contrast image) producing parameterizations that do not reproduce the image correctly.

3.2.4. Ward's mapping

Another simple model was proposed by Ward[47] who tried to find a mapping function that would exhibit similar behaviour as the mapping of Tumblin and Rushmeier[45], but that would be less computationaly demanding. His model attempts to match the just noticeable differences on the display device and in the real world and to produce a linear mapping function. He tries to find a constant of proportionality between display luminance and world luminance that would yield a display with roughly the same contrast visibility as the actual scene. This means that Ward's tone mapping operator tries to preserve the perceived contrast, as opposed to the Tumblin-Rushmeier's tone mapping operator that preserves perceived brightness of an image.

Ward's operator consists of a linear formula

$$\Delta L_{\mathrm{d}} = m L_{\mathrm{w}}, \tag{5}$$

where L_{d} is the luminance to be displayed and L_{w} is the real world luminance of the image pixel. As the goal of the operator is to preserve the perceived contrast, the scale factor m binds together the minimum discernible luminance changes, *i.e.* the just noticeable luminance differences, ΔL, at the display and world adaptation levels L_{da} and L_{wa}:

$$\Delta L(L_{\mathrm{da}}) = m \Delta L(L_{\mathrm{wa}}), \tag{6}$$

When the parameter m is valid for the given viewing conditions, the luminance mapping according to Equation (6) will map the differences that are just visible in the real world to the differences that are just visible on our display. In order to determine a good approximation of m for the given picture and usual viewing conditions, we have to supply two luminance values: the world adaptation luminance, L_{wa}, and the maximum display luminance, $L_{\mathrm{d,max}}$. The display adaptation luminance L_{da} is estimated from the maximum display luminance as

$$L_{\mathrm{da}} = \frac{1}{2} L_{\mathrm{d,max}} \tag{7}$$

and the scale factor m can be computed as

$$m = \frac{1}{L_{\mathrm{d,max}}} \left[\frac{1.219 + \left(\frac{L_{\mathrm{d,max}}}{2} \right)^{0.4}}{1.219 + L_{\mathrm{wa}}^{0.4}} \right]^{2.5}. \tag{8}$$

3.2.5. Proposals of Ferwerda, Pattanaik et al.

Ferwerda et al.[11] proposed another tone reproduction technique that is also based on the concept of matching just noticeable differences for a variety of adaptation levels. Their model captures the changes in threshold visibility, colour appearance, visual acuity (sometimes called contrast sensitivity) and time-course of adaptation of the human visual system.

As their work is based on experimental threshold detection data, they based their model on Ward's tone mapping operator. Their extension consists of taking into account not only the cone characteristics (as Ward did), but also the rod threshold data. Moreover, their model also counts for the whole range of viewing conditions from scotopic over mesotopic to photopic level. Using the adaptation dependent visual acuity function they are able to remove those image details that cannot be noticed by the human observer at a given adaptation level. Also, using the data measured for the time-course of the light and dark adaptation they can reproduce images so that they correspond to what the human observer would have seen when entering a dark or bright environment.

In the followup to this work[12], the authors study the effects of texture masking on masking out different artifacts in the resulting image. Their model consists of four stages similar to that of common models of early phases of human vision process: They first transform the spectral radiances into responses in a colour space, producing a colour representation of the image. In the next stage this colour representation is decomposed into pattern representations that account for different spatial frequency and orientation triggered detection mechanisms of the human visual system. Then an appropriate masking function is applied to each of these detection mechanisms to account for the effect of visual masking. In the last stage the results of masking are compared by the detection method to determine whether the input artifacts would be visible or whether the masking behaviour of the human visual system will mask them out and the human observer will not notice them at all.

Their most recent work[35] presents a computational model of adaptation and spatial vision for realistic tone reproduction. The model allows the operator to address the two major problems in realistic tone reproduction: wide absolute range and high dynamic range scenes can be displayed, and the displayed images match our perceptions of the scenes at both threshold and suprathreshold levels to the degree possible given a particular display device.

3.2.6. Visibility preserving mapping of Larson et al.

Larson et al.[23] present a technique that combines advantages of the previous ideas allowing us for correct display of scenes with wide dynamic ranges for different adaptation levels. The mapping operator is uniform and tries to preserve the visibility of objects in the image.

The task is achieved by using more then a single view adaptation level, which is in fact coherent with the mechanism of human visual system adaptation. The authors noticed that image luminance levels tend to be clusterized rather than to be uniformly distributed throughout the whole dynamic range of the picture. They state that as long as the bright areas are displayed with higher luminances than the dim areas, the absolute difference in the luminance is not so important.

Larson et al. first compute a luminance histogram and cumulative distribution function of all local adaptation luminances found in the image and discover clusters of similar adaptation levels. The adaptation luminances are computed for commonly accepted 1° field of view.

The histogram is then adjusted in order to minimize the visible contrast distortions. The clusters are mapped to the display values preserving local contrast of the cluster. The method uses also knowledge about veiling luminance (*i.e.* glare effects), colour sensitivity and visual acuity to count for imperfections of the human vision.

3.2.7. Miniumum loss methods

Neumann, et al.[30] present an interesting and simple extension to the linear mapping principle. In their mapping they look for such a colour interval in the image histogram that can be linearly mapped with the given clipping contrast in a way that the minimum information is lost from the input data due to clipping. They present two slightly different approaches to this method: In the first approach they try to minimize the number of colour histogram bins that are clipped, in the second one they try to minimize the number of image pixels affected by clipping.

3.3. Perceptual difference metrics

Nowadays, the mean-squared and root mean-squared error (MSE,RMSE) of the colour values is still widely used as a convergence metric for radiosity computation. The fact that this error metric may produce totally misleading results when judging the similarity of two images perceived by the human observer has been widely discussed in the image evaluation literature in the past years[13, 15, 24, 37]. The only work from the field of global illumination that does not ignore this fact is the recent paper by Myszkowski[26].

3.3.1. Contrast sensitivity function

It is well known that the human visual system is not able to resolve arbitrary details in the perceived image. The visibility limit for low contrast patterns can be predicted using a measure called *contrast sensitivity*[9, 37], which is defined as inverse of the contrast threshold for a given spatial frequency. This measure tells us how faded or washed out can image details be before they become indistinguishable from a uniform field. It is a function of the size of image features, or the spatial frequency of visual stimuli produced by these features. The *contrast sensitivity function* (CSF) plots the contrast sensitivity for all spatial frequencies and given viewing conditions. Two popular CSFs are plotted in Fig. 1.

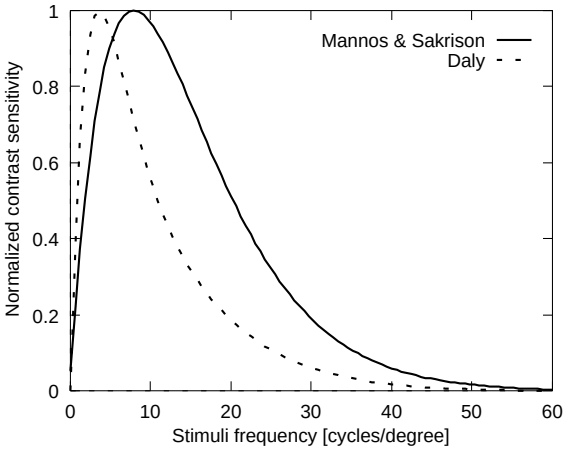

Figure 1: *Normalized Mannos and Sakrison's contrast sensitivity function*[24] *and normalized Daly's contrast sensitivity function*[9] *computed for adaptation level 50* cd/m^2.

