

Using the Visual Differences Predictor to Improve Performance of Progressive Global Illumination Computation

VLADIMIR VOLEVICH

University of Aizu

KAROL MYSZKOWSKI

Max-Planck-Institute for Computer Science

and

ANDREI KHODULEV and EDWARD A. KOPYLOV

Russian Academy of Sciences

A novel view-independent technique for progressive global illumination computing that uses prediction of visible differences to improve both efficiency and effectiveness of physically-sound lighting solutions has been developed. The technique is a mixture of stochastic (density estimation) and deterministic (adaptive mesh refinement) algorithms used in a sequence and optimized to reduce the differences between the intermediate and final images as perceived by the human observer in the course of lighting computation. The quantitative measurements of visibility were obtained using the model of human vision captured in the visible differences predictor (VDP) developed by Daly [1993]. The VDP responses were used to support the selection of the best component algorithms from a pool of global illumination solutions, and to enhance the selected algorithms for even better progressive refinement of image quality. The VDP was also used to determine the optimal sequential order of component-algorithm execution, and to choose the points at which switchover between algorithms should take place. As the VDP is computationally expensive, it was applied exclusively at the design and tuning stage of the composite technique, and so perceptual considerations are embedded into the resulting solution, though no VDP calculations were performed during lighting simulation.

The proposed global illumination technique is also novel, providing intermediate image solutions of high quality at unprecedented speeds, even for complex scenes. One advantage of the technique is that local estimates of global illumination are readily available at the early stages of computing, making possible the development of a more robust adaptive mesh subdivision, which is guided by local contrast information. Efficient object space filtering, also based on stochastically-derived estimates of the local illumination error, is applied to substantially reduce the visible noise inherent in stochastic solutions.

Authors' addresses: V. Volevich, University of Aizu, Tsuruga, Ikki-machi, Aizu Wakamatsu, 965-8580, Japan; K. Myszkowski, Max-Planck-Institute for Computer Science, Im Stadtwald, Saarbrücken, D-66123, Germany; email: karol@mpi-sb.mpg.de; A. Khodulev, E. A. Kopylov, Keldysh Institute of Applied Mathematics, Russian Academy of Sciences, 4 Miusskaya Sq., Moscow, 125047, Russia.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 0730-0301/00/0400-0122 \$5.00

Categories and Subject Descriptors: I.3.7 [**Computer Graphics**]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing, and texture*; I.4.1 [**Image Preprocessing**]: Digitization and Image Capture—*Sampling, Radiometry, Reflectance*; I.4.3 [**Image Preprocessing**]: Enhancement—*Filtering*; I.6.8 [**Simulation and Modeling**]: Types of Simulation—*Monte Carlo*; I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems—*Intensity, color, photometry, and thresholding*

General Terms: Algorithms, Human Factors

Additional Key Words and Phrases: Adaptive mesh subdivision, density estimation, human perception, Monte Carlo photon tracing, progressive refinement, view-independent solutions

1. INTRODUCTION

Current thought on global illumination solutions is that they require long hours of computing for complex scenes, and because of this are used only at the final stages of the scene design and rendering process. This may result in severe obstruction of the design process due to substantial differences in appearance that can occur between the final image and the intermediate images, since the latter are usually based on simplistic lighting computations (Figure 1). While it is clear that the complete global illumination solution for complex geometry, luminaires, and material reflectances is time consuming, the perceptual distance between the intermediate and final images could be reduced by better perception-motivated use of physically-based partial solutions. For example, human eye sensitivity to absolute luminance levels is rather poor, so the appearance of intermediate images can be practically indistinguishable from the final one, well before energy-based stopping criteria are met [Myszkowski 1998b]. This means that the unit of time spent for lighting computations using a given algorithm may have a significantly different impact on the appearance of the image, depending on the stage of computation. Obviously, the same statement is also true when we compare the performance of different algorithms. Thus, the problem of minimizing perceivable differences between the intermediate images and the final image can be stated as the selection of the optimal algorithm at every stage of computation. It is assumed that all algorithms considered are physically based and converge to the correct solution within some random and/or systematic error tolerances.

In this paper we attempt to address this problem experimentally. We first formulate the perception-based framework for estimating the performance of the standard global illumination algorithms we selected. The performance is measured at various stages of computation in terms of perceivable differences between the intermediate and final images. As an objective and quantitative metric we use the visible differences predictor (VDP) developed by Daly [1993], which is based on advanced models of the human visual system (HVS). On the basis of the results, we designed a novel global illumination technique that is a hybrid of appropriately ordered standard techniques. We applied our perception-based framework




Fig. 1. An illustration of perceptual distance between various lighting-computating conventions: (a) point and (b) spotlight sources using the OpenGL-native lighting and rendering (the spotlights approximate goniometrical diagrams); (c) single and (d) multipass (note mirror reflections) OpenGL rendering for precalculated global illumination [Diefenbach 1996].

again to enhance global illumination solutions (in terms of intermediate image quality) by tuning the switchover points between, and introducing algorithmic improvements into, its component techniques. Note that while the perceptually-based framework was used to design our technique and tune its parameters, the HVS models were not incorporated in the resulting technique. Thus, there is no overhead caused by the complex computation required by these models during lighting simulation.

In this research our goal was to design a practical global illumination technique suitable for applications involving rapid generation of images

from arbitrary views, such as production of high-quality virtual sets for broadcasting and computer games. This requires storing the precalculated lighting function as view-independent illumination maps. The design requirements we imposed on our lighting simulation and rendering solution can be summarized as follows:

- The solution must provide the user with an immediate and perceptually meaningful rendering response, even for complex scenes at any stage of computation (the initial response delay due to “warm-up” computing required by some algorithms is not acceptable: e.g., computing links in the hierarchical radiosity framework [Sillion and Puech 1994]). The solution must predict, on a physically sound basis, the distribution of light for any specified material reflectance functions and goniometrical diagrams that describe light sources.
- The simulation error must be controllable at all stages of computing.
- The solution must be view-independent to make possible walkthrough inspection of any portion of a scene upon user demand.

In the following section we discuss previous work on global illumination techniques from the perspective of these requirements, and also some attempts to incorporate perception models into a realistic rendering framework. This is followed by a detailed description of our novel technique. We also show some adaptive mesh subdivision enhancements for the mesh used to store illumination maps, which became possible within the framework of our technique. Finally, we present some implementation details and report on some experimental results. In the Appendix we describe a technique for filtering illumination maps that significantly enhances the image quality at early stages of progressive computation, and is based on stochastically-derived estimates of local illumination and its variance.

2. PREVIOUS WORK

In this section we discuss existing progressive global illumination techniques that are suitable for generating view-independent illumination maps. Since one of the main motivations for this work is to optimize the speed and quality of the visual responses given by a progressive global illumination solution, we also discuss some previous work on perceptually-driven global illumination techniques and image fidelity metrics used in rendering.

2.1 Progressive Global Illumination Techniques

Recent progress in global illumination algorithms and the increasingly better performance of computers and graphics hardware has made it possible to develop progressive and highly interactive realistic rendering techniques. In this section we narrow our discussion to those techniques that best match our design requirements. In particular, we do not discuss techniques that operate solely in the image space and cannot provide us with solutions in the object space. We also omit discussion of multipass

solutions that require expensive rendering (usually ray tracing), although the view-independent solutions we discuss are widely used as the first pass in multipass techniques.

Progressive, view-independent lighting calculations can be performed using radiosity techniques [Sillion and Puech 1994]. Interactive updates of global illumination for so-called incremental radiosity solutions for dynamic environments are reported [Sillion and Puech 1994], which make local recalculation of lighting possible for some simple scene changes, such as moving a light source or an object, e.g., Drettakis and Sillion [1997]. While we would like to take some advantage of dynamic solutions, in general, radiosity does not meet our design requirements well. Radiosity techniques are costly and algorithmically complex for environments with general BRDFs [Sillion and Puech 1994]. Also, it is difficult to bound the radiosity simulation error tightly and conservatively with computationally acceptable overhead [Lischinski et al. 1994; Veach 1997; Walter 1998]. The radiosity solution usually proceeds iteratively, and the resulting iterations may be extended in time for complex scenes. Now, if radiosity computations are stopped at an arbitrary moment (before a full iteration), the solution can be significantly biased, i.e., misleading for the user, since deterministically-guided computations could be focused on selected scene regions only. The radiosity performance depends strongly upon *how* the scene is represented (e.g., the number of polygons tessellating curved surfaces) instead of *what* the scene represents [Veach 1997]. The latter problem can be partially reduced by hierarchical cluster-based radiosity techniques [Smits et al. 1994; Sillion 1995]. However, these techniques require that complex data structures of links must first be built between all pairs of interacting surfaces and clusters. This requires significant storage and, because of this, links are not usually stored in databases of scenes but must be recalculated when scene data is retrieved from archives. Link computation is performed at the initial step in computing, which may substantially delay the first possible image inspection by the user. There were some successful attempts in getting rid of links [Myszkowski and Kunii 1995; Stamminger et al. 1998]. In this work we compare our new global illumination technique with our former progressive, linkless hierarchical cluster-based radiosity [Myszkowski and Kunii 1995].

Stochastic global illumination techniques have many advantages over radiosity solutions in terms of our design requirements, discussed in the introductory section. They scale well with the complexity of the scene, handle reflection functions of arbitrary complexity with ease [Heckbert 1990; Walter et al. 1997; Veach 1997; Jensen and Christiansen 1998], and the solution error can be estimated using statistical measures [Veach 1997; Walter 1998]. In this discussion we mostly focus on stochastic photon tracing from light sources toward the scene that fits our requirements the best, e.g., Appel [1968]; Arvo [1986]; Shirley [1990]; or Heckbert [1990]. As a result of photon tracing, the lighting function is available implicitly as the density of photon-hitting points and its explicit form can be reconstructed using density estimation techniques [Heckbert 1990].

Using a simple histogram method, in which object surfaces are subdivided into buckets and the number of photons hitting every bucket is stored, lighting reconstruction is almost immediate [Heckbert 1990]. A serious drawback to naive photon-bucketing is noisy images at the early stages of computation. This problem may also persist through later stages for small polygons, e.g., curved surfaces, leaves of plants, which may be hit by a small number of photons or not hit by any photon. Another serious problem with photon-bucketing into fixed buckets is discretization error in the reconstructed lighting function, which is difficult to control. There were some attempts at developing adaptive bucketing schemes [Heckbert 1990; Tobler et al. 1997], but a certain percentage of photons is discarded as buckets are refined (e.g., 25% in Tobler et al. [1997]).

Loss of photons can be prevented and better control of the discretization error can be achieved using techniques recently proposed by Walter et al. [1997; 1998]. The basic algorithm consists of three consecutive phases: (1) photon tracing and storage (usually on the hard disc); (2) lighting reconstruction at vertices of the fine uniform mesh; and (3) mesh decimation using an empirical model of luminance perception. Processing decomposition into three stand-alone phases simplifies many tasks such as memory management, software maintenance, and efficient mapping of computations into parallel architectures. However, it is not suitable for interactive applications because immediate image display based on a partial solution becomes difficult. Images generated using this method are of excellent quality, but reported timings are on the order of hours for scenes of medium complexity. Walter chose to compute the complete global illumination solution within the density estimation framework, which has some advantages in terms of estimating simulation errors and implementation simplicity. However, a huge number of photons must be traced to reconstruct lighting patterns of high spatial frequencies and high contrasts that are typical for direct illumination [Shirley 1990]. The density estimation phase, based on kernel methods with adaptive support, is also costly, and takes 10%–20% of the total processing time [Walter 1998]. Since photons are sorted on the basis of surfaces hit by these photons, shading quality problems for small polygons may appear. To overcome this, Walter [1998] proposes sharing hit information between neighboring polygons that share vertices. However, Walter does not discuss any solution to find neighboring polygons quickly when information on connectivity between polygons is missing, e.g., for polygons modeling independent objects and for polygons physically isolated in space. Rapid access to such information is crucial for proper selection of the kernel support size, thus affecting the efficiency of density estimation.

