

Report on an Investigation of the Application and Limits of Currently Available Programme Types for Photorealistic Rendering of Light and Lighting in Architecture

**The Kimbell Art Museum as a Case Study
for Lightscape, Radiance and 3D-Studio MAX**

Based on initial work commissioned by the Licht Akademie

Kurt Altmann
Planungsbüro Altmann
mail@kurt-altmann.de

Peter Apian-Bennewitz
Technische Visualisierungen
info@pab-opto.de

April 7, 2001

revisions:

March 2000

initial report (German)

April 2001

updated and translated version (English): added 3D-Studio details and newer Radiance developments

Summary

This report is intended for planners who are interested in the topic of exact visualisation of light distribution within their professional context. Beginners will find information allowing them to estimate the effort needed and the achievable results, whereas the detailed descriptions are primarily for experts.

In the reported investigation, three criteria were particularly significant: practical relevance, photorealistic rendering and quantitative information about light.

The model building, which serves as the test object for the simulation programs, is presented in chap. 1. Chapter 2.1 describes the modelling and structuring of this object as a CAD model, which is then adopted in the simulation programs (chap. 3). The simulation process itself is described in chap. 4.

Those cases where the operation of a program was unclear, leading to problematic results, are identified with the this symbol in the margin.



Detailed descriptions, which can be skipped during initial reading, are marked with the this symbol.



If you are reading a printed copy of this report, you may find that there is now an extended version available at the following web sites:

<http://www.licht-akademie.de>

http://www.pab-opto.de/render_vergleich

You will also find the illustrations in high quality at these addresses.

Unfortunately, the reproduction quality of the illustrations in this report depends on the output device in most cases: In particular, print-outs from colour laser printers are subject to fluctuations. If you view this file as PDF on a monitor, the reproduction also cannot be standardised due to the nature of the light-emitting image and the type of monitor used.

Copyright for all rendered images is owned by ©Kurt Altmann (for the examples with Lightscape) and ©Peter Apian-Bennewitz (for the examples with Radiance). All rights reserved. Reproduction and further use require the permission of the authors.

Contents

Summary	i
1 The model	1
1.1 Object description	1
1.2 The modelled zone	4
1.3 The test points and their implementation in the model	6
2 Importing the CAD data	9
2.1 Model structuring	9
2.2 Overview	9
2.3 LIGHTSCAPE	10
2.4 3D-Studio MAX	10
2.5 RADIANCE	11
3 Paths of Simulation	13
3.1 Lightscape	13
3.1.1 Geometric refinement	13
3.1.2 Blocks	13
3.1.3 Light sources	15
3.1.4 Sky models and daylight	17
3.1.5 Definition of materials	17
3.1.6 Simulation parameters	19
3.1.7 Surface properties	19
3.1.8 Finishing touches to preparation mode	20
3.2 Radiance	20
3.2.1 File organisation	22
3.2.2 Interactive geometric representation	22
3.2.3 Definition of materials	25
3.2.4 Sky models and daylight	25
3.2.5 Colour rendering	26
3.2.6 Fundamentals and simulation parameters	26

4	The simulation process	33
4.1	Lightscape	33
4.1.1	Initialisation	33
4.1.2	Iterations and radiosity procedure	34
4.1.3	Subsequent modifications	34
4.1.4	Analysis	35
4.1.5	Ray-tracing	35
4.1.6	Batch mode	35
4.1.7	Distributed computation	37
4.2	Radiance	37
4.2.1	Starting the simulation	37
4.2.2	Analysis	38
4.2.3	Further processing for output devices	41
4.2.3.1	Output to special devices	45
4.2.4	Distributed computation	45
5	Conclusions	49
5.1	Degree of modelling detail in a scene	49
5.2	Lightscape	49
5.2.1	Manageability of Lightscape	49
5.2.2	The test points	51
5.2.2.1	Model size, subsequent model processing	51
5.2.2.2	Treatment of light	52
5.2.2.3	Light distribution via the reflector plate	52
5.2.2.4	Limits of the material definitions	53
5.2.2.5	Computation times	54
5.2.2.6	Manageability of luminaires	54
5.3	Radiance	54
5.3.1	The test points	54
5.3.1.1	Model size	54
5.3.1.2	Light distribution	54
5.3.1.3	Light distribution via the reflector plate	54
5.3.1.4	Material definitions	55
5.3.1.5	Computation times	55
5.3.1.6	Manageability of luminaires	56
5.4	Summary	56
	References	57
	About the authors	59

<i>CONTENTS</i>	v
The hardware and software	61
Index	62

Chapter 1

The model

Which model, with which degree of detail and which geometrical components, is best suited for a comparison of different lighting simulation programs?