The contrast sensitivity values were determined as follows: At the given contrast level, a sine-waved stripe pattern with the given spatial frequency were presented to the human observer. When the stripes were very thin, *i.e.* the spatial frequency of the stimuli was very high (above 60 cycles per visual degree), the test subject was not able to distinguish between particular stripes. As the stripe frequency dropped, the threshold contrast, above which the stripes were distinguishable, dropped too. However, after achieving a peak value of approximately 4–8 cycles per visual degree, depending on the adaptation level, the threshold contrast starts to grow again. An example figure of *Campbell-Robson Contrast Sensitivity Chart*[34] illustrating this phenomenon is depicted in Fig. 2.

Mannos and Sakrison[24] studied the effects of contrast sensitivity on image coding. In their paper a proposal to the functional fit to the experimental data is presented. Their CSF has the form of

$$\tilde{A}(f_r) = (0.05 + 0.2964 f_r) \cdot \exp[-(0.114 f_r)^{1.1}] \quad (9)$$

and predicts the perceptual sensitivity to a visual stimuli of spatial frequency f_r. This CSF form has been proven to be applicable for image comparison[37, 13]. Also rather popular

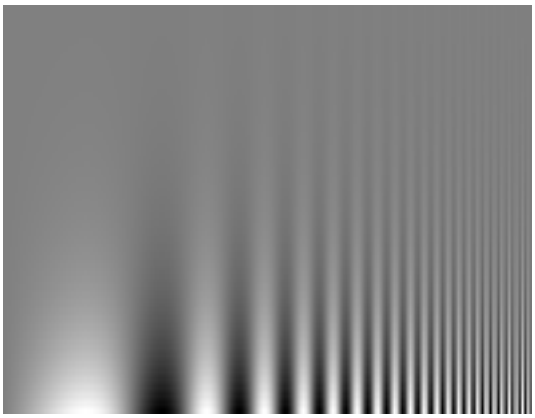

Figure 2: *Campbell-Robson contrast sensitivity chart.*

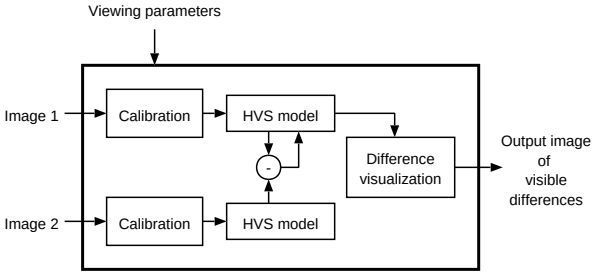

Figure 3: *The block structure of the visible differences predictor.*

is the adaptation level dependent CSF proposed by Daly[9], which forms a part of his perceptually-driven image comparison method.

Another measure of ability to perceive spatial details is known as *visual acuity*[11,23]. The visual acuity tell us the maximum resolvable spatial frequency at the given adaptation level.

3.3.2. Metric proposals

To develop a image comparison metric that would provide results corresponding to the results obtained from human observers is a very difficult task. There are however approaches that are already able to provide us with to some extent meaningful results. This section will briefly discuss perceptually driven comparison metrics that are used in computer graphics nowadays.

Visible Differences Predictor. The *visible differences predictor* (VDP) introduced by Daly[9] is an algorithm for describing the human visual response. Its goal is to determine the degree to which the physical differences between two images become visible to the human observer. The block structure of the algorithm is depicted in Fig. 3.

The VDP is a relative metric — it does not describe the absolute image quality, but rather describes the visibility of differences between two input images. The algorithm consists of three major components: a calibration component, used to transform the input to values understood by the second component which is a model of the human visual system (HVS). The difference of HVS responses is then visualized by the difference visualization component. Output of the VDP is an image map containing the probabilities of HVS detecting the differences between the input images for every pixel.

The calibration block obtains the image data and calibrates them so that they have any meaning to the subsequent HVS model block. The calibration process takes a number of input parameters, describing the conditions for which the VDP will be computed. The parameters include the viewing distance of the human observer, the pixel spacing and necessary values for the display mapping.

The HVS model concentrates on the lower-order visual system processing, *i.e.* on the visual cortex. The model addresses three main sensitivity variations of the human visual system: the dependence of sensitivity on the illumination level, on the spatial frequency of visual stimuli and on the signal content itself.

The variations in sensitivity as a function of light level are primarily due to light-adaptive properties of the retina — they are often referred to as the *amplitude nonlinearity* of the human visual system. The variations as a function of spatial frequency are expressed by the contrast sensitivity function (see Section 3.3.1). The dependency of sensitivity on the signal contents is due to the postreceptoral neural circuitry and is usually described as masking (see also [12]).

The difference visualization block allows for two different visualization techniques, displaying the difference probabilities either as a gray scale value on a uniform field of gray or displaying the probabilities in colour in the context of the reference image — in this case it is easier to judge the correspondence between the predicted differences and the differences actually observed between the two input images.

Teo and Heeger. Teo and Heeger[43] present a perceptual distortion measure based on so the called *normalization model* of early phases of human vision. Their model fits the empirical measurements of the neuron response properties in the primary visual cortex and the psychophysics of the spatial pattern detection.

They use a set of linear sensors that are tuned to different spatial orientation and frequencies as a model for early stages of the human vision. The result of the combined sensor response is squared and transformed to normalized contrast values — this yields a response that is close to the response of neurons in the visual cortex. The simulated neuron responses to the both images are then compared using the usual mean square error mechanism.

Comparison experiments of Rushmeier et al. Rushmeier et al. stated basic rules for perceptual metric behaviour[37]:

1. If we are comparing the image with itself, die predicted difference has to be zero, $M(A,A) = 0$.
2. The difference between two images must not depend on their order in the comparison, $M(A,B) = M(B,A)$.
3. $M(A,C)/M(A,B) >> 1$ for A and B appearing similar and A and C appearing different
4. $M(A,C)/M(A,B) \approx 1$ for A, B and C appearing similar to one another
5. $M(A,B)/M(C,D) \approx 1$ for the difference between A and B appearing similar to the difference between C and D

Two of the three metrics proposed by [37] gave promising results. Both of the metrics transform the image luminances with fast Fourier transform (FFT) into the frequency space and use CSF in image luminances to account for the eye sensitivity on luminance variations.

The first promising model has been derived form the visual fidelity criterion of Mannos and Sakrison[24]. The image luminances are first normalized by the image mean luminance. A cubed root is then applied to the normalized values in order to account for nonlinearity of the human perception. In the next step, a FFT of the resulting values is computed transforming the image data into the spatial frequency space. The FFT result is filtered with the contrast sensitivity function $\tilde{A}(f_r)$ (9). The pixel-based MSE difference between the resulting filtered values is then used to obtain the metric value.