Jensen developed the photon maps technique [Jensen 1996] consisting of two phases: (1) stochastic photon tracing and construction of photon maps, and (2) ray tracing involving reconstruction of the lighting function via density estimation based on the photon maps. Although it is a multipass method, we include this method in our discussion due to its interesting features. Jensen shows that images of good quality can be obtained using a

significantly smaller amount of photons in comparison with Walter et al. [1997], when direct lighting is not derived from photon maps but is explicitly computed during ray tracing [Jensen and Christiansen 1998]. Also, the nearest neighbor density estimation method used by Jensen adapts well to local photon distribution, and efficiently removes noise from the lighting function (although there might be some problems with the accuracy of lighting reconstruction in using this method [Myszkowski 1997; Walter 1998]). Photons are accessed in 3D space via the balanced kd-tree data structure, which makes the photon organization completely separate from geometric representation. To make a rapid search of the nearest photons for every ray-object intersection point possible, all photon data must be kept in memory, which, in practice, imposes limitations on the reconstructed lighting accuracy that can be achieved on a given computer [Walter 1998]. The reported rendering rates are not interactive, due mostly to ray tracing costs. The lighting function reconstruction from the photon maps, which is performed for every pixel, incurs a non-negligible portion of these costs.

Note that all photon methods show much better performance for indoor scenes than outdoor scenes, in which case the probability that photons interacting with surfaces “escape” from the scene, without contributing to the solution, is high. In this paper we consider applications mostly involving indoor lighting simulation and rendering, so we are not concerned with this problem.

Another recent trend in rendering is to use standard graphics hardware in a nonstandard way to generate images with a realistic look. While the image may resemble some advanced lighting effects, global illumination is not in fact performed, e.g., Diefenbach [1996]. A notable exception is instant radiosity [Keller 1997], which is based on the concept of the deterministic pseudorandom walk, and takes advantage of graphics hardware to perform view-dependent global illumination computing with rendering rates of a few seconds for scenes of medium complexity. The basic algorithm can be extended to handle a general BRDF and to provide a view-independent solution by accumulating the results of rendering in textures. However, the performance of such an extended instant radiosity was not reported. To make such a solution practical, many technical problems must be addressed; e.g., how to decide texture resolution to avoid texture memory paging and secure high-quality shading. Also, while the algorithm ultimately converges to the correct solution, further research is required to make continuous monitoring of simulation errors possible.

Slusallek et al. propose the lighting networks concept [Slusallek et al. 1998], in which a composite global illumination algorithm is configured by the user through combining component algorithms to exploit their strengths for a given task. However, such configuring requires significant user knowledge and experience to select the proper algorithms and arrange them into a workable network. In this research we also consider a composite algorithm, but its components are preselected and automatic switchover

between them is performed based on considerations about perception, whose main goal is to improve progressive refinement of image quality.

Most techniques discussed above do not consider perceptual factors to improve image quality progression as perceived by the human observer. Taking these factors into account to enhance progressivity of global illumination solutions as a function of computation time is a major goal of this work. In the following section, we discuss previous work on perceptually-based rendering solutions.

2.2 Perceptually-Driven Global Illumination Algorithms

Models of visual perception have recently attracted more attention in realistic image synthesis [Greenberg et al. 1997]. In the context of global illumination techniques, some research was done on perception-driven radiosity [Gibson and Hubbard 1997; Martin et al. 1997] and meshing [Myszkowski et al. 1994; Gibson and Hubbard 1997; Hedley et al. 1997; Walter et al. 1997] (see Prikryl and Purgathofer [1998] for a more complete survey of similar solutions). All these techniques use perceptual error metrics at the atomic level (e.g., every light interaction between patches, every mesh element subdivision), which put a certain amount of overhead on the procedures that are repeated thousands of times in the course of computation. This imposes severe limitations on the complexity of the HVS models, which in practice are restricted to models of brightness and contrast perception.

Recently, more advanced (and costly) HVS models were used in rendering to develop higher level perceptual error metrics that operate on complete images. The main motivation for application of such models is the poor performance of commonly used mean-squared error metrics in predicting the differences between images that can be perceived by the human observer [Daly 1993; Rushmeier et al. 1995; Gaddipatti et al. 1997]. Work done by Rushmeier et al. [1995] is one of the first attempts to provide perceptually adequate metrics of the differences between images in the context of realistic image synthesis and global illumination algorithms. Martens and Myszkowski [1998] studied the applicability of the VDP [Daly 1993] to some typical rendering tasks via human psychophysical experiments with 11 subjects. The experiments show a good match with the VDP predictions for shadow and lighting pattern-masking by textures, and in comparing the perceived quality of images generated at subsequent stages of the progressive radiosity solution.

There are also some successful attempts to embed such advanced image quality metrics directly into rendering and global illumination solutions. Bolin and Meyer [1998] developed an efficient approximation of the Sarnoff visual discrimination model (VDM) [Lubin 1995], which make it possible to use this model to guide samples in a rendered image. Because samples are only taken in areas where there are visible artifacts, some savings in rendering time compared to the traditional uniform or adaptive sampling are reported. Myszkowski [1998b] shows some applications of the VDP to

decide upon stopping conditions of global illumination solutions and to drive adaptive mesh subdivision, taking into account visual masking of the mesh-reconstructed lighting function by textures. Ramasubramanian et al. [1999] developed their own image quality metric, which they applied to predict the sensitivity of the human observer to noise in the indirect lighting component. This made possible more efficient distribution of indirect lighting samples by reducing their number for pixels with higher spatial masking (in areas of images with high frequency texture patterns, geometric details, and direct lighting variations). All computations are performed within the framework of the costly Monte Carlo particle tracing algorithm [Kajiya 1986], and a significant speedup of computations is reported compared to the sample distribution based on purely stochastic error measures.

2.2.1 Discussion. Embedding advanced HVS models into global illumination algorithms is very attractive because computing can be perception-driven, specifically for a given scene. However, the costs incurred by such models introduce an overhead in computing the actual lighting, which may become more significant the more rapid the lighting computation becomes. The potential gains of such perception-driven computing can be easily canceled by this overhead, depending on many factors, such as scene complexity, performance of a given lighting simulation algorithm for a given type of scene, image resolution, and so on. The HVS models can be simplified to reduce the overhead, e.g., Ramasubramanian et al. ignore spatial orientation channels in their visual masking model, but then underestimation of visible image artifacts becomes more likely. To prevent such problems and to compensate for ignored perceptual mechanisms, more conservative (sensitive) settings of the HVS models should be applied, which may also reduce gains in the lighting computation driven by such models.

It seems that keeping the HVS models at some high level of sophistication and embedding them into rendering algorithms, which are supposed to provide a meaningful response rapidly, e.g., in tens of seconds or single minutes, may be a difficult task. For example, full processing of the difference map between a pair of images at a resolution of 256×256 pixels using the VDP model [Daly 1993] takes about 20 seconds on a R10000, 195 MHz processor, and such processing should be repeated a number of times to get reasonable monitoring of progress in image quality. In this paper we explore an alternative approach, in which the advanced HVS models are used only at the design stage of the global illumination algorithms and the tuning of their parameters. Thus, the resulting algorithms can spend 100% of their computation time in lighting simulation, and the costs of HVS processing (which is performed offline) are of secondary importance.

In this research we decided to use the VDP model [Daly 1993], considered one of the leading HVS models for predicting the differences between images perceivable by the human observer [Li et al. 1998]. Our choice of

the VDP is based on the positive results of the VDP validation in psychophysical experiments for tasks similar to our current application of the VDP [Martens and Myszkowski 1998], and the VDP integrity experiments [Myszkowski 1998b] (extensive documentation of these experiments can be found on our web page [Myszkowski 1998a]). The robustness of the VDP in predicting perceivable differences between images (both natural and synthetic) is also independently reported by Rushmeier et al. [1995] and Li et al. [1998].

3. ALGORITHM DESIGN SUPPORTED BY VDP RESPONSES

In this section we describe our procedure for designing a novel global illumination algorithm with the objective of achieving good progress in image quality in the course of computing. We consider a pool of basic algorithms and estimate the progress of computation for every algorithm in terms of minimizing perceived differences between the intermediate and final images as a function of time. We assume the final image to be the one corresponding to the converged solution within some threshold of negligible error. We apply the VDP to get quantitative measures of the perceived differences for all basic algorithms we investigate. We use the VDP response to rank the basic algorithms on the basis of their performance in minimizing the differences at various stages of computing. We repeat such ranking experiments for a number of indoor scenes. Based on such extensive ranking, we develop a hybrid algorithm in which a basic algorithm with the highest rank at a given stage of computing is always activated. While the rank-driven ordering of the basic algorithms is the same across all scenes tested, the optimal selection of switchover points between the sequentially executed algorithms depends on the given scene characteristics. Ideally, the switchover points should be selected automatically, based on the performance of the component algorithms for a given scene, which could be measured by online VDP computation. However, doing the VDP computation at global illumination runtime is not acceptable due to the high costs of VDP processing. To overcome this problem, we decided to elaborate a robust heuristic for selecting switchover points, which provides good progress in image quality for a wide range of indoor scenes. For this purpose, we designed another experiment involving VDP offline (our experimental setting is shown in Figure 2). This setting is of general use and can be easily applied to any set of global illumination algorithms for selecting the best basic algorithm for a given task and computation stage.

In the following section we briefly discuss the representative algorithms selected by us for the algorithm pool, we then evaluate their performance using the VDP. On the basis of the results, we introduce some algorithmic improvements in the original techniques and propose a more efficient global illumination solution that uses these improved techniques in a nonstandard way. Finally, we discuss the derivation of our heuristic for the selection of switchover points.

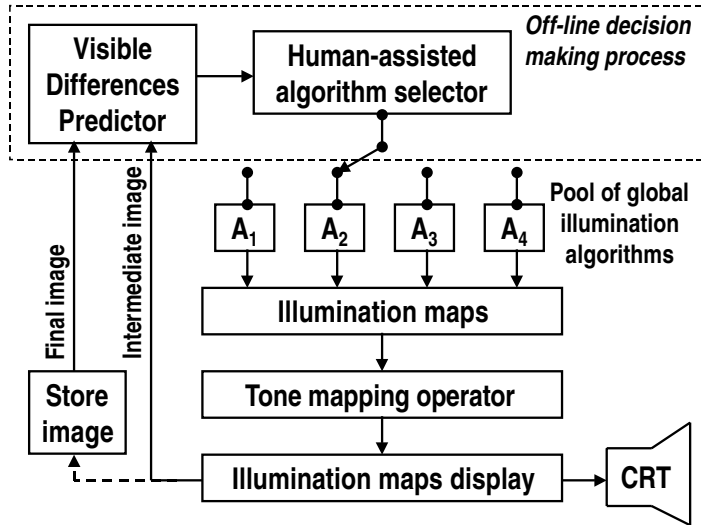


Fig. 2. Experimental setting for evaluating image quality progression and selection of switchover points between global illumination algorithms (human-assisted selection is based on minimizing the perceptual distance between the intermediate and final images).

3.1 Pool of Basic Algorithms

In this paper we consider the following view-independent algorithms:

- Deterministic direct lighting (DDL) computing with perceptually-based adaptive mesh subdivision [Myszkowski et al. 1994];
- Hierarchical (linkless and cluster-based) progressive radiosity (HPR) [Myszkowski and Kunii 1995] for indirect lighting computation. By default, a precalculated fixed mesh is used to store the resulting lighting. Optionally, this mesh can be further refined when solution for the fixed mesh is completed, but it may be costly for complex scenes [Lischinski et al. 1993];
- Density estimation photon tracing (DEPT) from light sources with photons bucketed into a nonadaptive mesh. By Direct DEPT (DDEPT) we denote buckets with photons coming directly from light sources, and by Indirect DEPT (IDEPT) we denote a different set of buckets with photons via at least one reflection.