This question was the starting point of the investigation. As the performance of Personal Computers improved, not only did the users' expectations increase, but also the features promised by software manufacturers. However, it often turned out that software, which appeared to be appropriate for smaller projects, met its limits when applied to larger projects from the planner's practice, due to the internal programming or the operating system.

At the same time, the reader of this report should be able to relate to the problem described. Exotic information obtained by using rendering farms with 100 or more machines, or high-end graphics work stations costing several ten thousand dollars, was not the subject of this investigation.

It was also assumed that a simple, rectangular, furnished office can be treated by all of the programs considered in the investigation without any difficulty. Therefore, a real object was sought from practical experience, for which the geometry could be scaled and which allowed the complexity within the model to be increased by steps. Further, the model should be able to accommodate additional "loads" in the form of other geometrical objects if necessary, without overloading the resulting scene.

Light should play a dominant role: daylight and artificial lighting should be equally represented. The choice thus fell on the Kimbell Art Museum in Fort Worth, Texas (architect: Louis I. Kahn). Real photos and plans were obtained from the sources cited in the text and referenced in the appendix.

1.1 Object description

The Kimbell Art Museum was planned and constructed in the period from 1966 to 1972. It is located about 3 kilometres from the centre of Fort Worth, surrounded by a park (fig. 1.1).

The building design is dominated by 16 parallel concrete vaults running north-south, which form the roof of the museum. Each of the vaults spans a unit of about 31.7 m x 6.7 m, and rests on 4 square corner columns. The columns have an edge length of about 0.6 m and are about 3.7 m high. Between the vaults, there are intermediate zones, about 2.4 m wide. The building is divided into three elements. The northern and southern wings, each consisting of 6 vaults, flank the central unit, which comprises 4 vaults.

Three open portica to the west constitute the main facade of the museum (fig. 1.2).

The building has two storeys. The museum administration, artists' studios, workshops and a shop are located in the lower storey. The galleries are in the upper storey (fig. 1.3).




Figure 1.1: Site plan [Bra92]



Figure 1.2: Aerial photo from the west during construction, June 1972 [Fra99]




Figure 1.3: Floor plan of the upper storey [Bra92]. The area surrounded by the red dashed line is that chosen for modelling in the following sections.

1.2 The modelled zone

The southern gallery in the upper storey was selected for modelling.

It extends over 5 vaults, consisting of an exhibition area of 43.1 m x 31.7 m unimpeded by columns, subdivided by two interior courtyards, or atria, and two emergency staircases (fig. 1.7).

The unimpeded height is app. 6.1 m to the vault apex, and about 3.1 m in the intermediate zones (fig. 1.4).




Figure 1.4: Cross-section through a vault with the adjacent intermediate zone [Bra92]. The path taken by daylight, with two reflections, is also indicated.

The larger of the two atria is constructed as a solid light well, which passes through the upper storey and supplies daylight to the studios in the lower storey.

The small atrium is completely glazed on the eastern and western sides (fig. 1.5) and provides daylight to this exhibition area from the side. Shading elements are installed along the inside of the glazing to block direct radiation. These were not considered in the modelling.

The ceiling vaults and the columns are constructed of structural concrete. The masonry walls are clad with travertine panels, inside and outside. The floor in the intermediate zones is covered with travertine pavers. There is parquet flooring under the vaults.

The exhibits are displayed on pedestals, in glass cases or are hung in front of the walls. The exhibition capacity can be increased by installing mobile room dividers, which are either free-standing or are mounted with a special fixture from the ceiling of the intermediate zones.

The special feature of Louis Kahn's design is his use of daylight for natural illumination of the exhibits. At the apex of the each cycloidal roof vault is an aperture for light, which is app. 76 cm wide and extends along the entire vault length. The light slot is closed with a curved skylight of PMMA ¹.

