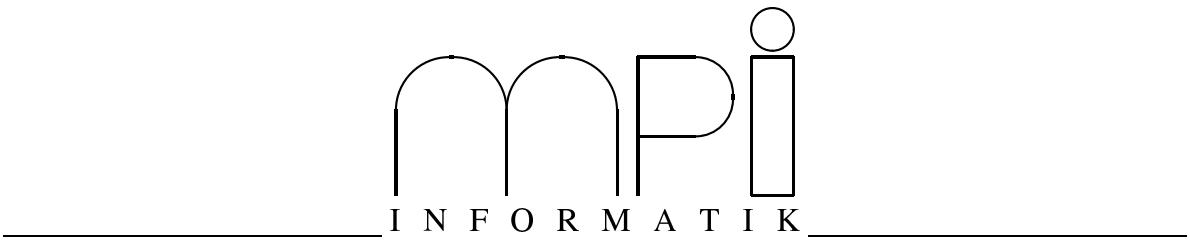

Efficient Light Transport Using Precomputed Visibility

Katja Daubert, Wolfgang Heidrich, Jan Kautz, Jean-Michel Dischler, Hans-Peter Seidel

MPI–I–2001–4-003 April 2001

FORSCHUNGSBERICHT RESEARCH REPORT

MAX-PLANCK-INSTITUT
FÜR
INFORMATIK

Stuhlsatzenhausweg 85 66123 Saarbrücken Germany

Author's Address

Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
Germany
daubert@mpi-sb.mpg.de

Keywords

Illumination, Ray Tracing, Monte Carlo Techniques, Frame Buffer Algorithms, Texture Mapping, Reflectance & Shading Models, Volume Rendering

Abstract

Visibility computations are the most time-consuming part of global illumination algorithms. The cost is amplified by the fact that quite often identical or similar information is recomputed multiple times. In particular this is the case when multiple images of the same scene are to be generated under varying lighting conditions and/or viewpoints. But even for a single image with static illumination, the computations could be accelerated by reusing visibility information for many different light paths.

In this report we describe a general method of precomputing, storing, and reusing visibility information for light transport in a number of different types of scenes. In particular, we consider general parametric surfaces, triangle meshes without a global parameterization, and participating media.

We also reorder the light transport in such a way that the visibility information is accessed in structured memory access patterns. This yields a method that is well suited for SIMD-style parallelization of the light transport, and can efficiently be implemented both in software and using graphics hardware. We finally demonstrate applications of the method to highly efficient precomputation of BRDFs, bidirectional texture functions, light fields, as well as near-interactive volume lighting.

1 Introduction

Global illumination algorithms usually spend the majority of time on visibility computations. It therefore seems natural to try and reuse visibility information computed at one point for different computations. For example, once the visibility between two points in the scene has been established, this information can be used for multiple light paths in which different amounts of energy are transported between the points. Without such a reuse, the situation is particularly unfortunate when multiple images of the same scene are to be generated under varying lighting conditions and/or viewpoints. In this case identical or at least similar visibility information needs to be computed multiple times.

There have been several approaches in the past where illumination information computed for one point in the scene has been reused for close by points. One example for diffuse scenes is the Irradiance Gradients method [28], and for scenes with specular objects there are techniques like photon maps [14] and density estimation [23]. Because these methods store illumination information (irradiance or incident radiance) at discrete points, it is not possible to reuse the information for changes of the light source. In addition, finding the desired information for one point in space requires a search through the data structure. This search can be performed in logarithmic expected time, but the resulting memory access patterns are irregular and can present a significant performance bottleneck.

We go a different way. Instead of storing and reusing illumination information, we directly reuse visibility information stored in a regular fashion that allows for constant time lookups. Our method is a generalization of the method by Heidrich et al. [13] for height fields to different kinds of geometry like general parametric surfaces, triangle meshes without a global parameterization, and volumes. Using the Method of Dependent Tests, a variant of Monte Carlo integration, we can access the visibility in a structured fashion, that allows for coherent, and therefore efficient memory access patterns in software implementations as well as for the use of graphics hardware for the light transport.

We demonstrate quality and performance of our approach by applying it to the computation of BRDFs, bidirectional texture functions (BTFs), light fields, as

well as to near-interactive volume lighting. Some examples of these applications can be seen in Figure 1.1.

Thus, the main contributions of this report are

- A generalization of the method by Heidrich et al. [13] to arbitrary parametric surfaces. This includes efficient visibility tests for indirect light and for shadows cast directly by light sources.
- A method for dealing with arbitrary triangle meshes without a global parameterization. This again includes efficient reuse of visibility information for both direct and indirect illumination.
- An efficient algorithm for computing the direct and indirect illumination in participating media. With the use of graphics hardware, this method is capable of generating the indirect illumination in volumes at almost interactive frame rates.
- The application of the mentioned methods to the efficient generation of walkthroughs, light fields, BRDFs, and bidirectional texture functions (BTFs).

The remainder of this report is organized as follows. In Section 2 we review the literature related to our approach. After briefly summarizing the work by Heidrich et al. for height fields in Section 3 we discuss our method for parametric surfaces, arbitrary triangle meshes, and volumes in Sections 4, 5, and 6. We finally present applications and results of the method in Section 7, and conclude with a discussion and conclusions in Sections 8 and 9.

Figure 1.1: *Different applications of our light transport algorithms. From left to right: non-parametric surface with shadows and indirect illumination, teapot with anisotropic BRDF, scattering in participating media, light field of a globally illuminated scene.*

2 Related Work

There have been a number of publications that describe the reuse of previously computed illumination information in global illumination algorithms. Irradiance Gradients [28] accelerate the computation of indirect light in diffuse scenes by reconstruction from irradiance samples that have been generated for other locations in close proximity of the desired surface point. The Irradiance Volume [11] represents a coarse volumetric representation of irradiance, from which the illumination at arbitrary locations can be reconstructed.

For scenes with specular objects, photons can be traced from the light sources through the scene, and stored on the objects. The incident light at arbitrary surface locations can then be reconstructed using techniques like density estimation [23] and the photon map [14].

Both the methods for diffuse and for specular surfaces store illumination information (irradiance or incident radiance) rather than visibility, and can therefore not be used to accelerate the computations in the case of changing light sources. Also the reconstruction process for any given point in the scene requires a search through the illumination data structure, which is typically the most costly part of the computation.

Other algorithms, like finite element methods for global illumination computations store visibility. In particular, the link structure in hierarchical and Wavelet Radiosity [12, 10] as well as Wavelet Radiance [3] can be interpreted as a cache for visibility information. However, since this structure only represents the most relevant parts of the visibility for a given illumination situation (BF-refinement), the information typically has to be recomputed if the illumination changes.