The second model is inspired by the first part of Daly's VDP[9] and is similar to the previous one. However, the image luminances are not normalized — instead of that, a pixel-based transformation of luminances that counts for both adaptation and nonlinearity of human perception is used. Then, as in the previous model, a FFT of the transformed values is computed and the result is filtered using the adaptation level dependent CSF computed for adaptation level 50 cd/m^2, that has been proposed by Daly in his paper. The results for both images are then compared using the MSE.

Wavelet comparison metric. Gaddipatti et al.[13] use a wavelet-based metric for the decision whether the rendering of two subsequent images has brought some perceptually noticeable difference or not. Since the comparison involves the wavelet transformation, image comparison operates on different "levels of detail" in each step. Using this comparison we are able to compare the rough features of both images first and to refine the comparison on details of the wavelet pyramid in further steps. This fact allows us to avoid the undesired features of the mean-squared error approaches, especially the sensitivity to the high frequency blur and brightness shifts.

Boker's experiments. Boker[6] states that the human visual system performs near to an optimum value for an ideal trichromatic system composed of three linear components. His experiments with an unconstrained four factor maximum model seems to fit significantly better than a three factor unconstrained model, suggesting that a colour metric is better represented in four dimensions than in a three dimensional space.

He mentions that the visual system would attempt to preserve, as much as possible, the covariances present in the distribution of photon energies generated by the product of the illuminant spectra with reflectance spectra from objects present in the environment. The visual system adapts to differing brightness and overall spectral content of illumination sources such that a perception of *colour constancy* is maintained within a wide range of environmental lighting conditions. If colour constancy is to be achieved, the interaction between the mean and variance of each wavelength must be removed in order to preserve an invariant pattern of covariances between wavelengths reflected from objects in the environment. It is this pattern of covariances that would define the perceived colour of an object.

Stochastic metric by Neumann et al. Neumann et al.[31] present a stochastic approach to image comparison. Their basic idea is to place a limited number of random-sized rectangle pairs in both compared images and to determine the average colour in each rectangle. For every rectangle pair the average colour difference is computed and weighted according to the rectangle size and the contrast sensitivity function. The result of the image comparison is then obtained by combining the CSF weighted colour differences for all tested rectangles. The use of quasirandom numbers makes the final method deterministic.

4. Perceptually-driven radiosity

By ensuring that only those features of the radiosity solution that are perceptually important are computed to a great accuracy, we hope to be able to produce visually acceptable images faster and more efficiently.

Because the mechanism by which we perceive the radiosity values is nonlinear, any metric used inside the radiosity system, that is based solely on radiometric comparisons cannot guarantee that additional errors will not be introduced during the display process. Differences in radiosity values that pass the refinement test (and therefore cause subdivision) may in fact be undetectable after the display transformation has been performed.

In case that we are more concentrated on the resulting images then on the actual radiometric values obtained during the global illumination simulation, one of the ways of improving the existing radiosity methods is to *compute only the necessary things*, that means to concentrate our computational effort on places where the errors are visible and disturbing. To achieve the maximal effect nowadays solutions use

- perceptual approaches to hierarchy subdivision that include perceptual error metrics as a part of the subdivision oracle,
- perceptual convergence metrics instead of radiometric metrics based on the mean squared error.

4.1. Existing solutions

If we exclude importance-driven radiosity methods that are sometimes classified as perceptually based, applications of human perception-aware techniques to the radiosity systems are still rather rare. Just recently there appeared several papers that use human perception-aware techniques to guide the computations. Techniques described in them can be subdivided into two groups.

The first group of perceptually driven radiosity approaches works directly in object space. Both algorithms belonging to this group estimate the perceptual transformation from luminance to display intensity before the radiosity simulation starts. This makes it possible to use perceptually driven oracle functions in hierarchical radiosity or to cull perceptually unimportant discontinuity lines (Gibson and Hubbold[14], Hedley et al.[19]).

The second group of algorithms is image-space based. The two algorithms presented in this group compare some features present in images resulting from the radiosity simulation. This information is then used to drive the hierarchical refinement of patches or to estimate the convergence of the method (Martin et al.[25], Myszkowski[26]).

4.1.1. Perceptually-driven hierarchical radiosity

Gibson and Hubbold[14] present an improvement to the hierarchical radiosity method that uses perceptually based measures to control the generation of view independent radiosity solutions.

In their paper they describe a new oracle that stops patch refinement once the difference between successive levels of elements becomes perceptually unnoticeable. The authors also show how the perceived importance of any potential shadow falling across a receiving element can be determined. This importance is then used to control the number of rays that are cast during visibility computations and allows to significantly reduce the total number of rays required for a form factor solution without significant loss in visible image quality. They also discuss how a method similar to the refinement oracle can be used to join elements of the radiosity mesh when the differences between elements at successive levels become unnoticeable.

The improvements are achieved using an a priori estimate of the real world adaptation luminance, making it possible to incorporate a tone mapping algorithm into the simulation process. Then, using Tumblin-Rushmeier's tone reproduction operator[45] in perceptually uniform $L^*u^*v^*$ colour space to transform luminance values to display colour values, those aspects of the displayed solution that are visually important can be determined. As the above mentioned model does not perform well under scotopic and mesotopic conditions — which are actually the viewing conditions in usual computer laboratories — more complete tone reproduction operators (such as that of Ferwerda et al.[11] or Larson et al.[23]) should be probably used to represent images at low levels of illumination more accurately.

In the following paragraphs we will discuss the three perceptually driven parts of the Gibson and Hubbold's algorithm in more detail.

Adaptive refinement. The underlying hierarchical radiosity method used by the authors works with CIE XYZ colour values. However, this colour space is not perceptually uniform — colours which have the same numerical distance do not have to be perceptually equidistant. This means that the comparison of two colour differences in the XYZ colour space may suggest both differences are the same, even if one would be perceived as smaller than the other. For performing colour comparisons, it is therefore better to transform the colour values from XYZ to $L^*a^*b^*$ or $L^*u^*v^*$ colour spaces, which are more perceptually uniform (CIE $L^*u^*v^*$ has a nonuniformity ratio of approximately 6:1, but the nonuniformity ration of CIE XYZ space is 80:1).

For an adaptive refinement test, one needs to construct a numerical measure that bounds the perceived difference in intensity between successive levels of element subdivision. There are two different cases that have to be distinguished: a receiving element being fully or partially visible to a source element.

When a receiving element is *fully visible* to a planar convex light source, the luminance function over that element will be unimodal and continuous in all derivatives. Consider a point v on a receiving element e, with luminance L_v and diffuse reflectivity ρ_e, and a source patch s with luminance L_s. The perceived colour at v after the radiosity from s has been transferred will be[14]

$$C_v = \zeta(L_v + \rho_e[F_{vs}L_s + L_a]) \tag{10}$$

where F_{vs} is the unoccluded point to area form factor between v and s, and ζ is the tone mapping function. The inclusion of the ambient correction term L_a is necessary to account for the effect of later iterations — when not used, the first few iterations will be transferring light to completely dark receivers, the perceptual importance of the transfers will be artificially high which would trigger excessive refinement. The ambient luminance L_a is used during element refinement and shadow testing and it is recalculated after every iteration.