Obviously direct and indirect lighting computation techniques are complementary, but in practice the following combinations of these basic algorithms are used: DDL+HPR, DDL+IDEPT, and DDEPT+IDEPT (DEPT for short). In this research we propose more elaborate compositions of these basic techniques to enhance progress in image quality. Before we move to this point, however, we summarize advantages and disadvantages of our basic techniques, which usually generalize well to the conventional wisdom about other hierarchical radiosity and density estimation solutions.




Fig. 3. Examples of templates of mesh element subdivision: (a) for a shadow covering an internal segment of an edge, and (b) for a shadow crossing two edges. Samples at the vertices and edge centers are computed first, and when perceivable lighting discontinuities are suspected, an additional one or two samples are inserted along every edge. Based on the location of suspected discontinuities, a template is selected, which may require some extra samples along inserted edges, as in (a).

3.1.1 Deterministic Direct Lighting. The DDL algorithm is based on adaptive meshing [Sillion and Puech 1994], which is used to reconstruct perceptually important lighting patterns [Myszkowski et al. 1994; Gibson and Hubbard 1997] (we provide more details on our perception-based oracle that drives mesh subdivision in Section 4). Computation is performed iteratively, and the goal of every iteration is progressive improvement in the reconstructed lighting quality using the minimal number of mesh elements possible (only triangles are used). Obviously, the initial mesh cannot be too coarse in respect to the meaningful details of direct lighting that are supposed to be reconstructed. The lighting function is sampled up to five times along every triangle's edge to provide more exact information on the lighting function's behavior. This makes it possible to perform edge splitting at the sample point location in the proximity of suspected lighting function discontinuities, while the remaining samples can often be reused in subsequent iterations [Myszkowski et al. 1994]. To improve the alignment of subdivided mesh elements to such discontinuities, a number of element subdivision templates are provided, which correspond to the most common lighting patterns within an element. Examples of two such templates are shown in Figure 3. An appropriate template is selected based on the distribution of displayable RGB values, which are perceptually-derived from the luminance values at sample points along all three edges. If the distribution of RGB values is too complex, i.e., does not fit any template, then simple recursive subdivision with new vertices located at the edge centers is performed.

The DDL algorithm does not require lighting and mesh recomputation from scratch for the limited scope of light source changes because *negative* light can be shot for light sources that are turned down or changed in their positions [Sillion and Puech 1994]. The DDL technique performs well, even for many light sources, under the condition that their impact on the scene illumination is localized by the goniometrical diagrams (usually a reasonable assumption for practical scenes). For a huge number of globally

operating light sources the DDEPT could be a better choice, although the fixed mesh makes reconstruction of lighting patterns of a size smaller than the mesh grain impossible.

3.1.2 Hierarchical Progressive Radiosity. The HPR technique is based on the shooting iteration [Sillion and Puech 1994], which is used within the hierarchical radiosity with a clustering framework [Myszkowski and Kunii 1995]. The accuracy of the lighting simulation for this technique is extensively validated against measurement data, and rendering quality is successfully compared with photographs of real-world scenes [Myszkowski et al. 1997]. The HPR algorithm, like a majority of radiosity solutions, does not support a general reflectance function. Specular reflections are supported, but in practice simulation of specular reflection by the “image method” [Sillion and Puech 1994] can only be applied for big planar mirrors due to its prohibitive costs. HPR computation time may strongly depend on the input geometry, and the simulation error is difficult to control.

3.1.3 Density Estimation Photon Tracing. Our DEPT technique is similar to other solutions for handling particle tracing from light sources (for a detailed description, see Heckbert [1990] or Walter [1998]). We assume that every particle carries the same power and we use simple unstratified sampling, which is a reasonable choice for view-independent solutions [Walter et al. 1997]. We observe that the DDEPT and IDEPT techniques provide feedback extremely rapidly for distributing global illumination. This is possible because tracing tens of thousands of photons can be accomplished easily within a single second on modern computers using state-of-the-art ray tracing software (e.g., see timings for photon tracing in Jensen and Christiansen [1998]). Such a number of photons may be sufficient to provide meaningful information on lighting distribution. In our approach, we use a simple histogram method [Heckbert 1990], which makes immediate lighting reconstruction possible at mesh vertices based on the number of photons that hit neighboring mesh elements. To reduce the discretization error, we use texture-based bucketing for caustic photons, i.e., photons reflected/refracted by specular surfaces [Jensen 1996]. Computational overhead imposed by the histogram method is negligible, so periodic monitoring of the progress of the solution becomes cheap using graphics hardware. Such functionality is missing in both Walter’s approach [Walter et al. 1997] and in the photon maps technique [Jensen 1996], since reconstruction of the lighting function is deferred until the photon-shooting phase is completed.

The well-known drawback of traditional histogram methods is lack of adaptability of the bucket (mesh element) size, which manifests itself in annoying dark (or bright) spots for buckets hit by too few (or too many) photons, or not hit at all. This results in poor image quality at the initial stages of computation when too few photons are traced and the variance of the solution is high (Figure 4(a)). Also, at the later stages of computation, even if the global variance is reduced for small objects, the local variance may still be high and expensive to reduce. Figure 5(a) shows a relevant




Fig. 4. Reducing excessive noise at early stages of the DEPT computation: results of lighting reconstruction for the scene built of 116,600 mesh elements after 10 seconds of photon tracing (Pentium II, 400 MHz processor) in two cases: (a) without filtering, and (b) with filtering.

example of a scene for which unpleasant noise is still present after one hour of photon tracing.

To reduce the problem of noise, we propose a novel, inexpensive filtering technique, which operates directly on the photon buckets, and is activated exclusively during rendering of the illumination maps. To achieve the required level of accuracy of reconstructed lighting at a given vertex v , a certain number of photons hitting region S , which is centered at v , must be considered. Such an accuracy condition can be met by adjusting the size of S based on the local photon density. Effectively, the size of S controls the level of noise in reconstructed lighting, which can be reduced at the expense of an increasing bias in the solution [Silverman 1985]. To somehow find a good tradeoff between the random and discretization errors, we apply our stochastic estimates of the variance of local illumination, which are used to calculate the size of S . (In Appendix A we provide formal derivation of our mathematically sound measure of illumination accuracy, and a detailed description of our algorithm for selecting S .) Our filtering requires a rapid search of neighboring buckets around v . If we were to rely on the neighborhood relations between the mesh elements, the result would strongly depend on the geometrical model and would not provide complete information for separate objects. Instead, we build a static and balanced 3D kd-tree structure for all mesh elements, which makes searching the neighbor buckets extremely fast. Figures 4(b) and 5(b) show the improvement in image quality as the result of our filtering.

In this paper we limit our discussion to the photon-bucketing approach suitable for interaction and progressive rendering purposes. Since the final computations of direct lighting are performed using the DDL method (or




Fig. 5. Reducing excessive noise for small polygons that receive a small number of photons or no photon: results of lighting reconstruction after 1 (top row), 15 (middle), and 60 (bottom) minutes of photon tracing (Pentium 200 MHz processor): (a) without filtering, and (b) with filtering. The scene is built of over 130,000 polygons and illuminated by 129 light sources.

ray tracing for high-quality still images), in practice, the quality of indirect lighting reconstructed via photon-bucketing is good. If indirect lighting reconstruction of higher quality is required by some applications, then along with photon-bucketing, all photons could be dumped into a file on the hard disc for deferred processing, e.g., using more advanced density estimation methods developed by Myszkowski [1997] and Walter [1998].

3.2 Measuring the Performance of the Basic Algorithms

We investigated the performance of our basic techniques (the DEPT technique with and without illumination map filtering), and their simple combinations (the DDL+HPR, DDL+IDEPT, and DDEPT+IDEPT techniques) in terms of perceived differences between the intermediate and final images using the VDP responses. We call the increasing similarity between intermediate and final images “perceptual” convergence, which is quantitatively predicted by the VDP as a function of time. The VDP response provides the probability of difference detection between a pair of




Fig. 6. Test scene SPOT: (a) full global illumination solution, (b) indirect lighting only, (c) direct lighting only.

images, which is estimated for every pixel. We measured the difference between images as the percentage of pixels for which the probability of difference detection is over 0.75, which is the standard threshold value for discrimination tasks [Daly 1993]. In all tests performed, we use images with 512×512 resolution. The diagonal of the images displayed on our CRT device was 0.2 meters, and we assumed that images were observed from the distance of 0.5 meters. All timings in this paper, unless explicitly stated otherwise, are measured on a Pentium II, 400 MHz processor.

In this paper we report results obtained for three scenes of various geometrical complexities with drastically different lighting designs, which we will refer to as the SPOT, KITCHEN and ROOM test scenes:

- SPOT: most of the scene is illuminated by indirect lighting only (shown in Figure 6). The scene is illuminated by 3 luminaires, built of about 5,000 polygons, and the original scene geometry was tessellated into 30,200 mesh elements.
- KITCHEN: most of the scene is illuminated by direct lighting (shown in Figure 7). The scene is illuminated by 8 luminaires and built of over 131,700 polygons (tessellated into 350,600 mesh elements). The scene represents a complete two-room apartment with furniture (Figure 7(b)). For the VDP processing we selected a single view of the kitchen, but lighting computations are performed for the whole apartment.
- ROOM: both direct and indirect lighting contribute significantly to scene illumination (shown in Figure 8). The scene is illuminated by 14 luminaires and built of over 50,100 polygons (tessellated into 121,900 mesh elements). The scene also represents a complete apartment with furniture (Figure 8(b)). We selected a single view featuring many small isolated polygons in the foreground.

At first we used the VDP to compare the performance of our illumination map filtering, which was developed to improve the quality of intermediate images obtained using the DEPT technique (see Section 3.1.3). We assume



Fig. 7. Full global illumination solution for the KITCHEN scene (a) view selected for VDP processing, and (b) overall view showing the scene complexity.

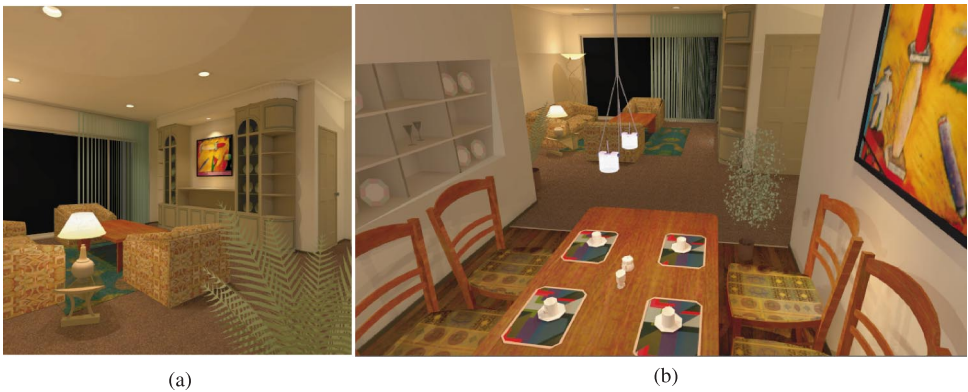


Fig. 8. Full global illumination solution for the ROOM scene (a) view selected for VDP processing, and (b) overall view showing the scene complexity.

that the final images used for VDP computation are based on DEPT solutions with and without filtering (for a given method, we consider the final image generated with this particular method) that converge within some negligible error tolerance. The graphs in Figure 9 depict the VDP-based quantitative measures for the perceptual convergence of image quality as a function of time for our SPOT, KITCHEN and ROOM scenes. As expected, filtering improves the image at the initial stages of computation, and then its importance gradually decreases as solution variance is reduced. Indeed, the size of filter support is reduced automatically as the variance of local illumination decreases at subsequent stages of computation (see Appendix A for details). Variance reduction might require substantial computation time for scenes with many small, isolated polygons (see the graphs in Figure 9(c)). Note that filtering does not introduce any additional bias into the converged solution, since the original estimate of illumination at vertices is used. Due to these favorable features, in all the following experiments, we use the DEPT algorithm with filtering exclusively.