Precomputed visibility that is completely separated from illumination and light source positions has been studied for special cases such as height fields. Horizon maps [20, 26] represent the visibility information required for computing shadows and masking from direct light sources in hight fields and bump maps, and have also been implemented using graphics hardware [13, 24]. There have also been some solutions for shadows in more general geometry like folded cloth [27]. These approaches are based on sampled representations of the visibility. Other,

analytic representations for general scenes like the visibility skeleton [6] suffer from a combinatorial explosion of the information with the scene complexity.

The use of precomputed visibility for the purpose of computing the light transport between different parts of the scene has so far been limited to the case of height field geometry [13]. We expand this idea to more general geometry and participating media. Also using the Method of Dependent Tests [7, 25], a variant of Monte Carlo integration, we access the precomputed visibility information in a structured form. This way we achieve regular memory access patterns for efficient software rendering, as well as the possibility to compute the light transport using graphics hardware.

We demonstrate the quality and performance of our method with a number of applications ranging from near-interactive computation of indirect light in participating media over the simulation of BRDFs (similar to the methods described by Cabral et al. [2] and Westin et al. [30]) to the computation of light fields [19, 9] and BTFs [4, 5].

3 Light Transport in Height Fields

Before we introduce our light transport algorithm for computing indirect light in general polygonal scenes and participating media, we briefly review the approach presented by Heidrich et al. [13] for height fields.

The fundamental idea of this method is to calculate the visibility in a precomputation step, and to store it in a set of scattering textures S_i. In order to do this, a fixed set $D = \{d_i\}$ of sample directions on the sphere is chosen. Then a ray is shot from each grid point in the height field into each of the directions d_i and intersected with the height field geometry. A scattering texture S_i holds the intersections for all rays starting at any point in the height field in one particular direction d_i. Each of these intersections is uniquely characterized by a 2D texture coordinate.

During rendering, this information is chained together to generate a multitude of different light paths, enabling the implementation of many existing Monte Carlo algorithms. In particular, Heidrich et al. [13] apply the Method of Dependent Tests, a Monte Carlo variant in which the same sampling pattern is used for all points in the height field, to map the light transport operation to graphics hardware.

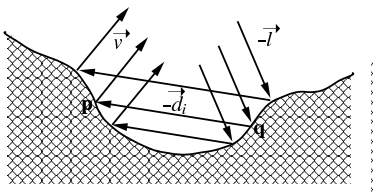

Figure 3.1: *When using the Method of Dependent tests, the light paths for computing the illumination at all points of the height field are identical.*

As an example, consider Figure 3.1, which illustrates the light paths with two reflections. Light arrives at the height field from direction $\vec{l}$, is reflected at each point in direction $-\vec{d_i} \in D$ and finally leaves the surface in the direction of the viewer $\vec{v}$. Most of the vectors are either constant across the height field, or vary

only slowly (like $\vec{l}$ and $\vec{v}$ for the case of point lights and perspective views). In fact, the only strongly varying parameters in the computation are the normal at $\mathbf{p}$, the point $\mathbf{q} := S_i[\mathbf{p}]$ visible from $\mathbf{p}$ in direction $\vec{d_i}$ and the normal at $\mathbf{q}$.

The computation is divided up into two parts, corresponding to the reflections at $\mathbf{q}$ and later at $\mathbf{p}$. First the direct illumination of the height field in viewing direction $-\vec{d_i}$ with light arriving from $\vec{l}$ is computed by a bump mapping step and stored in a texture L_d. Afterwards the second reflection is computed in a similar manner. This time the light direction is $\vec{d_i}$ and the viewing direction is $\vec{v}$, however the incoming radiance needs to be looked up in the direct illumination texture L_d. For each surface point $\mathbf{p}$ the visible point $\mathbf{q} = \mathbf{S_i}[\mathbf{p}]$ is looked up in the scattering texture corresponding to $\vec{d_i}$. $\mathbf{q}$ is then used as an index into the direct light texture L_d, yielding the light arriving at $\mathbf{p}$ in direction $-\vec{d_i}$.

This way indirect illumination of height fields is implemented as two steps of coherent light transport along one direction for all points at the same time. The mapping of this algorithm to hardware is straightforward using hardware bump mapping and dependent texture lookups (as defined for example by DirectX version 8) for the indirection.

4 General Parametric Surfaces

In the following, we will now extend the above algorithm to compute the indirect illumination in parametric surfaces. In order to precompute the scattering textures we first need to discretize the surface. We do this by using the parameterization to texture the surface. Then, given a texture of a certain resolution, each pixel in the texture can be easily mapped to a point on the surface. Similar to the height field precomputation step we now generate rays originating from each of these points in each of the global sample directions and intersect them with the surface. The intersection points can again be characterized by their parameter values, which we store as 2D floating point texture coordinates in separate textures for each sample direction. To compute indirect illumination, these scattering textures can now be used in the same way as for the height fields.

That is, given the scattering textures S_i and a per-texel normal, we first have to generate a texture-space representation of the direct illumination L_d. Then the indirect illumination can be computed completely in texture space using a sequence of table lookups for the light transport as depicted in Figure 4.1.

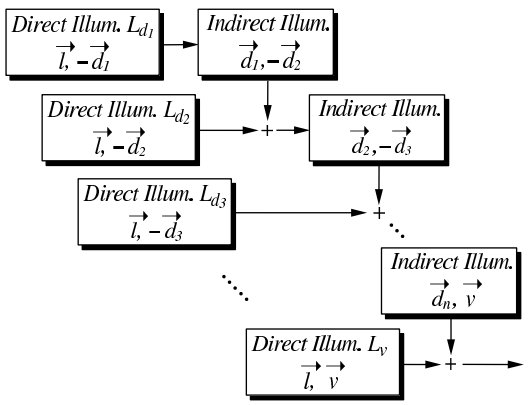

Figure 4.1: *Each box corresponds to a texture generated by a bump mapping step. Multiple scattering is computed by combining the textures as shown.*

The question then arises how to compute the direct illumination L_d, for which a shadow test is required. For the hight field case, Heidrich et al. [13] propose a hardware-friendly representation of horizon maps [20] that allows for efficient shadow tests for arbitrary light directions. Obviously, a horizon approach, no matter in which representation, will not work in the case of general parametric surfaces, since the light directions can consist of several disjoint regions. Similarly, representations like the shadow map [31] will not work, because these are valid only for specific light positions. Finally, analytic representations like the visibility skeleton [6] will be infeasible due to the combinatorial explosion in the complexity.