The subdivision oracle for a triangle patch and linear basis functions can evaluate Eq. (10) at the corner vertices and edge midpoints to obtain a tone-mapped colour C_v for six

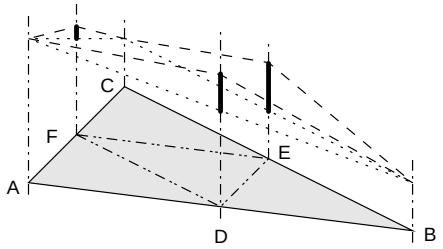

Figure 4: *Determining the perceived difference in luminance between successive levels of element refinement. The dotted line shows the luminance representation for the parent element, the dashed line shows the luminance represented by the successive level of elements. After Gibson and Hubbold*[14].

points v, $v \in \{A,B,\ldots,F\}$ of the triangle (see also Fig. 4). If the linear element at level i was rendered (*i.e.* triangle ABC) instead of four linear elements at level $i+1$ (*i.e.* triangles ADF, BED, CFE, and DEF), the colours C_{iD}, C_{iE}, and C_{iF} at positions D, E and F would be equivalent to linearly interpolating the corner luminances

$$C_{iD} = \frac{C_A + C_B}{2}, \tag{11}$$

$$C_{iE} = \frac{C_B + C_C}{2}, \tag{12}$$

$$C_{iF} = \frac{C_A + C_C}{2}. \tag{13}$$

If C_D, C_{iD}, C_E, C_{iE}, C_F and C_{iF} are represented in CIE $L^*u^*v^*$ colour space, we can calculate the perceived difference Δ^*_{uv} for each colour pair a and b as

$$\Delta^*_{uv} = \sqrt{(L^*_a - L^*_b)^2 + (u^*_a - u^*_b)^2 + (v^*_a - v^*_b)^2}. \tag{14}$$

The largest of the six Δ^*_{uv} values gives a numerical bound on the perceived difference in luminance obtained with and without the extra subdivision. This bound is then tested against a refinement threshold, and the element is subdivided if the threshold is exceeded.

Selecting an appropriate subdivision threshold for a hierarchical radiosity algorithm is a difficult task. The perceptual radiosity method allows for easy selection of the threshold by setting it equal to the just noticeable difference in perceived luminance — *i.e.* setting it to the minimum value of Δ^*_{uv} that will be visible to the user at the given adaptation level. Authors report a threshold of $\Delta^*_{uv} = 2$ suited the computation well. As $\Delta^*_{uv} = 100$ for a difference between reference black and reference white in the $L^*u^*v^*$ colour system, this value corresponds to the commonly used 2% threshold visibility level.

If the source element is only *partially visible* to the receiver, the luminance function over the receiving element can exhibit many changes in continuity. In addition to this, calculating the exact portion of the source that is visible from the receiving element is an expensive operation. Since the method employs a ray cast estimate of visibility (see below), the refinement test for these situations needs to be more conservative than that used when the source is totally visible to the receiver. Once partial visibility is detected the perceptual difference between the vertex m of the receiver with minimum luminance

$$C_m = \zeta(L_m + \rho_e L_\mathrm{a}) \tag{15}$$

and the vertex n with the maximum luminance plus the radiosity transferred from the source patch s

$$C_n = \zeta(L_n + \rho_e[F_{ns}L_s + L_\mathrm{a}]) \tag{16}$$

is computed. The notation is the same as for Eq. (10). This assumes that vertex m is occluded from the source, and so does not receive any radiosity, and vertex n is totally visible. Then, both C_m and C_n are transformed to $L^*u^*v^*$ colour space, Δ^*_{uv} calculated and tested against the subdivision threshold.

Shadow testing. Computing patch-to-patch form factors using Monte Carlo integration[46] involves casting a number of rays between the two patches, intersecting each ray with any potentially blocking objects, and counting the number of rays that are found to connect unoccluded parts of the patches. Visibility detection is often the most computationally intensive phase of the form factor computation. Reducing the number of rays would reduce the computation time, but it can also introduce unpleasant artifacts in the radiosity solution, such as aliasing or noise caused by point sampling errors. If the perceptual impact of light being transferred from the source patch to the receiving one could be determined, also the perceptual impact of any shadow casted over the receiver could be determined. This information can then b used to select an appropriate number of rays to use for the visibility testing. Shadows that are not likely to be noticed could be tested for very cheaply, and more effort focused towards those that are more visible.

Gibson and Hubbold use the colours at a vertex before and after light from a source is considered, C_old and C_new, to estimate the importance of a possible shadow:

$$C_\mathrm{old} = \zeta(L_e + \rho_e L_\mathrm{a}), \tag{17}$$

$$C_\mathrm{new} = \zeta(L_e + \rho_e[F_{vs}V_{vs}L_s + L_\mathrm{a}]). \tag{18}$$

The notation is the same as in Eq. (10). The results have to be converted to $L^*u^*v^*$ colour space in order to determine the perceived difference, Δ^*_{uv} between these quantities. This difference gives a measure of the perceptual importance of any potential shadow boundary caused by the light from the source arriving at the vertex. Also in this case the ambient correction has to be included in the transformation in order to estimate the importance of the shadow taking in account the radiosity received in later iterations.

Having obtained the Δ^*_{uv} value, the next step is to determine the number of rays that shall be cast between the elements in the course of the form factor computation. The method described by Gibson and Hubbold works with lower and upper error thresholds, $\Delta^{\min}_{uv}$ and $\Delta^{\max}_{uv}$, that again depend on just noticeable difference. Provided $\Delta^*_{uv} < \Delta^{\min}_{uv}$, the test for shadows consists of casting a single ray between the receiver and source. For changes above the upper threshold $\Delta^{\max}_{uv}$, a user specified number of rays, N, is used for shadow testing. For changes inbetween, the number of rays linearly dependent on Δ^*_{uv} given by

$$1 + \left\lfloor (N-1) \cdot \frac{\Delta^*_{uv} - \Delta^{\min}_{uv}}{\Delta^{\max}_{uv} - \Delta^{\min}_{uv}} \right\rfloor \tag{19}$$

is used.

This model of shadow importance can also be improved, as it does not take in account that the just noticeable differences vary according to adaptation level and viewer distance[9, 11].

Mesh optimization. The final application of perceptual awareness discussed by Gibson and Hubbold is the process of mesh optimization. They noticed that the resulting mesh from a progressive radiosity solution will be far from optimal, since element refinement triggered by some source may become unnecessary when also the light arriving at the receiving element during later iterations contributed to its radiosity. The authors show how perceptual measures can be used to detect where mesh refinement has become unnecessary, and how to apply this knowledge either as a post-process, or during the course of the radiosity computation in order to conserve memory occupied by the radiosity mesh.

In order to reduce the number of elements in a mesh a posteriori, one can compare the tone-mapped intensity representations of the parent and children elements in the hierarchy in the same way, as it was done during adaptive refinement in cases of full visibility. When the difference Δ^*_{uv} between the two representations is less than the *just noticeable difference*, the leaf elements can be removed without visual impact on the image.

The authors state that they have found it beneficial to perform this optimization after every primary light source patch, except the first one, has distributed its light, and then once again at the end of the simulation.