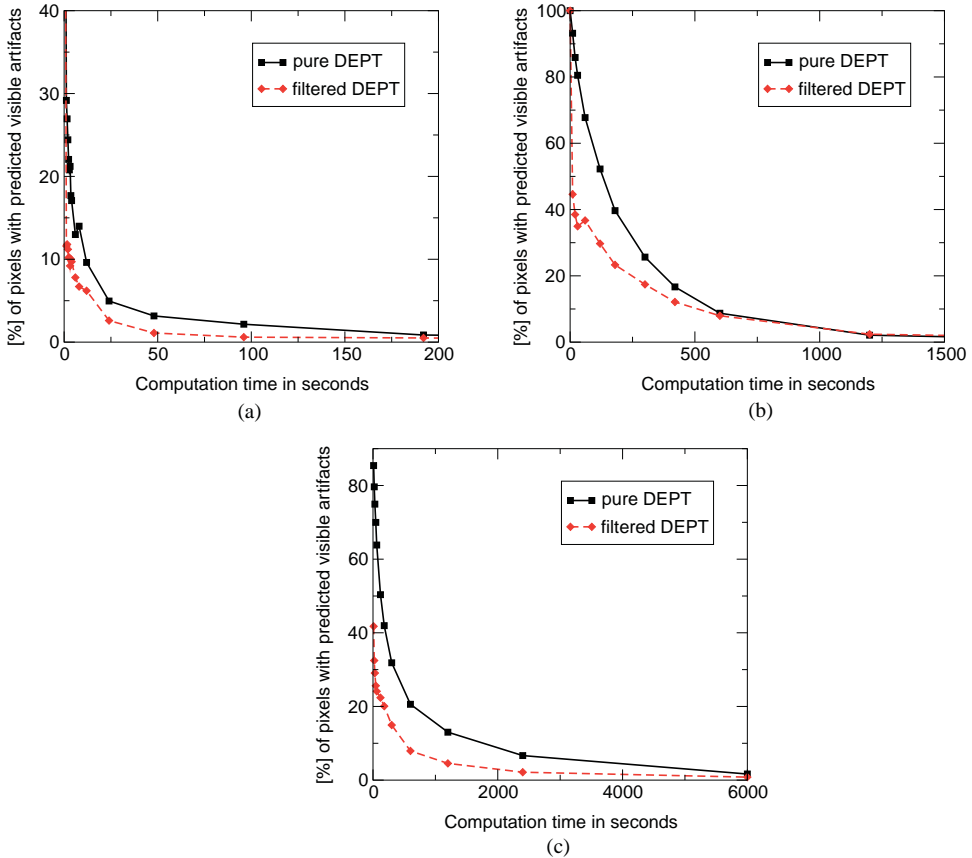


Fig. 9. Plots of the VDP results (predicted differences between the intermediate and final images) measuring the performance of the DEPT algorithm with and without filtering for tests (a) SPOT, (b) KITCHEN, and (c) ROOM.

Our experiments focus on comparing the performance of the following global illumination algorithms: DDL+HPR, DDL+IDEPT, and DDEPT+IDEPT. We assume that the final images used for the VDP computations are based on the DDL+HPR and DDL+IDEPT global illumination solutions, which converge within some negligible error tolerance. The final images obtained using these methods are usually only slightly different (minor discrepancies can be explained by various approximations assumed by each of these completely different algorithms, e.g., different handling of the visibility problem, the lighting function discretization during computations used by the HPR technique). To eliminate the influence of these differences on the VDP response, for a given method we consider the final image generated using this particular method. The only exception is the DDEPT+IDEPT method, for which we use the final image generated using the DDL+IDEPT technique because it provides more accurate direct lighting reconstruction for a given mesh/bucket density. In this study we use scenes with purely diffuse reflectance properties, which

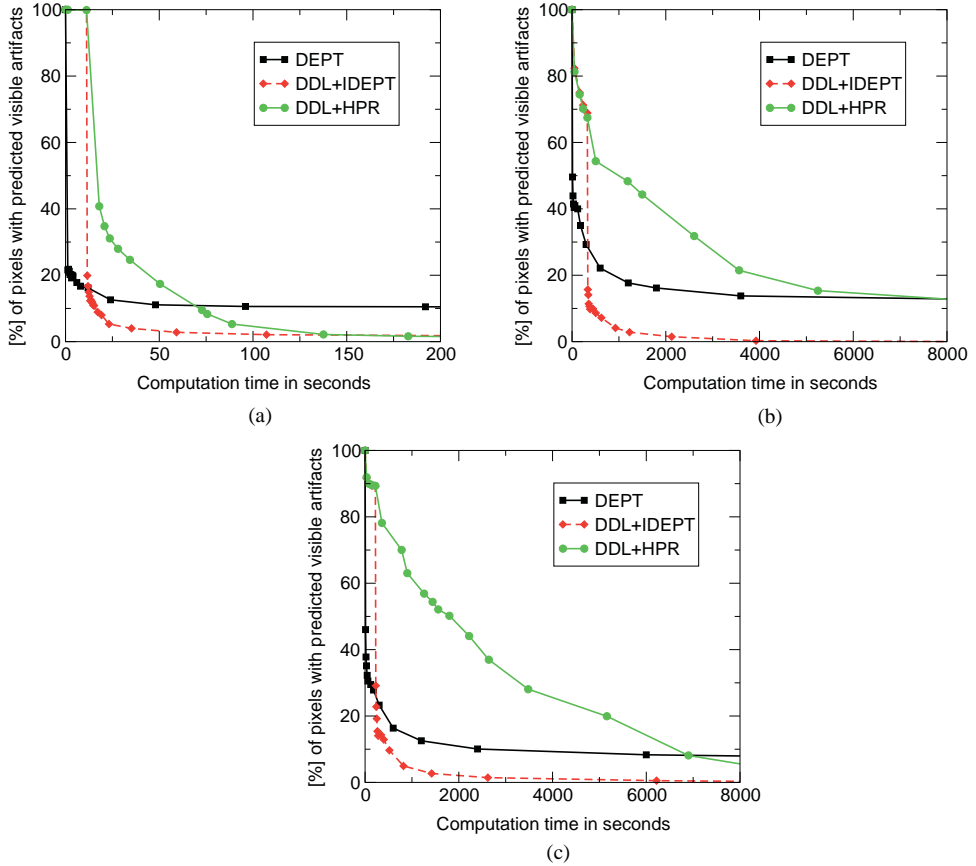


Fig. 10. Plots of the VDP results (predicted differences between the intermediate and final images) measuring the performance of the DEPT, DDL+IDEPT, and DDL+HPR algorithms for tests (a) SPOT, (b) KITCHEN, and (c) ROOM.

provide similar final images for the DDL+HPR, DDL+IDEPT, and ray-tracing techniques. This makes possible to double-check direct lighting quality using the DDL and ray-tracing techniques. Also, the quality of indirect lighting can be confirmed using two independent HPR and IDEPT methods.

The graphs in Figure 10 show that the perceptual convergence of the indirect lighting solution for the HPR technique is slower than the IDEPT approach (direct lighting is computed using the same DDL method). In our implementation, both methods are based on the same ray-tracing kernel used to perform visibility computations required to evaluate form factors and tracing photons. We did not use the ambient light approximation or overshooting techniques because we are interested in physically-sound intermediate results. In our experience, the difference in performance between the IDEPT over HPR methods is far more significant for complex scenes. The HPR technique shows better performance for simple scenes only. Based on these results, we use the DDL+HPR technique for scenes

built of fewer than 500 polygons. For scenes of more practical complexity, we consider the DDL, DDEPT, and IDEPT techniques to optimize the progressive refinement of image quality.

The graphs in Figure 10 show that at the initial stages of computing the combination of DDEPT+IDEPT provides the best performance, and rapidly gives meaningful feedback to the user. At later stages, the DDL+IDEPT hybrid shows faster perceptual convergence to the final image. In both cases, we used the same fixed mesh to bucket photons. Due to the basic mesh-element granularity, many subtle details of direct lighting distribution could not be captured well using the DDEPT technique. For example, small and/or narrow lighting patterns may be completely washed out. Also, when shadows are somehow reconstructed, they can be distorted and shifted with respect to their original appearance, and their boundaries can be excessively smooth. The problem of excessive discretization error, which is inherent in our DDEPT method, is reduced by adaptive mesh subdivision used by the DDL technique.

Based on these observations of “perceptual” performance, we propose a new composite technique that uses DDEPT, IDEPT, and DDL:

- (1) First, stochastic computation of direct and indirect lighting should be performed.
- (2) Second, the stochastically computed direct component should gradually be replaced by its deterministically computed counterpart to reconstruct the fine details of the lighting function.
- (3) Finally, stochastic indirect computation should be continued until some stopping criterion is reached, e.g., a criterion that is energy-based in terms of the solution variance (some engineering applications may require precise illumination values), or perception-based in terms of perceivable differences between the intermediate and final images [Myszkowski 1998b].

All algorithms discussed use mesh vertices to store the results of direct and indirect lighting computations separately, so switching between them can be performed easily. The mesh is adaptively refined to fit the lighting distribution better in the case of the DDL technique only (see Section 3.1.1), but then indirect lighting computed using the IDEPT can be interpolated at the new vertices. To make this solution workable, switchover points between the basic algorithms should be chosen to optimize the progress in image quality. We address this issue in the following section.

3.3 Selecting Switchover Points between Basic Algorithms

We investigated the problem of switchover point selection among our basic algorithms DEPT (DDEPT+IDEPT), DDL and IDEPT experimentally. We again used the VDP to get quantitative measures of progress in image quality as a function of time points T_i at which switching was performed. At first, we assumed only two switchover points T_1 and T_2 at which DEPT

Table I. Integral Values of VDP-Predicted Differences for Time Span $[0, T]$, DEPT, DDL+DEPT, and Composite T_1 and T_2 , and T_1, \dots, T_N Algorithms. Only minimal integral values are shown for two points (T_1 and T_2) and multiple-points (T_1, \dots, T_N) switching strategies of our composite algorithm

Scene name	T [s]	Integral values [%]			
		DEPT	DDL+IDEPT	T_1 and T_2	T_1, \dots, T_N
SPOT	36	16.883	42.179	13.594	13.224
KITCHEN	675	34.444	48.450	24.324	24.256
ROOM	540	28.649	48.378	22.463	22.403

is replaced by DDL and then DDL is replaced by IDEPT, respectively. We investigated various choices of T_1 , which effectively measures the duration of the initial DEPT computation. We assumed that T_2 is decided automatically when the DDL computations are completed. The composite algorithm performance at initial stages of computing for tests SPOT, KITCHEN, and ROOM is shown in Figure 11 (parts a, b, and c). The thick lines between T_1 and T_2 depict possible performance gains for various choices of T_1 . To get quantitative measures of the overall algorithm performance, we integrated the area under the VDP plots for the DDL+IDEPT, DEPT, and our composite algorithm. The integration was performed during the time span $[0, T]$, where T is bigger than the longest T_2 that we considered for a given scene. The integral values were normalized by the value of T , and are summarized in Table I. As can be seen, our composite algorithm performs much better than the standalone DDL+IDEPT or DEPT methods. Figure 12 (parts a and b) depicts the integral values for various choices of T_1 in our composite algorithm. $T_1 = 0$ means that the DEPT computations are not activated at the initial stage of processing, so, effectively, only the DDL+IDEPT computations are performed. The integral values increase slowly when T_1 is bigger than the optimal switchover point (the minimum value of the integral in the graphs), which means that the choice of T_1 is not extremely critical. However, if T_1 is too small, the variance of indirect lighting remains high, which may cause poorer performance of our composite algorithm, especially for scenes with dominating indirect lighting. This effect is visible in Figure 11 (parts a and c) for $T_1 = 0.5$ and $T_1 = 10$ seconds, respectively, where the DDL computations cannot reduce the perceptual distance between the corresponding intermediate and final images. In both cases, the standalone DEPT performs better as a function of time, until the DDL computations are completed in our composite algorithm and IDEPT is activated to reduce the variance of an indirect lighting solution.