We therefore propose the following shadow algorithm which is similar in spirit to the horizon map in that it represents an approximation of the shadowing information for arbitrary light directions and positions, but for arbitrary geometry. In a precomputation step we partition the sphere of possible light directions into several regions by choosing some uniformly distributed directions and defining the regions around them. Then we compute the fraction of solid angle not blocked by other parts of the surface for each region. In order to do this we can reuse the visibility information already computed. For each of the directions we determine to which of the regions it belongs, and then compute the fraction of these directions that do not hit other parts of the surface. Figure 4.2 illustrates this process.

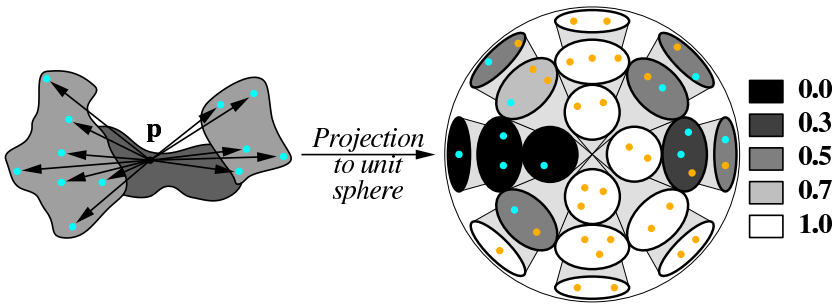

Figure 4.2: *Left: results of scattering precomputation for p (only hits are drawn). Right: Projection of hits (blue) and misses (yellow) to unit sphere. Color of shadow regions (cones) corresponds to value of fraction – dark: high number of hits, light: high number of misses.*

After having computed the fractions for all points and all shadow regions, we can store the results in a texture with one channel per region. During rendering, the shadow test for a given light direction can be performed by computing a weighted sum of the fractions for all directions to avoid quantization artifacts. For the weights we use cosine powers of the angle between the true light direction and the various reference directions.

These weights are chosen to be easily implementable using graphics hardware such as the register combiner feature on new PC boards [21]. For hardware rendering, we code the shadow information into RGBA textures, in such a way that we have one texture for four directional regions. Depending on how many simultaneous textures and combiner stages a given graphics hardware supports, we can check a number of directions at once. For every shadow texture we also need one vector of weights, and by loading two of these textures and weight vectors into a single combiner stage, we can compute the weighted sum for eight directions in one stage using two textures. Usually several passes are needed to compute the whole sum, the results of which are added using the blending operation. In our implementation we used 32 shadow directions on hardware that supports two simultaneous textures, and can therefore compute self shadowing for a given direction in four passes. The result is a texture with values ranging from zero (totally shadowed) to one (fully lit) for each point on the surface. This value can then be used to attenuate the result of a direct light computation.

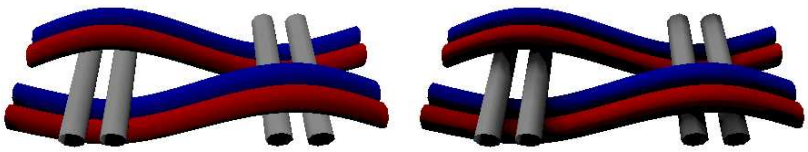

Figure 4.3: *Piecewise parametric surface without shadows (left) and with shadows computed by our shadow algorithm (right).*

Figure 4.3 shows a piecewise parametric surface without (left) and with shadows (right), computed by our algorithm. In both images the light source is located above and to the left of the object.

5 Arbitrary Triangle Meshes

One advantage of this shadow algorithm is that once we have the scattering information, the shadow computation takes place in some texture space. It is therefore well suited also for application to arbitrary triangles meshes, provided we find a way to efficiently index surface locations on these meshes.

One possibility is to reduce the problem to parametric surfaces by finding a parameterization for the triangle mesh. For example, we can use the MAPS algorithm [18] and first reduce the fine mesh to a coarse approximating triangle mesh. These coarse triangles have an inherent parameterization of their own. Then the vertices are reinserted, thereby assigning them parameter values dependent on their position on the coarse triangles. After completion the mesh consists of as many global parameterizations as there were coarse triangles.

We can then use the local parameterizations and merge them into a single large parameter space to hold the scattering information. At this point it is possible to apply the same shadowing algorithm as for parametric surfaces to generate the direct illumination map L_d for the whole mesh. From there on, we again use only texture space computations to calculate the indirect illumination using scattering textures S_i that are parameterized in the given global texture space for the mesh.

On the other hand, if we have a very fine, uniform mesh to start with, it may be sufficient to compute the illumination only at the vertices. In this case it is not necessary to generate a global parameterization and resample the surface into texture space. Instead of using a true texture for representing the samples, we use a simple lookup table, in which each vertex in the mesh is mapped to one table entry as depicted in Figure 5.1.

The shadow data structure will then contain the same information as described in Section 4, but now for every vertex in the mesh, that is, for every entry in the table.

For the scattering information, the visibility is computed in a similar fashion as for the parametric surfaces, with one important difference. Since we no longer have a parameterization for the surface, we cannot store exact intersections of rays with the surface, but rather have to quantize the intersections to entries represented

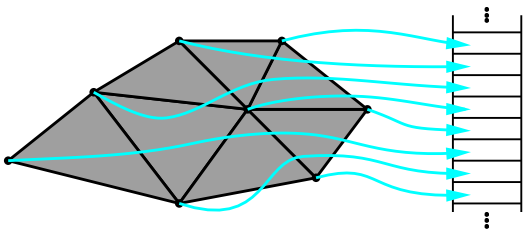

Figure 5.1: *Each vertex is mapped to an entry of a lookup table.*

in the table. This corresponds to storing the vertex closest to the ray intersection rather than the true intersection point. This will also slightly alter the direction d_i, a fact we can ignore if the mesh is fine enough. As the error depends heavily on the density of the tessellation, we can apply local refinements like subdivision or simple vertex insertion on the triangle to compensate for under-tessellated regions in the original mesh.

To use graphics hardware, we code the one-dimensional tables representing the shadow information and the scattering information into two-dimensional texture maps. Since we do not have the connectivity information that we get from a parameterization of the surface, we cannot interpolate in the illumination textures L_d during the light transport phase. However, once we have computed the indirect illumination by using table lookups as before, we have obtained per-vertex illumination that can be interpolated across the triangle mesh using Gouraud shading.

6 Volumes and Participating Media

As a final scenario for our approach to light transport we would like to consider an algorithm for computing the scattering of light in participating media. To this end, we generalize the texture based volume rendering for emission-absorption volumes described by Cabral et al. [1].

The original texture-based volume rendering algorithm works as follows. The volume is sliced with planes parallel to the image plane, yielding a stack of polygons. Each vertex receives texture coordinates corresponding to the 3D location of the vertex within the volume. Using 3D texture mapping hardware, each polygon is thus textured with the corresponding volume slice. These slices are rendered back-to-front using alpha blending to simulate the emission and absorption in the volume (see [1, 29] for details).