4.1.2. Discontinuity culling

In order to improve the visual quality of the radiosity simulation, discontinuity meshing is often used to improve the meshing quality in areas where discontinuities in the radiosity distribution occur. However, this method usually produces a too dense mesh of elements that slows down further computations and may pose different numerical problems. Different discontinuity culling techniques are therefore used, that make it possible to discard those discontinuity mesh edges that do not improve the solution very much. Gibson and Hubbold[14] proposed using the knowledge about the perceived importance of light from a source element that arrives at a receiver to decide whether or not the element should be subdivided along the discontinuity lines.

At the same time, Hedley et al.[19] proposed a new approach to culling the discontinuity lines. They use a perception-based metric to determine which discontinuities are important and which can be safely ignored. They noted that even if the discontinuity line itself may not necessarily be visible, it may exert a perceptible influence on the mesh due to improving the general triangulation or stopping artifacts like shadow or light leaks.

In their approach they sample the illumination information along the discontinuity line and also in a small distance at either side of the discontinuity. They record radiance values before and after the current light source has contributed energy at several pairs of points lying aside the discontinuity line. Then, for every sample point they compute differences between tone-mapped colour values in the perceptually uniform $L^*u^*v^*$ colour space. If they find such a sample point where the difference exceeds a given threshold, they assume the discontinuity line makes visible difference to the mesh and therefore it should be included. Other discontinuity lines are then culled and the algorithm proceeds with the next shooting patch.

The perceptual metric for discontinuity culling presented in the paper uses either Tumblin-Rushmeier's operator[45] or the linear mapping introduced by Ward[47]. The world adaptation level is computed either with the approach similar to that of Gibson and Hubbold[14] (see also Section 4.1.1) or — for non-closed environments, where the energy lost during the radiosity computations would result in an overestimate of the world adaptation level — it can be specified manually.

4.1.3. Image-space refinement criterion

Martin et al.[25] present a refinement criterion for the hierarchical radiosity method, which tries to improve image quality taking into account the smoothness of the solution based on pixel intensity values instead of energy ones, and visibility changes along the surfaces for high gradient detection (sharp shadows).

Similar to the observations of Gibson and Hubbold, the authors noticed that most existing oracle functions are based on the computation of radiometric magnitudes in object space — such as form-factors and energy values. These oracle functions do not take the image space features such as the pixel intensities into account. The observations suggest that some refinements in the energy transfer will have a little noticeable effect on the result of the radiosity simulation. Martin et al. therefore propose a refinement criterion that progressively improves image quality working directly on the rendered Gouraud shaded image displaying the result of the radiosity computation.

Martin et al. propose two different oracles for receiver and shooter refinement. We will discuss them in more detail in the following two paragraphs.

Receiver oracle. The oracle for receiver refinement takes into account the smoothness of the Gouraud approximation and the possible visibility artifacts. Beside the shading constraint, the authors also assume that the radiosity system works with triangular patches.

If this is true, the luminance value L at a given pixel X is obtained by linear interpolation of the triangle vertex luminances for the triangle that covers the pixel X. The vertex intensity $I_i^k(\mathbf{T})$ for the k-th basis function and some vertex $\mathbf{T}$ that belongs to patch i is computed by applying a tone-mapping function $\xi()$ to the luminance value $L_i^k(\mathbf{T})$ at the vertex:

$$I_i^k(\mathbf{T}) = \xi(L_i^k(\mathbf{T})). \tag{20}$$

In the paper the tone-mapping operator of Tumblin and Rushmeier[45] has been used.

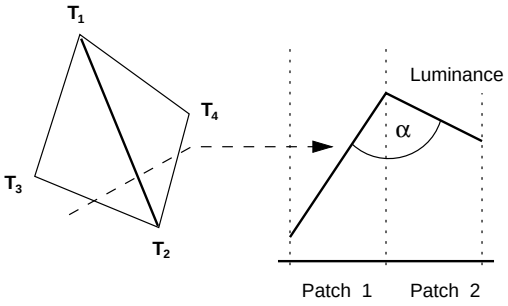

Figure 5: *The slope difference. The angle α measures the discontinuity in Gouraud shading. After Martin et al.*[25]

The subdivision oracle supposes that the transition in the Gouraud shading at the shared edge between two neighbouring patches of a same surface is continuous. This means the change in the intensity along a scan line across two coplanar adjacent patches is expected to be the same at the right side and the left side of the separating edge. As the reader can see in Fig. 5, the angle including the two linear intensity distributions at any point on the separating edge can be used as a measure of the shading discontinuity. If the surface is curved, the angle value depends on the geometrical angle containing both patches as well.

As the proposed subdivision criterion is based on continuity in Gouraud shading, another question appears: Is it possible that continuous Gouraud shading across two adjacent mesh elements appears also in cases when the receiver should be subdivided? Unfortunately the general answer is yes — this can really happen. The possible cases are depicted in Fig. 6.

The first difficulty is that many different patches acting as

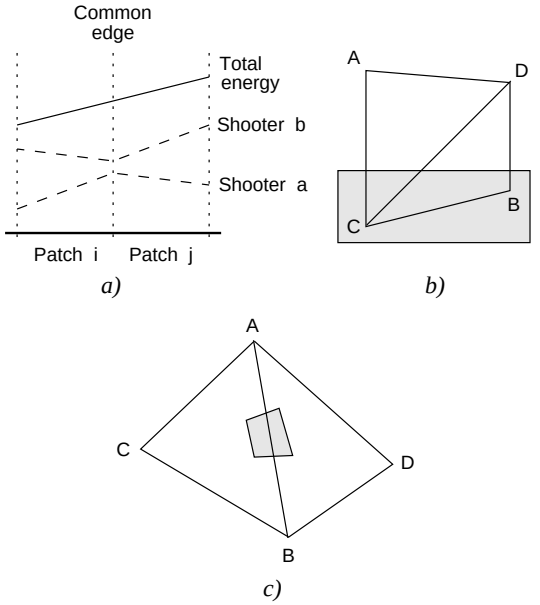

Figure 6: *Different cases of Gouraud shading failing in identification of refinement candidates. After Martin et al.*[25]

shooters contribute to a given radiosity value. These contributions may compensate each other when they are added up. Thus, non linear contributions of each patch separately may give a globally linear radiosity distribution. This is illustrated in Figure 6a, where two shooters *a* and *b* are contributing to receiver patches *i* and *j*. The radiosity function resulting from adding the two shooter contributions seem to be continuous across the edge separating the two receivers. However each isolated shooter contribution shows a high slope change across the shared edge. The contributions are opposite in sign though, and they compensate each other when they are added. Therefore, the oracle has to test the interpolation continuity for each shooter contribution separately. As a consequence, not all shooters interacting with a given patch may cause discontinuity and thus the receiver patch refinement should be realized only for the shooting patches which cause it.

A second aspect to be considered is illustrated in Fig. 6b: A shadow is projected over two adjacent patches in such a way that vertices *A* and *D* are outside the shadow area and vertices *B* and *C* are inside. If the radiosities at *A* and *D* are approximately the same and the radiosities at *B* and *C* also happen to be nearly the same, the Gouraud shading will show a smooth transition between the patches and fail in detecting that they should be subdivided. Fortunately, this case can be easily detected, as it corresponds to a discontinuity in the visibility along the edge. Therefore the oracle has to take into account the difference of visibility of the edge vertices towards the shooting patches leading eventually to a subdivision.