In all the experiments discussed so far it is assumed that there are only two switchover points, T_1 and T_2 . This means that DDL computations are performed continuously, until the final quality in direct lighting is achieved. However, a different strategy, which involves N switchover points T_1, \dots, T_N , is also possible. We investigated various choices of T_i ,

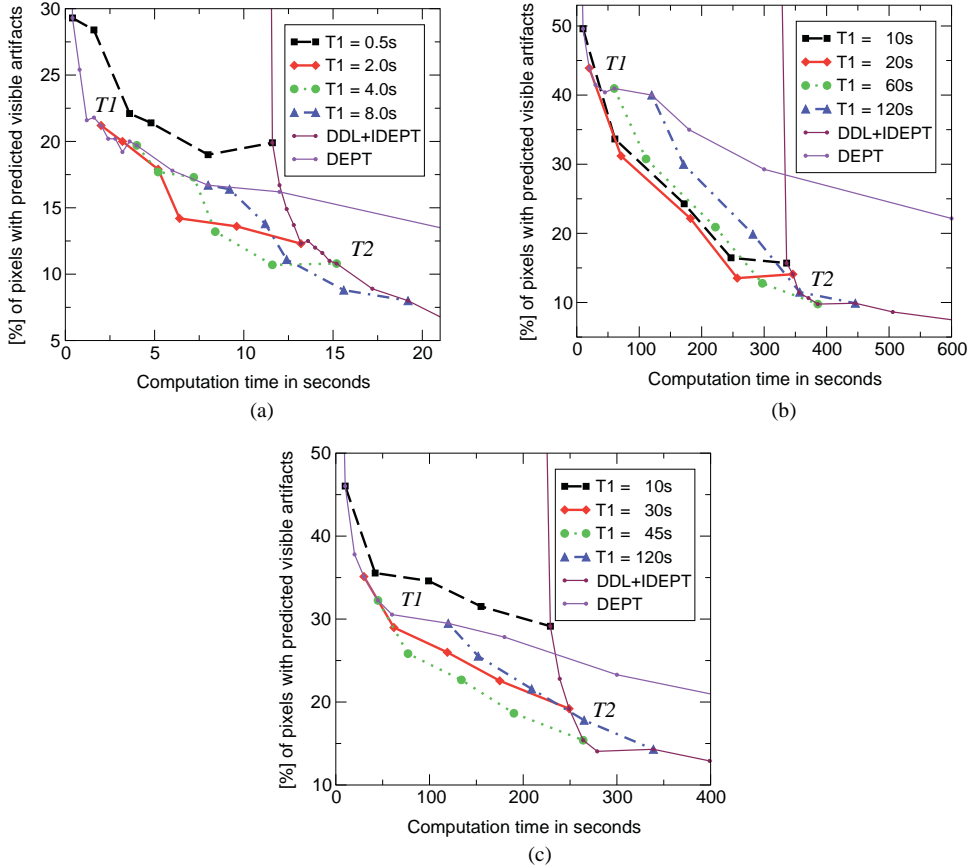


Fig. 11. Plots of VDP results (magnified from Figure 10) measuring the performance of DEPT and DDL+DEPT algorithms at the initial stages of computation for tests (a) SPOT, (b) KITCHEN, and (c) ROOM. The thick lines between two switchover points T_1 and T_2 depict possible performance gains (for various choices of switching time T_1) if the DEPT is replaced by the DDL at T_1 , and the IDEPT is activated at T_2 .

which controls switching between the DDL and IDEPT algorithms. For example, we performed the switching after completion of every single iteration of the DDL, or every two such iterations, and so on. We also changed T_1 , which effectively controls the initial DEPT computation time. As for the two-point switching strategy, we integrated the VDP-predicted differences between the intermediate and final images for various multiple points T_1, \dots, T_N switching strategies. In Table I, we show the global minimum value obtained for each test scene by considering all possible combinations of switchover point selections after every completed DDL iteration and for various T_1 . For example, in the KITCHEN test the optimal strategy involves four switchover points: the DEPT algorithm is switched to the DDL at $T_1 = 20$ seconds, after three iterations of DDL processing, the IDEPT is activated at $T_2 = 257$ seconds, then after 100 seconds of IDEPT

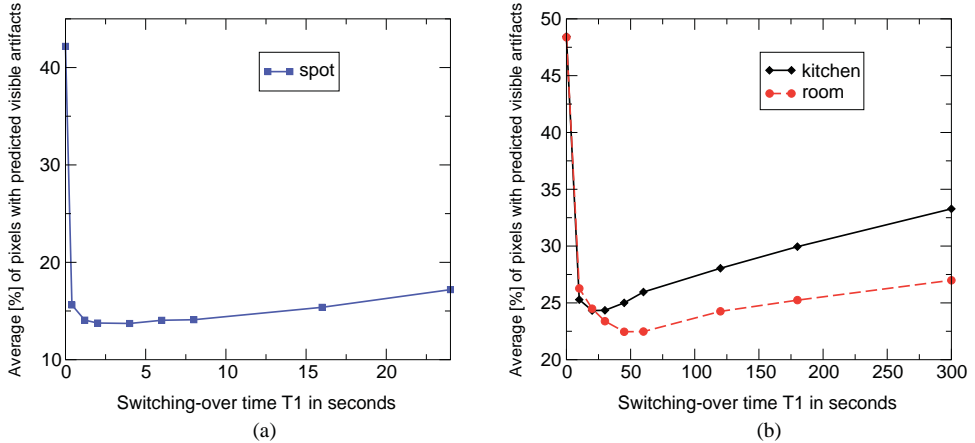


Fig. 12. Plots of VDP-predicted differences integrated during the same time span $[0, T]$ as a function of switchover point T_1 selection for tests (a) SPOT, (b) KITCHEN, and ROOM. For every test scene, T is chosen so it is bigger than the largest T_2 and used to normalize the integral values in the graph.

computing, the last iteration of the DDL is initiated at $T_3 = 357$ seconds, and after it completes at $T_4 = 446$ seconds, the IDEPT is activated. This strategy results in slightly better image quality than the two-point switching approach (Table I). Since the final iteration of DDL processing contributed little in perceivable image quality enhancement, it was more efficient to switch earlier to indirect lighting computation instead of proceeding with this iteration.

In general, gains in performance using the T_1, \dots, T_N strategy are negligible compared to the strategy based on well-chosen switchover points T_1 and T_2 (see Table I). Thus, for simplicity, in our composite algorithm we decided to use just the two switchover points. This decision is also justified by the fact that, in our implementation, additional iterations of the adaptive mesh subdivision can be requested by the user at any stage of the DDL or IDEPT computation without any penalty.¹ Hence the user can effectively execute the T_1, \dots, T_N switching strategy based on his judgement of the current image quality.

Since using the VDP to decide on selecting a switchover point online during lighting computations is too costly (see Section 2.2.1), we decided to design a heuristic to select the T_1 switchover point. We based our heuristic on extensive offline use of the VDP to evaluate the perceptual convergence of a global illumination solution for a number of scenes. As we have shown (see Figure 12 (parts a and b)), the choice of T_1 is not critical in terms of

¹The user may also change the thresholds controlling mesh subdivision to improve the quality of reconstructed lighting (or to reduce the complexity of a mesh). Mesh recalculation from scratch is not required. The computed lighting values at the mesh vertices are tested using the updated thresholds, and on this basis some vertices are inserted into (or removed from) the mesh.

progressive refinement in image quality. However, a too short T_1 may result in poor quality indirect lighting, which cannot be improved during the DDL computation. On the other hand, a too long T_1 may result in an undesirable delay in reconstruction of shadows and other shading details. Because of this, the upper bound for T_1 should be comparable to the computation time of the first iteration T_{i0} in DDL processing, after which the first rendering of a complete direct lighting distribution becomes possible. We can estimate T_{i0} well by measuring the timings of pilot photon tracing and by knowing the number of initial mesh vertices, the number of light sources, and estimating the average number of shadow feelers for linear and area light sources. Although T_{i0} is not a direct measure of the complexity of the lighting function to be reconstructed, it is a good predictor of such complexity.

Our heuristic for the selection of T_1 proceeds as follows. At first, we run the DEPT computations for time $T_\alpha = \alpha T_{i0}$ (where $\alpha = 0.1$, and $T_\alpha \geq 0.5$ seconds, since in our implementation we assume that 0.5 seconds is the minimal interval for sampling DEPT solution errors). We then estimate the *RMS* error \tilde{E} of the indirect lighting simulation (we provide a derivation of the *RMS* error measure for the DEPT algorithm in Volevich et al. [1999]). Based on the results of DEPT computations for multiple scenes, we assume that a reasonable approximation of indirect lighting can usually be obtained for the *RMS* error threshold value $E_{thr} \approx 15\%$. Taking into account the basic properties of stochastic solution convergence [Sillion and Puech 1994], we estimate the required computation time T_{thr} to reach the accuracy level E_{thr} as

$$T_{thr} = T_\alpha \frac{\tilde{E}^2}{E_{thr}^2},$$

and finally, we set T_1 as

$$T_1 = \min(T_{thr}, T_{i0}).$$

For simplicity, our heuristic relies on the energy-based criterion of indirect lighting accuracy. Obviously, in the perceptual sense this criterion does not guarantee the optimal T_1 switchover point selection. However, we found that this heuristic provides stable progressive refinement of rendered image quality for all tests we performed with various scenes. The robust behavior of our heuristic can be explained by the relative insensitivity of our composite algorithm to T_1 selection, and the strong low-pass filtering properties of our lighting reconstruction method at the initial stages of computation.

Figure 13 (parts a and d) shows the intermediate (computed at 3, 20, and 346 seconds) and final images of the KITCHEN scene, computed using our composite technique. Images in the second and third rows correspond

exactly to the switchover points $T_1 = 20$ seconds and $T_2 = 346$ seconds among the DEPT, DDL, and IDEPT algorithms. The composite algorithm performance in terms of the VDP responses for such a selection of T_1 and T_2 is shown in Figure 11(b). Only the DEPT computations were performed for images in the two top rows, and the 20 and 326 seconds of the DEPT and DDL computations were done, respectively, for images in the third row. Images in the second column (Figure 13(b)) show the absolute differences (normalized by the image mean) between the intermediate images (Figure 13(a)) and the final image (Figure 13(d)). The last column of images (Figure 13(c)) shows the corresponding responses of the VDP. Images in Figure 13 (parts b and c) were obtained via alpha blending of the original image (shown in grey-scale) with color, which encodes the absolute difference and difference-detection probability values, respectively, for every pixel. Color scales for encoding the differences in Figure 13(b) and probabilities in Figure 13(c) are shown in Figure 13(e). For example, the red and purple pixels in the VDP responses indicate difference-detection probability values greater than 0.75 (standard threshold value for discrimination tasks [Daly 1993]). As can be seen by comparing images in Figure 13 (parts b and c), VDP response provides a more reliable picture of the perceived differences by ignoring some artifacts of the lighting function that are of low contrast or are well masked by details in texture and geometry.

4. MAKING ADAPTIVE MESH SUBDIVISION MORE EFFICIENT

The appearance of images obtained at the initial DEPT stages of lighting simulation gives a good approximation of the final images, with the exception of views that contain many areas with strong direct lighting. DDL computations are performed to overcome these drawbacks and to guide adaptive mesh subdivision in improving the quality in the reconstruction of the lighting function. In our composite algorithm, DDL computing is performed after DEPT computing, which provides meaningful local estimates of global illumination at the early stages of processing. This feature of our algorithm contrasts with other global illumination methods. We take advantage of local estimates of global illumination to improve the performance of adaptive mesh subdivision.