We will extend this basic algorithm to direct and indirect illumination in participating media. In analogy to the surface case we will use the term "direct illumination" for light emitted from a point or directional light source that is scattered exactly once before it hits the eye. "Indirect illumination" is therefore light that is reflected more than once inside the medium. For this report, we will restrict ourselves to *isotropic* media [15], that is, to media whose phase function is given as $p(\cos\alpha)$, where α is the angle between the incident and the exitant light directions.

6.1 Illumination of Participating Media

For the following discussion, we will first assume we have a volumetric representation of $L_d(\mathbf{x}, \vec{l}_\mathbf{x})$, the radiance arriving directly from the light source located in direction $\vec{l}_\mathbf{x}$ of any given point $\mathbf{x}$. We assume that $L_d(\mathbf{x}, \vec{l}_\mathbf{x})$ is given on a regular grid that is axis-aligned with the original volume data, but does not need to have the same resolution. In Section 6.2 we describe how to generate such a representation.

Given a volumetric representation of L_d, the original texture based volume rendering algorithm can be modified to include direct light. To this end we sort

the slices back-to-front, and render every slice in two different passes. In the first pass, we have to take care of the absorption in the volume. This is done in the same way the original texture based volume rendering algorithm handles absorption: the slice is texture mapped with the local density $\rho(\mathbf{x})$, which is interpolated from a volume texture. In a back to front rendering of the slices, this attenuates the previously rendered parts of the volume by the absorbed parts.

In a second pass, the slice is rendered again, this time to add in the direct illumination. This local interaction of light is described as

$$L_o(\mathbf{x}, \vec{v}) = \rho(\mathbf{x}) \frac{1}{4\pi} \int_{S^2} p(\langle \vec{v} | \vec{\omega}_i \rangle) \cdot L_d(\mathbf{x}, \vec{\omega}_i) d\omega_i. \tag{6.1}$$

For the case of directional or point light sources that we want to consider here, we get $L_d(\mathbf{x}, \vec{\omega}_i) = \delta(\vec{\omega}_i, \vec{l}) \cdot L_d(\mathbf{x}, \vec{l})$, and Equation 6.1 simplifies to

$$L_o(\mathbf{x}, \vec{v}) = \rho(\mathbf{x}) \frac{1}{4\pi} p(\langle \vec{v} | \vec{l} \rangle) \cdot L_d(\mathbf{x}, \vec{l}). \tag{6.2}$$

Thus, the slice is texture-mapped with the 3D texture $L_d(\mathbf{x}, \vec{l_{\mathbf{x}}})$, weighted by the phase function $p(\cos \alpha)$, which is either constant (corresponding to directional light and orthographic viewer), or can be interpolated as an additional 1D texture or per-vertex color. The result of this first pass is added to the result of the first pass in the framebuffer.

To compute the actual indirect illumination in the medium, we first require visibility information for the scattering process, similar to the case of geometry. Again, we randomly define a global set of directions $D = \{\vec{d_i}\}$. Then we compute one *volume* of visibility information for each direction. Each point $\mathbf{x}$ in such a scattering volume S_i represents the coordinates of another point $\mathbf{y}$ with $\mathbf{y} = \mathbf{x} + k \cdot \vec{d_i}$. The distance k is determined by stochastic ray casting process that decides where a scattering event should take place. If we decide that light is scattered, we store the coordinate of the point as the point visible from the origin of the ray, as well as the percentage of light that arrives at the scattering point from the origin of the ray without being absorbed or scattered. Otherwise, we continue tracing the ray through the volume.

Assuming that we are given a direct illumination volume L_d along with the scattering volumes S_i, the actual light transport for the indirect light is then similar to the case of parametric surfaces, except that we have to use the phase function of the volume instead of a BRDF in order to model the local interactions of light with matter. All these computations take place by simple table lookups in texture space. The right side of Figure 6.1 shows the result of such a rendering, while the left side shows the result of direct illumination only.

The difference in the light transport for participating media to the algorithms for surfaces is mostly the fact that all tables, including the direct illumination L_d

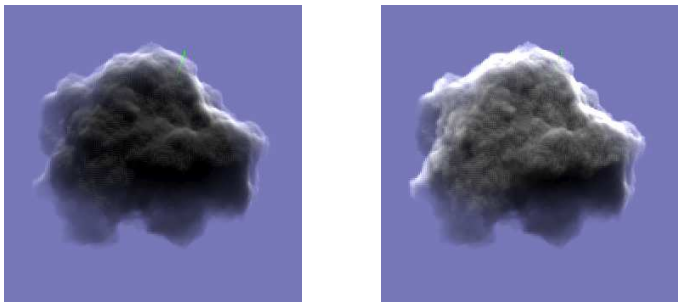

Figure 6.1: *Volume lighting. Left side: direct illumination only. Right side: direct illumination and one-bounce indirect illumination.*

and the scattering textures S_i are now 3-dimensional. As a consequence, the table lookups for the light transport also have to be 3-dimensional. Unfortunately, 3-dimensional dependent texture lookups are not available on current hardware or in specifications of future hardware like DirectX 8. Therefore it is at the moment necessary to implement the table lookups in software. In Sections 7 and 8 we discuss the performance implications of this. We will see that the method can be quite fast even in software implementations.

6.2 Direct Illumination in Participating Media

So far we have assumed that we are given a volumetric representation of the direct illumination in every point of the volume. Now we need to discuss how to arrive at such a representation.

To this end we need to compute, for every voxel in the volume, the direct light arriving from the light source, that is, the radiance leaving the light source into the direction of the voxel minus the absorbed and outscattered parts. A texture mapping approach to achieve this is to render the volume from the light source by intersecting it with planes as before, but here the planes are rendered front-to-back. Every newly rendered slice is weighted by the accumulated transparency of the previously rendered slices, which represent the portions of the volume closer to the light source. The result of the rendering of every individual slice is a sampled representation of the direct illumination arriving at this slice of the volume.

To avoid resampling and representing regions that are not covered by the actual volume, we align the volumetric representation of the direct illumination with the voxel grid of the original volume. This is achieved by aligning the slices in parallel to one of the axes of the coordinate system rather than to the image plane while rendering from the light source.