Finally, as with other radiosity methods, the initial mesh used by the method of Martin et al. must have a minimum density in order to avoid precision cases such as illustrated in Fig. 6c. The authors state that the initial mesh must be at least such fine that the maximum patch size is less than the minimum projection size of the smallest object in any scene surface. However, in practice, the initial meshing size will be fixed to some uniform value.

Taking in account the above mentioned cases, the oracle function proposed by Martin et al. tests the Gouraud shading continuity at the edges between adjacent receiving elements and checks the visibility of the edge vertices towards the shooting element. The oracle decides whether is is necessary to subdivide the receiver element and if so, the radiosity algorithm subdivides it and establishes new links. As the radiosity is represented by linear basis functions, new links are created between the shooter and the new sampling points.

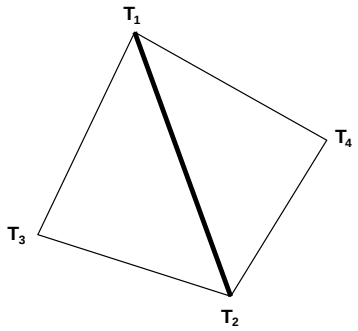

Figure 7: *Vertex geometry for computing the slope estimator. After Martin et al.*[25].

In the paper[25], the following receiver oracle is proposed. The oracle implements ideas outlined above and in addition uses view importance to influence the hierarchical refinement process. The method searches all the edges over each mesh and performs the following three steps for every shooter patch k that contributes to a pair of adjacent patches at the given hierarchy level:

1. For the edge shared by the pair of adjacent patches, the visibility difference ΔV_k between its vertices $\mathbf{T}_1$ and $\mathbf{T}_2$ and with respect to the shooter k is computed as

$$\Delta V_k = \left| V_{k,\mathbf{T}_1} - V_{k,\mathbf{T}_2} \right|, \tag{21}$$

   where $\Delta V_k \in \langle 0,1 \rangle$ and $\Delta V_k = 0$ if there is no change in the visibility between $\mathbf{T}_1$ and $\mathbf{T}_2$ and shooting element k. The value $\Delta V_k = 1$ indicates that one point of the pair $\mathbf{T}_1$, $\mathbf{T}_2$ is completely occluded from the shooter k, while the other is visible from the whole shooting element.
2. The value of V_k is tested against a given visibility threshold ε_{vis} in order to decide if the visibility change is small enough to consider the change of slope as a valid measure of smoothness:
   a. If $\Delta V_k < \varepsilon_{\text{vis}}$, the change in the slope itself can be used as a measure of smoothness. As this oracle is applied many times during the algorithm, an easy but non-conservative estimation of the slope change is used in the paper.
   Since the method considers values in image space, the slope change is computed using pixel intensity values instead of radiosity ones. The mapping to the pixel intensities is nonlinear which means that also the contribution of the other shooters must be taken into account in order to estimate the real change of slope produced by shooter k. The greater the energy of the other shooters the smaller the slope change.
   Using the geometry shown in Fig. 7, the slope change is computed as

$$\Delta I_k = \frac{1}{2} \left| \zeta(L_1 + L_{\min}) + \zeta(L_2 + L_{\min}) - \zeta(L_3 + L_{\min}) + \zeta(L_4 + L_{\min}) \right|, \tag{22}$$

   where $L_{\min}$ is the minimum luminance of the vertex luminances of the two adjacent patches, $L_1 \ldots L_4$ without considering the contribution of the shooting patch k. Note that L_1 and L_2 are the luminances of the separating edge vertices.
   b. If $\Delta V_k \geq \varepsilon_{\text{vis}}$, a more conservative estimation of the slope change is used, based on the difference between the maximum and the minimum pixel intensity values of the four vertices, again taking into account the contribution of other shooters

$$\Delta I_k = \max_{j=1\ldots4} \zeta(L_{\min} - L_j) - \min_{j=1\ldots4} \zeta(L_{\min} - L_j). \tag{23}$$

3. The estimator ΔI_k is tested against the slope threshold $\varepsilon_{\text{slope}}$, weighted by the view importances of the two adjacent patches, Y_{left} and Y_{right}.
   If $\Delta I_k \cdot (Y_{\text{left}} + Y_{\text{right}}) \geq \varepsilon_{\text{slope}}$, both the adjacent patches are subdivided and the links from shooter k are created on the lower level of hierarchy.

Shooter oracle. The above described criterion relies on accurate radiosity values at the vertices. If the vertex radiosities are not computed with enough accuracy is, the oracle will fail.

When a link is initially established between a shooter and a vertex (associated to a basis function at a given level), it is therefore necessary to decide if energy transfer at that level is accurate enough to guarantee a good approximation of the radiosity value at the vertex. If the accuracy of the energy transfer is low, the link should be refined.

This leads to another oracle that decides if a link at a given level is accurate enough. The authors use an energy based oracle that compares the energy carried by the link with a given threshold. If the transferred energy exceeds the threshold, the interaction has to be refined. As the input of the oracle, the maximum of the three products of shooter coefficients by

the transport coefficients, weighted by the importance of the receiver triangle, is used.

According to the paper, the current implementation results show that the above presented criteria significantly reduce the number of links that are needed for a given accuracy while keeping a reasonable memory requirements of the method.

4.1.4. Applications of VDP to radiosity

Myszkowski[26] studied recently the applicability of Daly's VDP[9] (see Section 3.3.2) to the global illumination process. His aim was to develop a suitable image-space criterion to judge about convergence of the global-illumination solution, which would help to predict when to stop the simulation process. He also experiments with the influence of texture masking on a non-uniform adaptive subdivision algorithm.

Perceptual convergence. Using the VDP, quantitative measures of differences between two images are generated. In addition, image regions where such differences will be noticed by the human observer are identified. If the images contain variance in indirect lighting, absolute difference metrics (as MSE) would immediately report those images as being different. The VDP responds more selectively to this phenomenon, taking in account the local contrast, spatial frequency of the variance and visual masking.

Terminating criteria for image synthesis. If the goal of the image synthesis process is the perceptual accuracy of the resulting image and one does not have to care about the actual physical accuracy, the global illumination computation should be stopped in the moment when the resulting image quality becomes indistinguishable from that of the fully converged solution measured by some norm of the radiosity values. Myszkowski states that approaches like mean squared error norm (MSE) do not suite to this task very well, as they predict rather large differences between images tone-mapped with Tumblin-Rushmeier's operator. When discussing the features of MSE, he also objects that MSE takes the global error for the whole scene and in fact the local error may be much higher at some place.

The author again proposes using the VDP in the image space as the termination criterion and tries to determine which images in the sequence of gradually converging images should be compared in order to reliably predict the actual convergence. In his experiments, he finds the VDP comparison between images obtained in time τ and 0.5τ of the simulation to be predict the convergence well.