The goal of the adaptive meshing is to reconstruct lighting functions without visible artifacts using a minimal number of mesh elements. In Myszkowski [1998b] it is shown quantitatively how the perception of artifacts in mesh-based shadow reconstruction is affected by the contrast at shadow borders. Based on our experimental results (see Myszkowski [1998a] for more details), we noticed that important savings in the number of mesh elements can be achieved if the meshing algorithm is based upon local estimates of the global illumination solution. Unfortunately, for many practical algorithms (including our former meshing solution [Myszkowski et al. 1994]) such estimates are not available, and adaptive meshing is performed based on direct lighting only. One notable exception is a meshing technique proposed by Gibson and Hubbold [1997] that uses an “ambi-

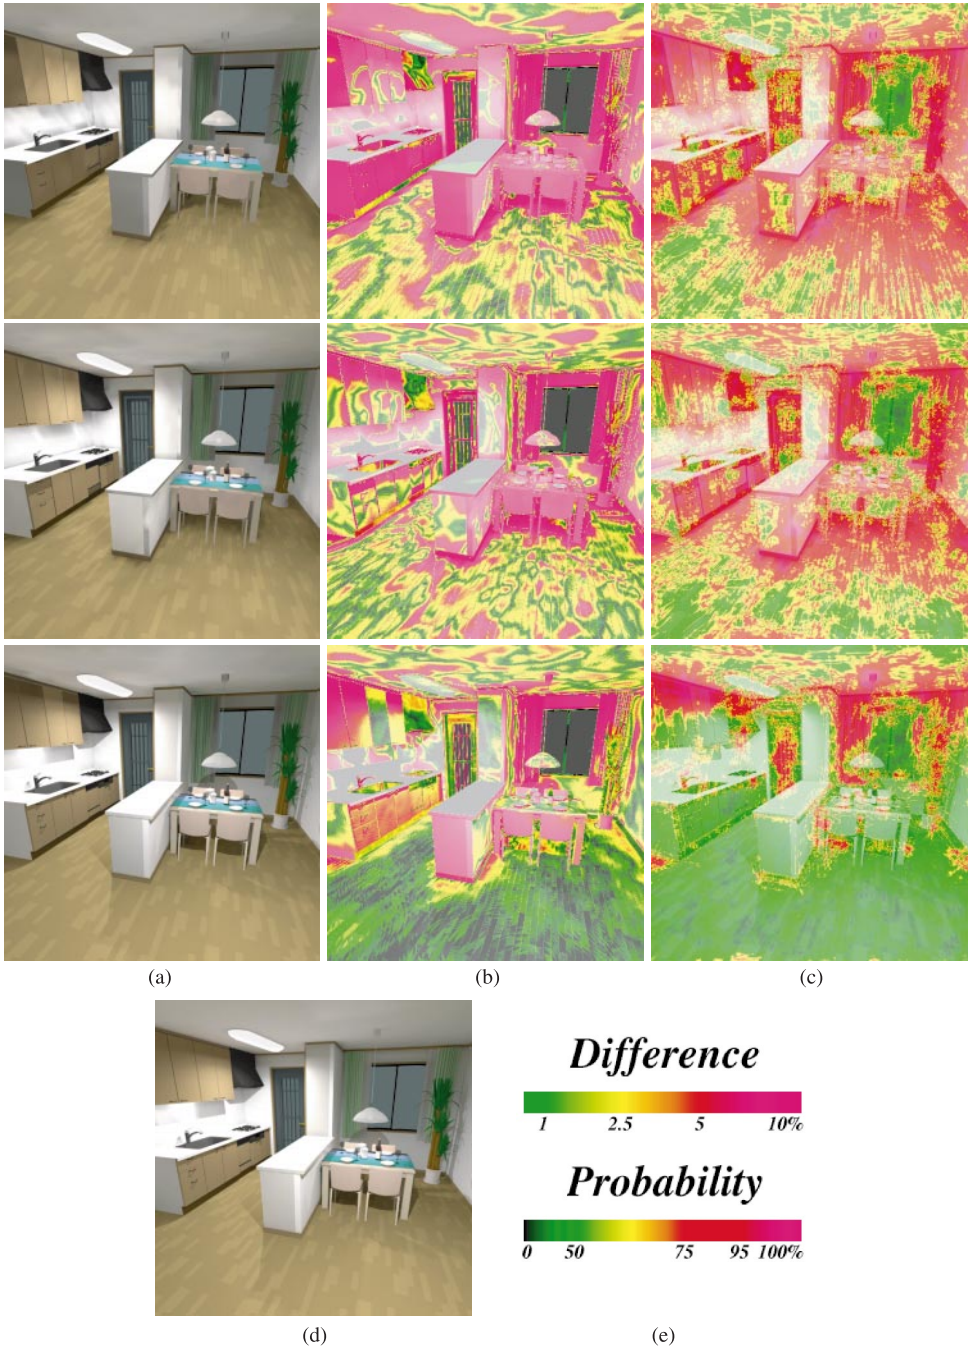


Fig. 13. Perceptual convergence of image quality as a function of computation time for our composite algorithm. Column (a) shows subsequent stages of computation at 3 (top row), 20 (middle), and 346 (bottom) seconds; column (b) depicts the absolute differences of pixel intensity between the current and fully converged solutions; column (c) shows the corresponding visible differences predicted by the VDP; column (d) shows the fully converged solution used as a reference. Color scales for encoding the differences in column (b) and probabilities in column (c) are shown in panel (e).




Fig. 14. Adaptive mesh subdivision for test SPOT: (a) based on direct lighting, (b) based on global illumination.

ent correction term” to compensate for the lack of knowledge of global illumination during the initial stages of the progressive radiosity solution. However, the ambient correction term is estimated globally for the whole scene on the basis of a simplistic global illumination approximation, which ignores visibility between surfaces and their orientations. This may lead to significant errors in local estimates of indirect lighting, which may affect the quality of mesh-based lighting reconstruction.

We observed that our meshing solution usually provides better results than those of Gibson and Hubbold because our local estimates of global illumination are based on the results of an actual lighting simulation. We use perceptually-based criteria to guide our adaptive mesh subdivision [Myszkowski et al. 1994]. For each sample point located along a given triangle edge, we transform the stimulus luminance values to predicted perceived brightness using Stevens’ power law [Tumblin and Rushmeier 1993], and a decision on edge splitting is made on the basis of the local differences in brightness. Obviously, the available local estimate of global illumination in the proximity of the processed edge makes possible a more reliable evaluation of the contrast at the shadow borders (which can be substantially suppressed by indirect lighting). This factor is of primary importance in avoiding excessive mesh subdivision that would slow down calculations.

In Figure 14(a) we show the results of our early adaptive mesh subdivision technique [Myszkowski et al. 1994], which is based on direct lighting only. The original uniform mesh is built of 30,200 triangles and subdivided into over 121,000 triangles. Significant reduction in mesh complexity (Figure 14(b)) is achieved when an estimate of indirect lighting is included in the original oracle [Myszkowski et al. 1994] that controls mesh subdivision. The resulting mesh is composed of just over 86,000 triangles. When we replaced estimates of indirect lighting with the constant ambient term as in Gibson and Hubbold [1997], over 97,000 mesh elements resulted.

Another unique feature of our composite algorithm is that it provides two independent estimates of direct lighting, which can be used immediately to improve the robustness of the mesh-based reconstruction of lighting patterns. DDL computation provides illumination values at selected sample points, e.g., mesh vertices. DDEPT computation provides similar estimates, which characterize the average illumination in certain regions surrounding a given vertex. By comparing the pointwise and regionwise estimates of direct illumination, we obtain additional characteristics of the lighting function within a mesh element. If the two estimates differ significantly we make the mesh subdivision threshold more conservative, since we can expect some complex lighting patterns of high contrast within element boundaries. Conversely, if the estimates are in good agreement, we reduce the number of sample points along the element edges, since we can expect more uniform lighting distribution. This extension of basic DDL computation is done at virtually no cost using our composite algorithm, since both estimates of direct illumination are available. In the tests that we performed we noticed that the quality of illumination maps using this technique is better, and we did not observe any significant changes in computation time. This means that more intensive mesh subdivision (forced by more conservative subdivision thresholds in the regions with anticipated complex lighting) is balanced by the reduced number of sample points in the regions with simple lighting distribution.

5. IMPLEMENTATION

We implemented our composite algorithm as a part of the Specter system (Integra, Inc.), a commercial lighting simulation and rendering software system. The Specter system features a graphical user interface, and we dedicated one of its control panels specifically for interactive use of the DDL and DEPT techniques (Figure 15). This panel also includes a control button for activating our composite algorithm, which at current settings is the default global illumination solution. The user may optionally activate the DEPT and DDL+IDEPT global illumination methods. The DDL and DEPT techniques are rather simple and intuitive for the user to control. The DDL computation requires two input parameters that can be changed interactively: the threshold values of RGB, used to make decisions on mesh subdivision along every edge, and the number of iterations that controls the depth of such subdivision. DEPT's basic control is limited to setting an acceptable level for a solution error, which is used as the stopping condition. The simplicity of the manual settings is an important feature of the DEPT algorithms from the user's point of view [Walter 1998]. As an option, the user may activate the illumination map filtering and control filtering accuracy. The user may also decide on the time for a periodical update of the illumination map rendering for a selected view. At the same time, an updated estimate of the lighting simulation *RMS* error is reported. The user may break-off computation at any moment, and be provided with an updated rendering and error estimate. Obviously, the solution is view-




Fig. 15. Control panel for three global illumination techniques: our composite algorithm, the DDL+IDEPT algorithm, and the DEPT algorithm.

independent, so rendering of any view of the scene is possible. The user may also decide to continue the calculation immediately, or may quit the Specter system and resume computation from that point some other time.

An important issue is color processing. In our approach, global illumination computation is performed exclusively in the RGB space for reasons of efficiency. This simplification may result in some color distortions, but they are usually not significant in terms of the perceivable differences in image quality for many practical applications (for a more formal justification of validity for RGB approximation in lighting computations, see Borges [1991]).

6. RESULTS

The initial examples illustrating previous sections show that intermediate solutions of the global illumination problem can be provided rapidly by our composite algorithm. For example, see the timings for photon tracing given for the scenes in Figures 4 and 13(a). (These timings include the costs of reconstructing the lighting function using the histogram density estimation method, which is implicitly performed during the bookkeeping on the number of photons that hit mesh elements.) However, the complete render-

Table II. DEPT Computation and Illumination Map Rendering Performance as a Function of Mesh Density for the SPOT Test. Rendering speed measured for Visual Workstation 320 graphics hardware (Silicon Graphics, Inc.)

No. of mesh elements	No. of vertices	Timings [seconds]			
		Photon tracing	Filtering	TMO+rendering	Total
30,200	23,600	1.0	0.9	0.4	2.3
91,300	59,400	1.0	2.1	0.9	4.0
314,500	181,800	1.0	5.4	2.2	8.6
1,155,100	623,700	1.0	15.2	6.4	22.6

ing requires the following further steps before an image finally appears on the screen:

- (1) filtering the illumination maps;
- (2) converting luminance values (the final product in reconstructing the lighting function) into displayable RGB units using a tone mapping operator (TMO) [Tumblin and Rushmeier 1993];
- (3) displaying the illumination maps using a graphics engine.

In practice, the overhead costs required by these three steps are negligible in respect to the total costs for the global illumination solution. However, at the initial stages of computation these costs (in particular, the filtering costs) may take a significant portion of the total rendering costs. For example, for the SPOT test, a meaningful image can be generated after just one second of DEPT computation (Figure 16). However, at this stage the solution variance is high, and the size of filter support must be large to reduce noise efficiently. This incurs a higher cost due to the search for a larger number of neighboring vertices in the kd-tree data structure. The cost can be reduced significantly if the illumination value computed for a vertex v can be assigned to its neighbors, v_i , located around v within a distance shorter than the threshold value $d_f R$ (measured as a fraction of the radius R of the filter support region). Table II summarizes filtering costs for the SPOT test after one second of photon tracing $d_f = 25\%$, and for various densities of mesh tessellation. When illumination is computed for every vertex ($d_f = 0\%$), filtering requires at least three times longer for all cases shown in Table II. Still, it is hard to see significant differences in the image quality compared to images for $d_f = 25\%$.