Using 3D texture mapping hardware, the direct illumination volume can there-

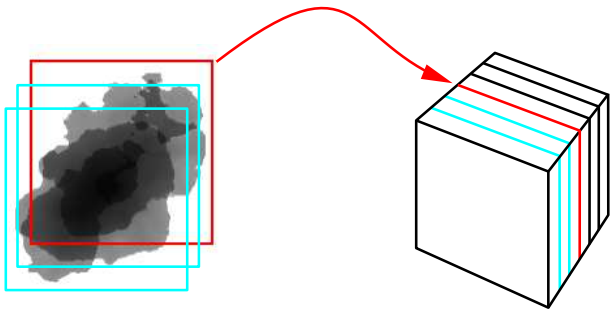

Figure 6.2: *The volume slices representing the direct illumination in the volume result from a front-to-back rendering from the light source.*

fore be generated as follows. First, we find the largest component of the light direction $\vec{l}$. Then, we slice the volume into planes perpendicular to the corresponding coordinate axis. We render these slices front-to-back, thereby multiplying them with the transparencies of the previous slices that have been accumulated in the framebuffer. After rendering every individual slice, copy the region of the framebuffer covered by it as a new slice into the 3D texture representing the direct illumination. This copying between framebuffer and texture memory is quite efficient on modern graphics hardware and does not result in the serious performance penalties of framebuffer readbacks to main memory. An example of the direct illumination generated with this method can be seen on the left side of Figure 6.1.

7 Applications and Results

For the implementation of our methods we chose two different platforms. The methods described in Sections 3, 4, and 5 were implemented on SGI workstations as well as PCs with GeForce class graphics cards, while the methods from Section 6 were only implemented on an Octane VPro, which supports 3D textures and provides enough texture RAM for our volume rendering algorithm.

As a first application of our method we consider the simulation of BRDFs. We used our methods for generating BRDF samples for several different micro geometries by first computing shadowing and indirect illumination in texture space as described above. We then rendered an orthographic image (since the light and viewing direction are parallel in BRDF computations) of the geometry from the viewing direction to handle occlusion. For the BRDF computation we assume periodic micro geometry, which means we also have to handle occlusions between several periods of the geometry. Rather than replicating the geometry to account for this kind of occlusion, we simply replicate the 2D image of one period, and composite multiple copies back to front. A BRDF sample is then obtained by averaging over the area covered by one copy of the micro geometry.

If we sample light and viewing directions over the sphere rather than the hemisphere, we can also account for transmission, yielding a bidirectional scattering distribution function, or BSDF. In this case it is usually advisable to also store a transparency value for each direction, which accounts for the Dirac peak of light passing straight through the material. To arrive at this transparency, we generate an alpha mask in the framebuffer during the rendering of the geometry. This mask represents pixels that are actually covered by the geometry. During averaging of the BRDF sample, the ratio of covered and uncovered pixels is taken into account for generating the transparency.

The computation time for one BRDF sample depends mostly on the texture size and on the number of directions used for computing the indirect light. In our case we used 128 sample directions for the integration and were able to compute samples for a 32x32 texture in 0.7 seconds, or for a 64x64 texture in 2.3 seconds. Due to the cost of traditional simulation algorithms, the usual approach of simu-

Figure 7.1: *Software rendering using BRDFs computed from the models on the right. a.) satin (cushion and tablecloth) b.) nearly transparent curtain over windows, c.) sofa, d.) red curtains.*

lating BRDFs with a virtual gonioreflectometer is to acquire only a small number of samples, and to project those into a basis like Spherical Harmonics [30] or cosine lobes [17] in order to arrive at a smooth BRDF representation. We found that this method blurred out a lot of the detail for some of our more complex microgeometry, and therefore obtained a more dense sampling with 10000 samples. For a model fitting into a 64x32 texture the simulation therefore takes slightly more than four hours. We took samples for all combinations of $100 - 200$ viewing and lighting directions distributed on the whole sphere.

The resulting tabular BRDFs were then used in a software renderer to generate the scenes in Figures 7.2 and Figure 7.1. The sofa's BRDF in Figure 7.1 was computed from the model depicted in Figure 7.1c, consisting of about 3400 vertices. The resulting BRDF is more or less diffuse with a slight color shift from green to blue for different viewing angles. The satin BRDF of the cushion and the tablecloth was computed from the model shown in Figure 7.1a. We used a specular value of $k_s = 0.3$ and an exponent of $N = 8$ for the micro BRDF that also shows up as a specular highlight in the simulated BRDF. The red curtains were made of a woven material modeled with the piecewise parametric surface shown in Figure 7.1d. We aligned the tangents of the curtain model in such a way that the grey cylinders of the micro geometry run horizontally across the curtain. The resulting BRDF is anisotropic and shows clear color shifts to red and blue, respectively, for grazing viewing angles. Also note how the BRDF becomes less transparent for these angles. This behavior is even more prominent for the BRDF generated from the micro geometry shown in Figure 7.1b, which we applied to

the almost transparent curtains in front of the windows. Note how the curtains are nearly invisible for orthogonal viewing directions and only become grey and less transparent for grazing angles. In Figure 7.2 we show the same scene for different lighting situations. In the left image the sun shines through the window, lighting the curtains from the back. Notice how the transparency varies for different viewing angles, letting light through for orthogonal views and blocking the light in the folds. In the right image no light is falling through the window. The curtains are lit mainly by a light source to the left of the viewer. Note that the transparent curtain's folds cast shadows onto the wall below the window.

Figure 7.2: *Closeup of the curtains in Figure 7.1. Left: light from the window shining through curtains. Right: light only from inside the room.*

We also used the method described by Kautz and McCool [16] to render the BRDFs in hardware. After reparameterizing the BRDF, this method approximates the 4D BRDF by a product of two 2D functions, each of which is sampled and stored in a texture map. During rendering the product is computed using the blending operation. The torus in Figure 7.3 was rendered with the same BRDF as the red curtains in Figure 7.2 (the transmission of the material is ignored here). Note the color shifts to red and blue for the grazing angles.

Model	Size	Time in sec		Memory in kB	
		Scat.	Shad.	Scat.	Shad.
c	64x64	82.46	85.47	2051	257
d	64x32	18.91	15.11	4099	513

Table 7.1: *Precomputation times for 128 scattering directions and 32 shadow regions for the models in Figure 7.1*c *and* d.

The timings for the precomputation of scattering textures and shadow fractions for the models c and d in Figure 7.1 can be taken from Table 7.1. The set of scattering directions consisted of 128 directions uniformly distributed over the

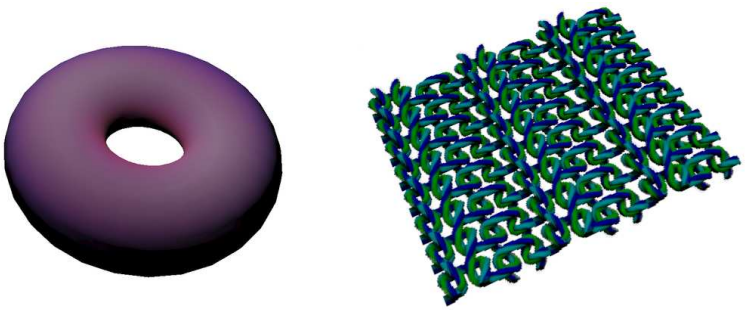

Figure 7.3: *Left: frame from hardware rendering of BRDFs as described in [16] using BRDF computed from model in Figure 7.1d. Note the color shifts at grazing angles. Right: rendering of light field computed with our methods.*

sphere. We divided the sphere into 32 regions to determine fractional visibility for the shadow computations. The column marked "Size" refers to the amount of texture space used up by the models.