Adaptive mesh subdivision. Myszkowski notices that the research on perceptually-aware meshing strategies has up to now been almost exclusively limited to using the tone-mapping operator of Tumblin and Rushmeier[14, 19, 25] to predict the influence of subdividing the current mesh element. He therefore proposes a three-step process that uses VDP to predict the masking effect triggered after the patch has been textured. The approach is based on an older method of the author, which uses nonuniform adaptive subdivision and is as follows:

Step 1: Choose candidates according to the traditional tone-mapped approach. Generate edges.

Step 2: Update the lighting recomputing only the lighting situation for the candidates from step 1. Compute VDP between the image from the last iteration step and the current image.

Step 3: Restore edges that do not produce visible differences between both images.

The author reports that this approach works well for areas where illumination changes only gradually, but in proximity of illumination discontinuities the mesh reduction is rather poor.

About authors

Jan Přikryl is a PhD. student at the Institute of Computer Graphics, Vienna University of Technology, Vienna, Austria. His main research topic is perceptual approaches to radiosity solutions.

Werner Purgathofer is Professor of Practical Informatics at the Vienna University of Technology, where he is the head of the Visualization and Animation Group at the Institute of Computer Graphics. His main research topic is global illumination, especially stochastic radiosity methods. His other interests include colour science, visualization and virtual reality.

References

1. P. Bekaert, L. Neumann, A. Neumann, and Y. D. Willems. Hierarchical Monte Carlo radiosity. In G. Drettakis and N. Max, editors, *Rendering Techniques '98 (Proceedings of the Ninth Eurographics Workshop on Rendering)*, pages 259–268. Eurographics, Springer-Verlag/Wien, 1998.

2. P. Bekaert and Y. D. Willems. A Progressive Importance-Driven Rendering Algorithm. In E. Ružický, P. Eliáš, and A. Ferko, editors, *Proceedings of the Tenth Spring School on Computer Graphics '94*, pages 58–67, Comenius University, Bratislava, Slovakia, June 1994.

3. P. Bekaert and Y. D. Willems. Importance-Driven Progressive Refinement Radiosity. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 316–325. Eurographics, Springer-Verlag/Wien, 1995.

4. P. Bekaert and Y. D. Willems. Hirad: A Hierarchical Higher Order Radiosity Implementation. In W. Purgathofer, editor, *Proceedings of the Twelfth Spring School on Computer Graphics '96*, pages 213–227, Comenius University, Bratislava, Slovakia, June 1996.

5. C.-A. Bohn and R. Garmann. A Parallel Approach to Hierarchical Radiosity. In V. Skala, editor, *Proceedings of the Winter School of Computer Graphics and CAD Systems '95*, pages 26–35, University of West Bohemia, Plzeň, Czech Republic, February 1995.

6. S. M. Boker. A Measurement of the Adaptation of Color Vision to the Spectral Environment. *Psychological Science*, 8(2):130–144, 1997.

7. K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman. Spatially Nonuniform Scaling Functions for High Contrast Images. In *Proceedings of Graphics Interface '93*, pages 245–253, May 1993.

8. M. F. Cohen and J. R. Wallace. *Radiosity and Realistic Image Synthesis.* Academic Press Professional, Boston, MA, 1993. ISBN 0-12-178270-0.

9. S. Daly. *The Visible Differences Predictor: An Algorithm for the Assessment of Image Fidelity*, chapter 14, pages 179–206. Digital Images and Human Vision. The MIT Press, 1993.

10. M. Feda. A Monte Carlo Approach for Galerkin Radiosity. *The Visual Computer*, 12(8):390–405, 1996.

11. J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A Model of Visual Adaptation for Realistic Image Synthesis. In *Computer Graphics Proceedings*, Annual Conference Series, pages 249–258. ACM SIGGRAPH, 1996.

12. J. A. Ferwerda, S. N. Pattanaik, P. Shirley, and D. P. Greenberg. A Model of Visual Masking for Computer Graphics. In *Computer Graphics Proceedings*, Annual Conference Series, pages 143–152. ACM SIGGRAPH, 1997.

13. A. Gaddipatti, R. Machiraju, and R. Yagel. Steering Image Generation with Wavelet Based Perceptual Metric. In D. W. Fellner and L. Szirmay-Kalos, editors, *Eurographics '97 Conference Proceedings*, volume 16 of *Computer Graphics Forum*, pages 241–251. Eurographics, September 1997.

14. S. Gibson and R. J. Hubbold. Perceptually Driven Radiosity. *Computer Graphics Forum*, 16(2):129–140, 1997.

15. B. Girod. *What's Wrong with Mean-squared Error?*, chapter 15, pages 207–220. Digital Images and Human Vision. The MIT Press, 1993.

16. A. S. Glassner. *Principles of Digital Image Synthesis.* Computer Graphics and Geometric Modeling. Morgan Kaufmann, San Francisco, CA, 1995. ISBN 1-55860-276-3.

17. S. J. Gortler, P. Schröder, M. F. Cohen, and P. Hanrahan. Wavelet Radiosity. In *Computer Graphics Proceedings*, Annual Conference Series, pages 221–230. ACM SIGGRAPH, 1993.

18. P. Hanrahan, D. Salzman, and L. Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics Proceedings*, Annual Conference Series, pages 197–206. ACM SIGGRAPH, 1991.

19. D. Hedley, A. Worall, and D. Paddon. Selective Culling of Discontinuity Lines. In J. Dorsey and P. Slusallek, editors, *Rendering Techniques '97 (Proceedings of the Eigtht Eurographics Workshop on Rendering)*, pages 69–80. Eurographics, Springer-Verlag/Wien, 1997.

20. A. Keller. The Fast Calculation of Form Factors using Low Discrepancy Sequences. In W. Purgathofer, editor, *Proceedings of the Twelfth Spring School on Computer Graphics '96*, pages 195–204, Comenius University, Bratislava, Slovakia, June 1996.

21. A. Keller. Quasi-Monte Carlo Radiosity. In X. Pueyo and P. Schröder, editors, *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 101–110. Eurographics, Springer-Verlag/Wien, 1996.

22. A. Keller. Instant Radiosity. In *Computer Graphics Proceedings*, Annual Conference Series, pages 49–56. ACM SIGGRAPH, 1997.

23. G. W. Larson, H. Rushmeier, and C. Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes, October 1997.

24. J. L. Mannos and D. J. Sakrison. The Effect of a Visual Fidelity Criterion on the Encoding of Images. *IEEE Transactions of Information Theory*, 20(4):525–536, July 1974.

25. I. Martin, X. Pueyo, and D. Tost. An Image-Space Refinement Criterion for Linear Hierarchical Radiosity. In *Graphics Interface '97*, pages 26–36, May 1997.

26. K. Myszkowski. The Visible Differences Predictor: application to global illumination problems. In G. Drettakis and N. Max, editors, *Rendering Techniques '98 (Proceedings of the Ninth Eurographics Workshop on Rendering)*, pages 223–236. Eurographics, Springer-Verlag/Wien, 1998.

27. A. Neumann, L. Neumann, P. Bekaert, Y. Willems, and W. Purgathofer. Importance-Driven Stochastic Ray Radiosity. In X. Pueyo and P. Schröder, editors, *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 111–122. Eurographics, Springer-Verlag/Wien, 1996.