The conversion of luminance values stored in mesh vertices into RGB values involves the TMO. The costs of TMO computation for every vertex can be significantly reduced if a look-up table with precomputed luminance ranges and the corresponding RGB values is used. The cost of displaying illumination maps depends directly on the performance of the graphics engine. Table II summarizes the total costs of filtering, tone map processing (we use TMO similar to Tumblin and Rushmeier [1993]), and rendering. These overhead costs depend roughly linearly on the number of mesh




Fig. 16. Results for the SPOT scene after one second of DEPT computation for various mesh tessellations: (a) 30,200; (b) 314,500; (c) 1,155,100 mesh elements (see Table II for additional information).

elements. Also, the photon (ray) tracing costs in the Specter system scale well with mesh complexity; i.e., the number of photons shot in a unit of time decreases only slightly with increases in the number of mesh elements. These two factors make the complexity properties of our rendering based on the DEPT solution very attractive. Figure 16 depicts the SPOT scene for selected mesh tessellations as specified in Table II. As can be seen, the quality of images obtained after one second of photon tracing is comparable for all mesh densities studied, which means that our filtering algorithm compensates well for the lower number of photons per mesh element.

The results of the experiments show that our DEPT technique depends only slightly upon *how* the scene is represented (as postulated in Veach [1997]). In the following tests, we emphasize scenes illuminated by a greater number of light sources, which are hard for the DDL technique. However, we expect meaningful and prompt response from the DEPT. Figure 17 depicts a scene, GYM HALL, illuminated by 2,600 light sources and built of over 10,200 polygons (tessellated into 30,800 mesh elements). The light sources are located along the balconies and side corridors, so that the scene is mostly lit by indirect light transferred through light shelves. The image obtained after just 20 seconds of the DEPT computation (Figure 17(a)) resembles the corresponding image well (Figure 17(b)), which required 30 and 89 minutes of DEPT and DDL computation, respectively (see Table III for more detailed statistics on DEPT and DDL processing for scenes discussed in this section). The GYM HALL scene is extremely easy to compute using our DEPT technique because changes in the lighting function are rather slow, and sharp shadow boundaries do not need to be reconstructed. This means that DDL computation and adaptive mesh subdivision do not contribute much to image quality refinement, and a meaningful image can be obtained after just a couple of seconds of DEPT computation. On the other hand, this kind of scene can be hard to compute for many global illumination methods, e.g., progressive radiosity, due to the huge number of similar light sources that illuminate the scene locally.

Table III. DDL and DEPT Statistics for GYM HALL and THEATER Scenes. *RMS* error was estimated (see Volevich et al. [1999] for derivation of *RMS* error for the DEPT algorithm) in the context of global illumination and indirect illumination only (shown in parenthesis with *). For the DDL+IDEPT solution the direct lighting photons and the *RMS* error for global illumination are ignored

Scene name	Shown in Figure	<i>RMS</i> error [%]	Hit points #		Timings [h:m:s]		
			DDEPT	IDEPT	DEPT	DDL	
GYM HALL	17a	(16.1*)	12.8	563,301	352,012	00:00:20	–
	17b	(1.6*)	–	–	33,000,488	00:30:00	01:29:00
THEATER	18b	(22.2*)	5.9	703,689	156,863	00:00:30	–
	18c	(3.5*)	–	–	6,331,505	00:20:00	01:08:00

Figure 18 shows another example of fast perceptual convergence of the intermediate solutions in terms of the perceived quality of the corresponding images. The THEATER scene is built of 17,300 polygons (tessellated into 22,300 mesh elements) and is illuminated by 581 light sources. Figure 18 (parts a and b) depicts nonfiltered and filtered illumination maps, obtained by DEPT after 30 seconds. Figure 18(b) closely resembles the corresponding image in Figure 18(c), which took 20 and 68 minutes of DEPT and DDL computation time, respectively. The final antialiased image (Figure 18(d)) was rendered using ray tracing, which took 234 minutes (the image resolution is 960×740 pixels). In ray tracing computation, direct lighting was recomputed for every image sample. This solution is typical for multipass approaches, e.g., Jensen [1996]. Indirect lighting was interpolated based on the results (stored at mesh vertices) of IDEPT computation. Since all surfaces of the scene in Figure 18 exhibit the Lambertian properties of light reflection, the illumination maps (Figure 18 (parts b and c)) are of similar quality to that obtained using ray tracing computation (Figure 18(d)). Obviously, once calculated, illumination maps make walk-throughs of adequate image quality possible almost immediately, while the ray tracing approach requires many hours of computation if the viewing parameters are changed. This example shows the advantages of high-quality view-independent solutions for rendering environments with prevailing Lambertian properties.

7. CONCLUSIONS AND FUTURE WORK

In this paper we propose a novel composite global illumination technique designed specifically to reduce the perceptual distance between the intermediate and final images as a function of computation time. The technique exploits the particular strengths of different lighting simulation algorithms in terms of the progressive refinement of image quality as perceived by the human observer. To select the most effective component algorithm at every stage and to decide upon the switchover points between the component algorithms, important characteristics of the HVS captured in the VDP are taken into account. Since VDP processing is computationally expensive, all perceptual considerations are performed exclusively at the design stage of




Fig. 17. A progressive refinement in rendering for the GYM HALL scene: (a) photon tracing with filtering (20 seconds); (b) complete illumination maps (119 minutes).




Fig. 18. Comparison of various renderings for the THEATER: (a) photon tracing without illumination map filtering (30 seconds); (b) photon tracing with filtering (30 seconds); (c) enhanced accuracy of direct illumination (88 minutes); (d) ray traced image (234 minutes).

the composite technique, and do not introduce any overhead to lighting simulation. The resulting mixture of sequentially executed algorithms

provides intermediate images of high quality rapidly, usually within minutes, or even seconds, for complex scenes.

The primary advantage of this technique is the rapid and perceptually meaningful rendering response upon user request at any stage of lighting computations. This makes our technique especially suitable for applications requiring global illumination and involving intensive interaction with the user. Our technique could also be used to improve the efficiency of other well-established global illumination solutions by providing local estimates of illumination quickly. For example, in parallel with storing photons in the photon maps [Jensen 1996] or dumping them into a file on the hard disc as in Walter [1998], the DEPT solution could be pursued to identify well-illuminated surface regions featuring smooth distribution of lighting. For such regions, further collection of photons does not contribute to any meaningful improvement in the quality of lighting reconstruction because the number of photons stored so far may be sufficient in terms of variance criteria. Thus, instead of storing further photons, only weights of the photons stored so far could be properly assigned based on the number photons that have been bucketed for every region. This should lead to a reduction in the total number of stored photons and make a more focused collection of photons possible for regions with complex lighting patterns.

In our future work we plan to focus on improving the performance and extending the functionality of the component algorithms. In particular, the concepts of photon-bucketing and filtering of reconstructed lighting must be extended to efficiently capture the directionality of lighting stored in the illumination maps. In terms of perception, it would be tempting to validate our composite algorithm settings through systematic psychophysical experiments with subjects. So far, only informal validation has been performed by the authors of this paper and by the testers of the Specter system.

APPENDIX

A. FILTERING ILLUMINATION MAPS

In this appendix we describe in more detail the selection of the support area for filtering illumination maps. The main objective is to balance random and discretization errors, taking into account the local distribution of lighting [Myszkowski 1997; Walter 1998]. We propose a novel solution for this problem with the photon-bucketing method, which we use for density estimation. Our solution is designed to reduce the perceivable noise inherent in the early stages of DEPT computation, while not introducing any bias into the final converged solution. We base our approach on strict mathematical considerations, and try to avoid heuristics whenever possible. We first formulate our efficient estimate of illumination variance (accuracy), which is computed locally and is an important component of the procedure for selecting the filter support area. The procedure itself is described in the next section, along with our heuristics for selecting valid photons for local illumination estimates.

A.1 Local Measure of Illumination Accuracy

We make the following assumptions, which are usually valid in practical cases:

- All generated photons are mutually independent.
- The probability of hitting a surface area the size of the filter support by a particular photon is sufficiently small.
- The probability of hitting the filter support area again by a photon that has already hit this area is sufficiently small.

Let us denote

S as the filter-support area,

ϕ as the luminous flux reaching S ,

V_ϕ as the variance of ϕ ,

$H(p)$ as the random “hit function” for a photon p such that

$$H(p) = \begin{cases} 1 & \text{if photon } p \text{ hits the area } S \\ 0 & \text{otherwise,} \end{cases}$$

V_H as the variance of $H(p)$,

n as the number of photons hitting S ,

N as the total number of generated photons,

Φ as the total luminous flux emitted by all light sources in the scene,

$e = \Phi/N$ as the luminous flux of a single photon.

The exact values of ϕ , V_H , V_ϕ are unknown, and we calculate their estimates $\tilde{\phi}$, \tilde{V}_H , \tilde{V}_ϕ from the values that are known exactly during the DEPT computation: n , N , S , and Φ . Based on the variance definition, V_H can be expressed as

$$V_H = EH^2 - (EH)^2 = \frac{\phi}{\Phi} - \left(\frac{\phi}{\Phi}\right)^2 \quad (1)$$

where E denotes mathematical expectation. Since luminous flux reaching S can be estimated as

$$\tilde{\phi} = ne = \frac{n}{N}\Phi \quad (2)$$

and the $(\phi/\Phi)^2$ term in (1) is negligibly small, we may estimate \tilde{V}_H as

$$\tilde{V}_H = \frac{\tilde{\phi}}{\Phi} = \frac{n}{N} \quad (3)$$

Based on the mutual independence of all photons, V_ϕ can be formulated as

$$V_\phi = NV_H e^2 \quad (4)$$

which is a direct consequence of the fact that the variance of the sum of independent random functions is equal to the sum of variances of the functions. Then, by approximating V_H in (4) using \tilde{V}_H from (3), we can express \tilde{V}_ϕ as

$$\tilde{V}_\phi = N\tilde{V}_H e^2 = N\tilde{V}_H \frac{\Phi^2}{N^2} = \frac{n\Phi^2}{N^2}$$

We use the relative mean deviation $\sigma_\%$ (expressed in percents) as a local measure of the accuracy of the illumination estimate:

$$\sigma_\% = \frac{\sqrt{\tilde{V}_\phi}}{\tilde{\phi}} 100\% = \frac{\sqrt{n}\Phi/N}{n\Phi/N} 100\% = \frac{100\%}{\sqrt{n}} \quad (5)$$

(the case $n = 0$ is processed specially; see below).

We estimate the illumination of S as $\tilde{I} = \tilde{\phi}/S$. We also estimate the confidence interval $[I_{low}, I_{high}]$ for the area illumination I such that for any value $I \in [I_{low}, I_{high}]$ the probability that exactly n photons hit S is greater than some probability threshold (e.g., 1%). Thus, we consider the illumination values $I \in [I_{low}, I_{high}]$ as practically possible, and we exclude other values as of low probability.

A.2 Selecting the Filter Support Area

The filter support area for a vertex v is located inside a sphere of radius r and center v . Whenever possible, we try to restrict the support to the semiplane surface containing v . The following simple criteria of photon selection for illumination estimate are applied:

- All photons should be located inside a sphere with center v and radius r . More precisely, a mesh vertex that buckets photons from adjacent triangles should be inside the sphere.
- The normal vectors $\tilde{\mathbf{n}}_i$ at mesh vertices v_i should be roughly aligned with the normal vector $\tilde{\mathbf{n}}$ at v within some tolerance margin.
- Whenever applicable, only photons from the topologically-connected surface that contains v are accepted. Otherwise, all photons that satisfy the above criteria are accepted. However, we apply one more heuristic to

avoid summing-up photons belonging to parallel non-coplanar surfaces, which often obscure each other (a common situation in architectural models). We use the following heuristic: the angle between \vec{n} and the vector \vec{v}_i spanned between v and v_i should be close to the straight angle within some tolerance margin.