By slightly modifying the algorithm for computing BRDF samples sketched above, we can also compute samples for BTFs [4, 5]. Again we compute the direct and indirect light, as well as shadows and repeat the scene to account for occlusion. However, we store a whole image per combination of one viewing and one lighting direction, instead of only an averaged BRDF sample. Since BTFs are six-dimensional functions, a faithful representation is very demanding in terms of memory and bandwidth. For instance, writing a computed image for the model in Figure 7.1b takes nearly as long as the illumination computation. This makes the generation of BTFs almost a factor of 2 slower than the BRDF simulation.

Another application for our method is to generate light fields of globally illuminated scenes. Light fields consist of a high number of images of a scene for different view points. The information in these images can be used to render the scene at interactive rates, completely independent of its complexity. In comparison to the BRDF and the BTF, the illumination remains the same over all images, only the camera location changes. By reusing visibility information in a similar manner as for the computation of BTFs we can generate a light field of the resolution 256^2x16^2 in less than two minutes. The right image in Figure 7.3 shows a frame from a light field computed from the model in Figure 7.1c.

As mentioned above, the algorithm for participating media has been implemented for SGI VPro, which has enough texture memory for our purposes, and supports 3D textures. On this platform, computing the direct illumination L_d as described in Section 6.2 and displaying it as described in Section 6.1 is interactive, at about 12 frames per second for the 256^3 cloud data in Figure 6.1. If we assume a constant phase function $p(\alpha) = c$, L_d does not need to be recomputed if the

light source remains fixed. In this case, we achieve about 20 fps. for displaying the direct illumination in the volume as described in Section 6.1.

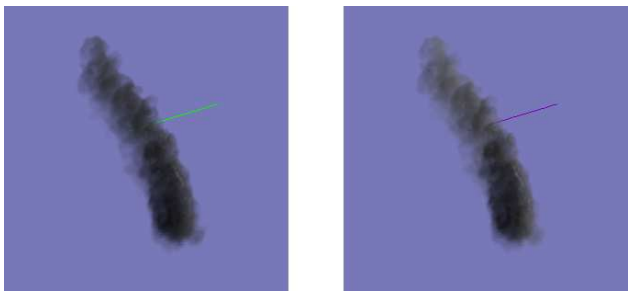

Figure 7.4: *Simulated smoke with direct (left) and indirect (right) illumination.*

For the computation of the indirect illumination it is advisable to reduce the resolution of the grid in order to save memory. For the cloud we used a grid size of 64^3, with precomputed visibility for 64 directions, yielding 64 MB of scattering information that can be generated in about 17 minutes by volume ray-casting on a 300 MHz R10k. As described in Section 6.1, we perform the light transport for the volume rendering algorithm in software. This means that for computing one-bounce indirect illumination in the participating medium, we have to traverse these 64 MB of data in software, and perform 16 million table lookups. This takes about 1.6 seconds, with about 0.7 seconds spent on cache misses (measured by performing the same computations again and again on the same data and comparing the times). We expect that these costs could be reduced by adding hardware support for this kind of 3D table lookup, because the graphics chips have a higher bandwidth to the texture RAM than the CPUs have to the main memory. In addition, graphics chips are able to employ more aggressive caching and prefetching schemes than general-purpose CPUs.

In the special case of a constant phase function we again only need to recompute the indirect illumination once for every change of the light source. In that case the indirect illumination can be merged into one volume with the direct illumination and displayed at 12 fps. Another example for illumination in participating media is depicted in Figure 7.4.

8 Discussion

As we have seen in the previous section, our approach allows for the efficient computation of illumination for different lighting and viewing situations. At the same time, the precomputation times for the visibility information are quite moderate. The efficiency of our algorithms is due to the fact that we reuse visibility information that would have had to be recomputed with other approaches.

The other reason for the performance of our method is that we use regular data structures, which allow us to reduce the light transport operator to simple, constant time table lookups. While this is already interesting for software implementations, we have argued that we can achieve additional performance benefits by mapping the light transport onto graphics hardware. This in turn is possible because similar table lookups are performed for all points on the surface or in the volume, which leads to a SIMD-style parallelism.

Since most computations take place in texture space, we do not make much use of specific hardware features like geometric transformations and scan-conversion of geometric primitives. In fact these are only used for the final display of the scene after the illumination has already been computed. The reason why graphics hardware is still interesting for this kind of algorithm, is that it is optimized for high bandwidths to texture and framebuffer RAM with specific caching schemes specially designed for this kind of lookup process.

Of course the final display of the illuminated scene could easily be replaced by other algorithms like interactive ray-tracing [22] while keeping our algorithms for light transport based on precomputed visibility. In this case we would only need a high-bandwidth architecture optimized for performing the table lookups.

Finally, our algorithm should be easy to parallelize on multiprocessor systems or clusters. One way would be to use different CPUs to compute different light paths. Each CPU then only needs to know the visibility textures corresponding to directions that comprise its paths. Another possibility could be to use the visibility data structure to cluster the geometry into mutually visible parts, similar to the form factor graph for parallel Radiosity described by Funkhouser [8].

9 Conclusion

In this report we have presented an approach for computing the indirect light in environments using precomputed visibility information. With this method, we are able to amortize the cost of generating the visibility information over many different light transports, for example to generate different light paths for a single image, or to compute many different images with varying illumination and changing camera position.

Another advantage of the method is the regular structure of the visibility information, which reduces the light transport operator to simple table lookups. This allows us to map this operation to graphics hardware, thereby exploiting the enormous memory bandwidths available in contemporary graphics systems. On the other hand, the reduction of the light transport to table lookups and the regularized memory access patterns yield a high performance even in software implementations.

We have demonstrated the feasibility, quality, and performance of the proposed method by applying it to the simulation of BRDFs, the computation of light fields and BTFs, and near-interactive illumination in participating media. We think that these applications show the possibilities of utilizing our method also for other scenarios, such as the efficient illumination of animation sequences.

Bibliography

[1] B. Cabral, N. Cam, and J. Foran. Accelerated Volume Rendering and Tomographic Reconstruction Using Texture Mapping Hardware. In *1994 Symposium on Volume Visualization*, pages 91–98, October 1994.