28. L. Neumann, M. Feda, M. Kopp, and W. Purgathofer. A New Stochastic Radiosity Method for Highly Complex Scenes. In G. Sakas, P. Shirley, and S. Müller, editors, *Photorealistic Rendering Techniques (Proceedings of the Fifth Eurographics Workshop on Rendering)*, pages 195–206. Eurographics, Springer-Verlag Berlin Heidelberg New York, 1995.

29. L. Neumann, K. Matković, A. Neumann, and W. Purgathofer. Incident Light Metering in Computer Graphics. *Computer Graphics Forum*, 1998. To appear.

30. L. Neumann, K. Matković, and W. Purgathofer. Automatic Exposure in Computer Graphics Based on the Minimum Information Loss Principle. In *Proceedings of CGI 98*, pages 666–677, Hannover, Germany, June 1998.

31. L. Neumann, K. Matković, and W. Purgathofer. Perception Based Color Image Difference. In *Eurographics '97 Conference Proceedings*, volume 17 of *Computer Graphics Forum*. Eurographics, September 1998. To appear.

32. L. Neumann, A. Neumann, and P. Bekaert. Radiosity with Well Distributed Ray Sets. In D. W. Fellner and L. Szirmay-Kalos, editors, *Eurographics '97 Conference Proceedings*, volume 16 of *Computer Graphics Forum*, pages C–261–C–270. Eurographics, September 1997.

33. L. Neumann, W. Purgathofer, R. F. Tobler, A. Neumann, P. Eliáš, M. Feda, and X. Pueyo. The Stochastic Ray Method for Radiosity. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 206–218. Eurographics, Springer-Verlag/Wien, 1995.

34. I. Ohzawa. Campbell-Robson Contrast Sensitivity Chart: A New Rendition. Available from `http:://totoro.berkeley.edu/izumi/CSF/A_JG_RobsonCSFchart.html`.

35. S. N. Pattanaik, J. A. Ferwerda, J. A. Fairchild, and D. P. Greenberg. A Multiscale Model of Adaptation and Spatial Vision for Realistic Image Display. In *Computer Graphics Proceedings*, Annual Conference Series. ACM SIGGRAPH, 1998. In print.

36. S. N. Pattanaik and S. P. Mudur. Computation of the Global Illumination by Monte Carlo Simulation of the Particle Model of Light. In A. Chalmers and D. Paddon, editors, *Third Eurographics Workshop on Rendering*, pages 71–83, Bristol, May 1992. Eurographics.

37. H. Rushmeier, G. Ward, C. D. Piatko, P. Sanders, and B. Rust. Comparing Real and Synthetic Images: Some Ideas about metrics. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 82–91. Eurographics, Springer-Verlag/Wien, 1995.

38. C. Schlick. Quantization Techniques for Visualization of High Dynamic Range Pictures. In G. Sakas, P. Shirley, and S. Müller, editors, *Photorealistic Rendering Techniques (Proceedings of the Fifth Eurographics Workshop on Rendering)*, pages 7–20. Eurographics, Springer-Verlag Berlin Heidelberg New York, 1995.

39. F. Sillion and C. Puech. *Radiosity and Global Illumination*. Morgan Kaufmann, San Francisco, CA, 1994. ISBN 1-55860-277-1.

40. B. Smits, J. Arvo, and D. Greenberg. A Clustering Algorithm for Radiosity in Complex Environments. In *Computer Graphics Proceedings*, Annual Conference Series, pages 435–442. ACM SIGGRAPH, 1994.

41. B. E. Smits, J. R. Arvo, and D. H. Salesin. An Importance-Driven Radiosity Algorithm. In *Computer Graphics Proceedings*, Annual Conference Series, pages 273–282. ACM SIGGRAPH, 1992.

42. D. Stuttard, A. Worrall, D. Paddon, and C. Willis. A Parallel Radiosity System for Large Data Sets. In N. M. Thalmann and V. Skala, editors, *Proceedings of the Winter School of Computer Graphics and CAD Systems '95*, pages 421–430, University of West Bohemia, Plzeň, Czech Republic, February 1996.

43. P. C. Teo and D. J. Heeger. Perceptual Image Distortion. In *Human Vision, Visual Processing and Digital Display V, IS&T/SPIE's Symposium on Electronic Imaging: Science & Technology*, 1994.

44. R. F. Tobler, A. Wilkie, M. Feda, and W. Purgathofer. A Hierarchical Subdivision Algorithm for Stochastic Radiosity Methods. In J. Dorsey and P. Slusallek, editors, *Rendering Techniques '97 (Proceedings of the Eigtht Eurographics Workshop on Rendering)*, pages 193–204. Eurographics, Springer-Verlag/Wien, 1997.

45. J. Tumblin and H. E. Rushmeier. Tone Reproduction for Realistic Images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

46. J. R. Wallace, K. A. Elmquist, and E. A. Haines. A Ray Tracing Algorithm for Progressive Radiosity. In *Computer Graphics Proceedings*, Annual Conference Series, pages 315–324. ACM SIGGRAPH, 1989.

47. G. Ward. *A Contrast-Based Scalefactor for Luminance Display*, volume IV of *Graphics Gems*, chapter VII.2, pages 415–421. Academic Press, 1994.

48. H. R. Zatz. Galerkin Radiosity: A Higher Order Solution Method for Global Illumination. In *Computer Graphics Proceedings*, Annual Conference Series, pages 213–220. ACM SIGGRAPH, 1993.

Appendix A: Determining the real world adaptation level

Most tone mapping operators require the knowledge about the value of real world adaptation luminance $L_{a,w}$ in order to work properly. The real world adaptation level for an image-based approach can be estimated using techniques described by Tumblin and Rushmeier[45], Ward[47] or Larson et al.[23]. However, in object-space based radiosity computations, an estimate of the real world adaptation level is usually needed a priori. This estimate can be obtained from an ambient luminance value of every patch with technique proposed by Gibson and Hubbold[14].

First, the area averaged reflectivity ρ_{ave} and emissivity L_{ave}^{out} are computed:

$$\rho_{ave} = \frac{\sum_{i=0}^{N} A_i \rho_i}{\sum_{i=0}^{N} A_i}, \tag{24}$$

$$L_{ave}^{out} = \frac{\sum_{i=0}^{N} A_i L_i^{out}}{\sum_{i=0}^{N} A_i}. \tag{25}$$

Here, ρ_i is the diffuse reflectivity of patch i with area A_i, and L_i^{out} is the emissivity of the patch. An ambient luminance value can be found that represents the average illuminance on a surface after an infinite number of reflections

$$L_{amb} = \frac{L_{ave}^{out}}{1 - \rho_{ave}} \tag{26}$$

The ambient illuminance L_{amb} allows the a priori estimation of the luminance of a patch i in the fully converged solution.

To obtain an estimate of the final adaptation level, Gibson and Hubbold take a logarithmic area weighted average of estimated surface luminances L_i, adjusted so that brightnesses are kept around the brightness constancy level of 8.4 dB:

$$\log_{10}(L_{wa}) = \frac{\sum_{i=0}^{N} A_i \log_{10} L_i}{\sum_{i=0}^{N} A_i} + 0.84. \tag{27}$$

Acknowledgements

This work has been supported by Austrian "Fonds zur Förderung der wissenschaftlichen Forschung" under project number P11545-MAT.