Calculating filter support is an iterative process. We start with some minimal value of r , and then increase its size until one of the following stopping criteria is met:

- The local measure of illumination accuracy $\sigma_{\%}$ becomes smaller than some threshold value, e.g., 5%, which corresponds to the maximal allowed noise level.
- Because the above stopping criterion is not optimal for the area of high illumination gradients, e.g., shadow borders, it is better to sacrifice the random error reduction for the sake of keeping the discretization error within some reasonable limits [Myszkowski 1997]. To detect such a case we calculate the intersection of the confidence intervals $[I_{low}, I_{high}]$ for all r values considered in the previous and current iterations. If the intersection is empty, then we assume it is due to high discretization error and we stop increasing r .
- The first stopping criterion is not applicable for dark areas, which are not hit by any photon (or hit by a few number of photons). In this case we cannot estimate $\sigma_{\%}$ using (5). Instead, we use the upper bound of the confidence interval I_{high} . If the I_{high} is converted to luminance and transformed using TMO corresponds to the dark display level, we stop the process.

ACKNOWLEDGMENTS

We thank the anonymous reviewers of a previous version of this paper for their comments and suggestions. We also thank Akira Fujimoto for his encouragement, and Przemek Rokita and Bill Martens for stimulating discussions and reviewing the manuscript.

REFERENCES

- APPEL, A. 1968. Some techniques for shading machine renderings of solids. In *Proceedings of the on AFIPS Spring Joint Computer Conference*, AFIPS Press, Arlington, VA, 37–45.
- ARVO, J. R. 1986. Backward ray tracing. In *Conference Proceedings on SIGGRAPH '86* (Dallas, TX, Aug. 18–22, 1986), D. C. Evans and R. J. Athay, Eds. ACM Press, New York, NY.
- BOLIN, M. R. AND MEYER, G. W. 1998. A perceptually based adaptive sampling algorithm. In *Proceedings of the 25th Annual Conference on Computer Graphics (SIGGRAPH '98, Orlando, FL, July 19–24, 1998)*, S. Cunningham, W. Bransford, and M. F. Cohen, Eds. ACM Press, New York, NY, 299–309.
- BORGES, C. F. 1991. Trichromatic approximation for computer graphics illumination models. *SIGGRAPH Comput. Graph.* 25, 4 (July 1991), 101–104.

- DALY, S. 1993. The Visible Differences Predictor: An algorithm for the assessment of image fidelity. In *Digital Images and Human Vision*, A. B. Watson, Ed. MIT Press, Cambridge, MA, 179–206.
- DIEFENBACH, P. J. 1996. Pipeline rendering: Interaction and realism through hardware-Based multi-pass rendering. Ph.D. Dissertation. University of Pennsylvania, Philadelphia, PA.
- DRETTAKIS, G. AND SILLION, F. X. 1997. Interactive update of global illumination using a line-space hierarchy. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97, Los Angeles, CA, Aug. 3–8)*, G. S. Owen, T. Whitted, and B. Mones-Hattal, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 57–64.
- GADDIPATTI, A., MACHIRAJU, R., AND YAGEL, R. 1997. Steering image generation with wavelet based perceptual metric. In *Proceedings of the 1997 Eurographics Forum on Computer Graphics*, 241–251.
- GIBSON, S. AND HUBBOLD, R. J. 1997. Perceptually-driven radiosity. *Comput. Graph. Forum* 16, 2, 129–141.
- GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P., ARVO, J., LAFORTUNE, E., FERWERDA, J. A., WALTER, B., TRUMBORE, B., PATTANAİK, S., AND FOO, S.-C. 1997. A framework for realistic image synthesis. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97, Los Angeles, CA, Aug. 3–8)*, G. S. Owen, T. Whitted, and B. Mones-Hattal, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 477–494.
- HECKBERT, P. S. 1990. Adaptive radiosity textures for bidirectional ray tracing. *SIGGRAPH Comput. Graph.* 24, 4 (Aug. 1990), 145–154.
- HEDLEY, D., WORRALL, A., AND PADDON, D. 1997. Selective culling of discontinuity lines. In *Proceedings of the 8th Eurographics Workshop on Rendering* (June 1997), 69–80.
- JENSEN, H. W. AND CHRISTENSEN, P. H. 1998. Efficient simulation of light transport in scenes with participating media using photon maps. In *Proceedings of the 25th Annual Conference on Computer Graphics (SIGGRAPH '98, Orlando, FL, July 19–24, 1998)*, S. Cunningham, W. Bransford, and M. F. Cohen, Eds. ACM Press, New York, NY, 311–320.
- JENSEN, H. W. 1996. Global illumination using photon maps. In *Proceedings of the 1996 Eurographics Workshop on Rendering Techniques* (Porto, Portugal, June 17–19, 1996), X. Pueyo and P. Schröder, Eds. Springer-Verlag computer science series Springer-Verlag, Berlin, Germany, 21–30.
- KAJIYA, J. T. 1986. The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (Aug. 1986), 143–150.
- KELLER, A. 1997. Instant radiosity. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97, Los Angeles, CA, Aug. 3–8)*, G. S. Owen, T. Whitted, and B. Mones-Hattal, Eds. ACM Press/Addison-Wesley Publ. Co., New York, NY, 49–56.
- LI, B., MEYER, G., AND KLASSEN, R. 1998. A comparison of two image quality models. In *Human Vision and Electronic Imaging III. SPIE Vol. 3299* 98–109.
- LISCHINSKI, D., SMITS, B., AND GREENBERG, D. P. 1994. Bounds and error estimates for radiosity. In *Proceedings of the ACM Conference on Computer Graphics (SIGGRAPH '94, Orlando, FL, July 24–29, 1994)*, D. Schweitzer, A. Glassner, and M. Keeler, Eds. ACM Press, New York, NY, 67–74.
- LISCHINSKI, D., TAMPIERI, F., AND GREENBERG, D. P. 1993. Combining hierarchical radiosity and discontinuity meshing. In *Proceedings of the ACM Conference on Computer Graphics (SIGGRAPH '93, Anaheim, CA, Aug. 1–6, 1993)*, M. C. Whitton, Ed. ACM Press, New York, NY, 199–208.
- LUBIN, J. 1995. A visual discrimination model for imaging system design and development. In *Vision Models for Target Detection and Recognition* World Scientific Publishing Co., Inc., River Edge, NJ, 245–283.
- MARTENS, W. AND MYSZKOWSKI, K. 1998. Psychophysical validation of the Visible Differences Predictor for global illumination applications. In *Proceedings of the Conference on Visual-*

- ization '98. *Late Breaking Hot Topics* (Research Triangle Park, NC, Oct. 18–23, 1998), T.-M. Rhyne and R. Moorhead, Eds. IEEE Computer Society Press, Los Alamitos, CA.
- MARTIN, I., PUEYO, X., AND TOST, D. 1997. An image-space refinement criterion for linear hierarchical radiosity. In *Proceedings of the Conference on Graphics Interface '97* (Kelowna, B.C., Canada, May 21–23, 1997), W. A. Davis, M. Mantei, and R. V. Klassen, Eds. Canadian Information Processing Society, Toronto, Canada, 26–36.
- MYSZKOWSKI, K. 1997. Lighting reconstruction using fast and adaptive density estimation techniques. In *Proceedings of the 8th Eurographics Workshop on Rendering* (June 1997), 251–262.
- MYSZKOWSKI, K. 1998. the Web page with documentation of the VDP validation experiments. <http://www.u-aizu.ac.jp/labs/csel/vdp/>
- MYSZKOWSKI, K. 1998. The Visible Differences Predictor: Applications to global illumination problems. In *Proceedings of the 1998 Eurographics Workshop on Rendering Techniques*, G. Drettakis and N. Max, Eds. 223–236.
- MYSZKOWSKI, K., KHODULEV, A., AND KOPYLOV, E. A. 1997. Validating global illumination algorithms and software. In *Proceedings of the Conference on Art and Interdisciplinary Programs SIGGRAPH '97* (SIGGRAPH, Los Angeles, CA, Aug. 3–8, 1997), L. Pockock, R. Hopkins, D. Ebert, and J. Crow, Eds. ACM Press, New York, NY, 156.
- MYSZKOWSKI, K. AND KUNII, T. 1995. An efficient cluster-based hierarchical progressive radiosity algorithm. In *Proceedings of the Conference on ICSC*, Springer-Verlag, New York.
- MYSZKOWSKI, K., WOJDALA, A., AND WICYNSKI, K. 1994. Non-uniform adaptive meshing for global illumination. *Mach. Graph. Vision* 3, 4, 601–610.
- PRIKRYL, J. AND PURGATHOFER, W. 1998. State of the art in perceptually driven radiosity. In *State of the Art Reports. Eurographics '98* Eurographic Seminars, Tutorials and Perspectives in Computer Graphics.
- RAMASUBRAMANIAN, M., PATTANAIK, S., AND GREENBERG, D. 1999. A perceptually based physical error metric for realistic image synthesis. In *Proceedings of the Conference on Computer Graphics* (SIGGRAPH 99, Aug. 1999), ACM Press, New York, NY, 73–82.
- RUSHMEIER, H., WARD, G., PIATKO, C., SANDERS, P., AND RUST, B. 1995. Comparing real and synthetic images: Some ideas about metrics. In *Proceedings of the 6th Annual Eurographics Workshop on Rendering* (June 1995), 82–91.
- SHIRLEY, P. 1990. A ray tracing method for illumination calculation in diffuse-specular scenes. In *Proceedings of the Conference on Graphics Interface* (Halifax, Nova Scotia, May 14–18, 1990), S. MacKay and E. M. Kidd, Eds. Canadian Information Processing Society, Toronto, Canada, 205–212.
- SILLION, F. 1995. A unified hierarchical algorithm for global illumination with scattering volumes and object clusters. *IEEE Trans. Visual. Comput. Graph.* 1, 3, 240–254.
- SILLION, F. AND PUECH, C. 1994. *Radiosity and Global Illumination*. Morgan Kaufmann Publishers Inc., San Francisco, CA.
- SILVERMAN, B. 1985. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, Ltd., London, UK.
- SLUSALLEK, P., STAMMINGER, M., HEIDRICH, W., POPP, J.-C., AND SEIDEL, H.-P. 1998. Composite lighting simulations with lighting networks. *IEEE Comput. Graph. Appl.* 18, 2, 22–31.
- SMITS, B., ARVO, J., AND GREENBERG, D. 1994. A clustering algorithm for radiosity in complex environments. In *Proceedings of the ACM Conference on Computer Graphics* (SIGGRAPH '94, Orlando, FL, July 24–29, 1994), D. Schweitzer, A. Glassner, and M. Keeler, Eds. ACM Press, New York, NY, 435–442.
- STAMMINGER, M., SCHIRMACHER, H., SLUSALLEK, P., AND SEIDEL, H.-P. 1998. Getting rid of links in hierarchical radiosity. *Comput. Graph. Forum* 17, 3, 17–3.
- TOBLER, R., WILKIE, A., FEDA, M., AND PURGATHOFER, W. 1997. A hierarchical subdivision algorithm for stochastic radiosity methods. In *Proceedings of the 8th Eurographics Workshop on Rendering* (June 1997), 193–204.
- TUMBLIN, J. AND RUSHMEIER, H. 1993. Tone reproduction for realistic images. *IEEE Comput. Graph. Appl.* 13, 6, 42–48.
- VEACH, E. 1997. Robust Monte Carlo methods for lighting simulation. Ph.D. Dissertation. Stanford University, Stanford, CA.

- VOLEVICH, V., MYSZKOWSKI, K., KHODULEV, A., AND KOPYLOV, E. 1999. Perceptually-informed progressive global illumination solution. Tech. Rep. TR-99-1-002. Aizu University.
- WALTER, B. 1998. Density estimation techniques for global illumination. Ph.D. Dissertation. Cornell University, Ithaca, NY.
- WALTER, B., HUBBARD, P. M., SHIRLEY, P., AND GREENBERG, D. P. 1997. Global illumination using local linear density estimation. *ACM Trans. Graph.* 16, 3, 217–259.

Received: January 1999; revised: December 1999; accepted: February 2000