[2] B. Cabral, N. Max, and R. Springmeyer. Bidirectional Reflection Functions From Surface Bump Maps. In *Computer Graphics (SIGGRAPH '87 Proceedings)*, pages 273–281, July 1987.

[3] P. Christensen, E. Stollnitz, D. Salesin, and T. DeRose. Wavelet Radiance. In *Fifth Eurographics Workshop on Rendering*, pages 287–302, June 1994.

[4] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and Texture of Real World Surfaces. *ACM Transactions on Graphics*, 18(1):1–34, January 1999.

[5] J.-M. Dischler. Efficiently Rendering Macro Geometric Surface Structures with Bi-Directional Texture Functions. In *Rendering Techniques '98 (Proc. of Eurographics Workshop on Rendering)*, pages 169–180, June 1998.

[6] F. Durand, G. Drettakis, and C. Puech. The Visibility Skeleton: A Powerful and Efficient Multi-Purpose Global Visibility Tool. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 89–100, August 1997.

[7] A. Frolov and N. Chentsov. On the calculation of certain integrals dependent on a parameter by the Monte Carlo method. *Zh. Vychisl. Mat. Fiz.*, 2(4):714 – 717, 1962. (in Russian).

[8] T. Funkhouser. Coarse-Grained Parallelism for Hierarchical Radiosity Using Group Iterative Methods. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 343–352, August 1996.

[9] S. Gortler, R. Grzeszczuk, R. Szelinski, and M. Cohen. The Lumigraph. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 43–54, August 1996.

[10] S. Gortler, P. Schröder, M. Cohen, and P. Hanrahan. Wavelet Radiosity. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, pages 221–230, August 1993.

[11] G. Greger, P. Shirley, P. Hubbard, and D. Greenberg. The Irradiance Volume. *IEEE Computer Graphics & Applications*, 18(2):32–43, March–April 1998.

[12] P. Hanrahan, D. Salzman, and L. Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, pages 197–206, July 1991.

[13] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating Micro Geometry Based on Precomputed Visibility. In *Computer Graphics (SIGGRAPH '00 Proceedings)*, pages 455–464, July 2000.

[14] H. W. Jensen. Global Illumination using Photon Maps. In *Eurographics Rendering Workshop 1996*, pages 21–30. Eurographics, 1996.

[15] J. Kajiya and B. Von Herzen. Ray Tracing Volume Densities. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, pages 165–174, July 1984.

[16] J. Kautz and M. McCool. Interactive Rendering with Arbitrary BRDFs using Separable Approximations. In *Rendering Techniques '99 (Proc. of Eurographics Workshop on Rendering)*, pages 247 – 260, June 1999.

[17] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg. Non-Linear Approximation of Reflectance Functions. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 117–126, August 1997.

[18] A. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. MAPS: Multiresolution Adaptive Parameterization of Surfaces. In *Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 95–104, July 1998.

[19] M. Levoy and P. Hanrahan. Light Field Rendering. In *Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 31–42, August 1996.

[20] N. Max. Horizon mapping: shadows for bump-mapped surfaces. *The Visual Computer*, 4(2):109–117, July 1988.

[21] NVIDIA Corporation. *NVIDIA OpenGL Extension Specifications*, October 1999. Available from http://www.nvidia.com.

[22] S. Parker, W. Marin, P.-P. Sloan, B. Smits P. Shirley, and C. Hansen. Interactive Ray Tracing. In *Symposium on Interactive 3D Graphics*, pages 119–126, April 1999.

[23] P. Shirley, B. Wade, P. Hubbard, D. Zareski, B. Walter, and D. Greenberg. Global Illumination via Density Estimation. In *Eurographics Rendering Workshop 1995*. Eurographics, June 1995.

[24] P. Sloan and M. Cohen. Hardware Accelerated Horizon Mapping. In *11th Eurographics Workshop on Rendering*, pages 281–286, June 2000.

[25] I. Sobol. The Use of ω^2-Distribution for Error Estimation in the Calculation of Integrals by the Monte Carlo Method. In *U.S.S.R. Computational Mathematics and Mathematical Physics*, pages 717–723, 1962.

[26] J. Stewart. Fast Horizon Computation at all Points of a Terrain with Visibility and Shading Applications. *IEEE Transactions on Visualization and Computer Graphics*, 4(1):82–93, March 1998.

[27] J. Stewart. Computing Visibility from Folded Surfaces. *Computers and Graphics*, 23(5):693–702, October 1999.

[28] G. Ward and P. Heckbert. Irradiance Gradients. *Third Eurographics Workshop on Rendering*, pages 85–98, May 1992.

[29] R. Westermann and T. Ertl. Efficiently Using Graphics Hardware in Volume Rendering Applications. In *Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 169–178, July 1998.

[30] S. Westin, J. Arvo, and K. Torrance. Predicting Reflectance Functions From Complex Surfaces. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, pages 255–264, July 1992.

[31] L. Williams. Casting Curved Shadows on Curved Surfaces. In *Computer Graphics (SIGGRAPH '78 Proceedings)*, pages 270–274, August 1978.

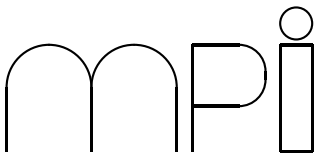

Below you find a list of the most recent technical reports of the Max-Planck-Institut für Informatik. They are available by anonymous ftp from `ftp.mpi-sb.mpg.de` under the directory `pub/papers/reports`. Most of the reports are also accessible via WWW using the URL `http://www.mpi-sb.mpg.de`. If you have any questions concerning ftp or WWW access, please contact `reports@mpi-sb.mpg.de`. Paper copies (which are not necessarily free of charge) can be ordered either by regular mail or by e-mail at the address below.

Max-Planck-Institut für Informatik
Library
attn. Anja Becker
Stuhlsatzenhausweg 85
66123 Saarbrücken
GERMANY
e-mail: `library@mpi-sb.mpg.de`

MPI-I-2001-4-002	H.P.A. Lensch, J. Kautz, M. Goesele, H. Seidel	A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web
MPI-I-2001-4-001	H.P.A. Lensch, J. Kautz, M.G. Goesele, W. Heidrich, H. Seidel	Image-Based Reconstruction of Spatially Varying Materials
MPI-I-2001-1-001	P. Krysta	Approximating Minimum Size 1,2-Connected Networks
MPI-I-2000-4-003	S.W. Choi, H. Seidel	Hyperbolic Hausdorff Distance for Medial Axis Transform
MPI-I-2000-4-002	L.P. Kobbelt, S. Bischoff, K. Kähler, R. Schneider, M. Botsch, C. Rössl, J. Vorsatz	Geometric Modeling Based on Polygonal Meshes
MPI-I-2000-4-001	J. Kautz, W. Heidrich, K. Daubert	Bump Map Shadows for OpenGL Rendering
MPI-I-2000-2-001	F. Eisenbrand	Short Vectors of Planar Lattices Via Continued Fractions
MPI-I-2000-1-005	M. Seel, K. Mehlhorn	Infimaximal Frames A Technique for Making Lines Look Like Segments
MPI-I-2000-1-004	K. Mehlhorn, S. Schirra	Generalized and improved constructive separation bound for real algebraic expressions
MPI-I-2000-1-003	P. Fatourou	Low-Contention Depth-First Scheduling of Parallel Computations with Synchronization Variables
MPI-I-2000-1-002	R. Beier, J. Sibeyn	A Powerful Heuristic for Telephone Gossiping
MPI-I-2000-1-001	E. Althaus, O. Kohlbacher, H. Lenhof, P. Müller	A branch and cut algorithm for the optimal solution of the side-chain placement problem
MPI-I-1999-4-001	J. Haber, H. Seidel	A Framework for Evaluating the Quality of Lossy Image Compression
MPI-I-1999-3-005	T.A. Henzinger, J. Raskin, P. Schobbens	Axioms for Real-Time Logics
MPI-I-1999-3-004	J. Raskin, P. Schobbens	Proving a conjecture of Andreka on temporal logic
MPI-I-1999-3-003	T.A. Henzinger, J. Raskin, P. Schobbens	Fully Decidable Logics, Automata and Classical Theories for Defining Regular Real-Time Languages
MPI-I-1999-3-002	J. Raskin, P. Schobbens	The Logic of Event Clocks
MPI-I-1999-3-001	S. Vorobyov	New Lower Bounds for the Expressiveness and the Higher-Order Matching Problem in the Simply Typed Lambda Calculus
MPI-I-1999-2-008	A. Bockmayr, F. Eisenbrand	Cutting Planes and the Elementary Closure in Fixed Dimension
MPI-I-1999-2-007	G. Delzanno, J. Raskin	Symbolic Representation of Upward-closed Sets
MPI-I-1999-2-006	A. Nonnengart	A Deductive Model Checking Approach for Hybrid Systems
MPI-I-1999-2-005	J. Wu	Symmetries in Logic Programs
MPI-I-1999-2-004	V. Cortier, H. Ganzinger, F. Jacquemard, M. Veanes	Decidable fragments of simultaneous rigid reachability

MPI-I-1999-2-003	U. Waldmann	Cancellative Superposition Decides the Theory of Divisible Torsion-Free Abelian Groups
MPI-I-1999-2-001	W. Charatonik	Automata on DAG Representations of Finite Trees
MPI-I-1999-1-007	C. Burnikel, K. Mehlhorn, M. Seel	A simple way to recognize a correct Voronoi diagram of line segments
MPI-I-1999-1-006	M. Nissen	Integration of Graph Iterators into LEDA
MPI-I-1999-1-005	J.F. Sibeyn	Ultimate Parallel List Ranking ?
MPI-I-1999-1-004	M. Nissen, K. Weihe	How generic language extensions enable "open-world" desing in Java
MPI-I-1999-1-003	P. Sanders, S. Egner, J. Korst	Fast Concurrent Access to Parallel Disks
MPI-I-1999-1-002	N.P. Boghossian, O. Kohlbacher, H.-. Lenhof	BALL: Biochemical Algorithms Library
MPI-I-1999-1-001	A. Crauser, P. Ferragina	A Theoretical and Experimental Study on the Construction of Suffix Arrays in External Memory
MPI-I-98-2-018	F. Eisenbrand	A Note on the Membership Problem for the First Elementary Closure of a Polyhedron
MPI-I-98-2-017	M. Tzakova, P. Blackburn	Hybridizing Concept Languages
MPI-I-98-2-014	Y. Gurevich, M. Veanes	Partisan Corroboration, and Shifted Pairing
MPI-I-98-2-013	H. Ganzinger, F. Jacquemard, M. Veanes	Rigid Reachability
MPI-I-98-2-012	G. Delzanno, A. Podelski	Model Checking Infinite-state Systems in CLP
MPI-I-98-2-011	A. Degtyarev, A. Voronkov	Equality Reasoning in Sequent-Based Calculi
MPI-I-98-2-010	S. Ramangalahy	Strategies for Conformance Testing
MPI-I-98-2-009	S. Vorobyov	The Undecidability of the First-Order Theories of One Step Rewriting in Linear Canonical Systems
MPI-I-98-2-008	S. Vorobyov	AE-Equational theory of context unification is Co-RE-Hard
MPI-I-98-2-007	S. Vorobyov	The Most Nonelementary Theory (A Direct Lower Bound Proof)
MPI-I-98-2-006	P. Blackburn, M. Tzakova	Hybrid Languages and Temporal Logic
MPI-I-98-2-005	M. Veanes	The Relation Between Second-Order Unification and Simultaneous Rigid *E*-Unification
MPI-I-98-2-004	S. Vorobyov	Satisfiability of Functional+Record Subtype Constraints is NP-Hard
MPI-I-98-2-003	R.A. Schmidt	E-Unification for Subsystems of S4
MPI-I-98-2-002	F. Jacquemard, C. Meyer, C. Weidenbach	Unification in Extensions of Shallow Equational Theories
MPI-I-98-1-031	G.W. Klau, P. Mutzel	Optimal Compaction of Orthogonal Grid Drawings
MPI-I-98-1-030	H. Brönniman, L. Kettner, S. Schirra, R. Veltkamp	Applications of the Generic Programming Paradigm in the Design of CGAL
MPI-I-98-1-029	P. Mutzel, R. Weiskircher	Optimizing Over All Combinatorial Embeddings of a Planar Graph
MPI-I-98-1-028	A. Crauser, K. Mehlhorn, E. Althaus, K. Brengel, T. Buchheit, J. Keller, H. Krone, O. Lambert, R. Schulte, S. Thiel, M. Westphal, R. Wirth	On the performance of LEDA-SM
MPI-I-98-1-027	C. Burnikel	Delaunay Graphs by Divide and Conquer
MPI-I-98-1-026	K. Jansen, L. Porkolab	Improved Approximation Schemes for Scheduling Unrelated Parallel Machines
MPI-I-98-1-025	K. Jansen, L. Porkolab	Linear-time Approximation Schemes for Scheduling Malleable Parallel Tasks
MPI-I-98-1-024	S. Burkhardt, A. Crauser, P. Ferragina, H. Lenhof, E. Rivals, M. Vingron	*q*-gram Based Database Searching Using a Suffix Array (QUASAR)