In order to avoid directly incident radiation, a reflector is mounted below the light opening, inside the vault. This has a fin which protrudes into the roof aperture, so that even sunlight incident under oblique angles is blocked. The two plates of the reflector consist of polished aluminium sheeting in the upper area, which is perforated in the lower area to create the impression of transparency. The opacity in the upper area prevents incident sunlight from shining directly onto the exhibits.

The reflector redirects the incident daylight onto the underside of the vault, where it is redirected and diffusely reflected into the gallery.

¹ Poly-Meta-Methyl-Acryl, also known under the trademark of "Plexiglas".



Figure 1.5: Interior view of the southern gallery, with light entering from the left through the smaller atrium [BL91]

Lighting rails are integrated into the ends of the reflector plates so that light fixtures can be positioned. Lighting rails are also embedded in the metal ceilings above the intermediate zones.

Lunettes, which are narrowest at the vault apices, are located in all of the end walls to the north and south. They are intended to emphasise the construction of the cycloidal roof vaulting. In the side walls to the east and west of the gallery, there are horizontal light slots above the travertine cladding. Here also, the intention is to emphasise the separation of the materials and the construction.

However, the openings in the walls and the ceiling do not provide any direct visual contact with the surroundings. It is only in the vicinity of the small atrium that a limited sky segment can be seen.

Fig. 1.5 illustrates the construction details described above, as well as the presentation of objects in an exhibition. Exactly this interior view is to be taken as the basis for visual comparison of results from the lighting simulation programs.

1.3 The test points and their implementation in the model

Various "special features" of the design itself and the modelling process should be helpful in localising weaknesses in the investigated programs:

As already mentioned, the most prominent characteristic of the museum is the distinctive way in which incident daylight from above is distributed in the room by the reflector surfaces and the cycloidal ceiling. Thus, the light must be redirected in the simulation at least twice, once via the reflector plate, and then via the structural concrete ceiling. The combination of the reflectivity of different material surfaces is the point which is tested here.

The complexity of this model was increased by introducing sculptures of very different types: The total number of relevant geometrical surface components increased to approximately 120,000 individual facets², a value which can well occur in practice³.

However, the model was reduced after initial test runs, as the waiting times for image composition with Lightscape and 3D-Studio MAX no longer allowed work to proceed smoothly.

After the viewpoint and the associated image segment (fig. 1.6) had been fixed, all components which were not visible and did not have any effect on the overall appearance were eliminated or represented at a very abstract level, thus creating a new model.

In addition, the model was reduced to three field sizes (fig. 1.7). The second, larger atrium was removed from the model. This simplified model still consists of app. 53,000 facets.

The wall and floor panels of travertine reflect the austere geometry of the whole design. To prevent a random textural pattern from arising in the later visualisation, the panels were individually modelled, which also ensures that the tessellation pattern from the joints is also accurately reproduced.

Particularly during the planning phase, it is advantageous if material properties can be exchanged quickly: As the individual panels were referenced during the modelling, it was possible to test whether the investigated programs accept this referencing, so that changes in individual panels have an effect on the total appearance.

Often, small image files based on pixels⁴ form the basis for the textures which are assigned to surfaces. If the texture extends over a larger area, the small texture image is usually replicated over the whole surface (fig. 1.8). In this way, the appearance from the individual image file can be retained, whereas there would be a tendency towards aliasing if the individual image were scaled up onto the larger total area.

However, it should be noted that very few image files are suitable for multiple, adjacent replication,

²The number of individual facets is not equal in all CAD and simulation programs, due to differing triangulation and polygonalisation procedures (cylinders, extruded polygons, subdivision of polygons with more than three vertices).

³Values of 300,000 were cited as being common.

⁴Pixel: smallest element of a computer image.




Figure 1.6: View of the model in AutoCAD from the perspective of fig.1.5

as the upper and lower edges, as well as the left and right edges of the image usually differ in their brightness values and/or colour variation. The "joint" in the textures then becomes clearly evident and, in the worst case, dominates the entire textured surface (fig. 1.8).

For this reason, the viewed segment was chosen to include also larger areas, so that the realism of the image could be checked.

As an alternative, textures can be generated which are defined by algorithmic procedures. We tested whether and how such textures can be created, and what their properties were with regard to visual perception and lighting technology.

In addition, initial test runs showed that undesired results were produced with Lightscape for some material properties: For instance, the lower surface of the upper reflector appeared to glow brightly, although this area should actually be opaque. To verify this problem, a second reflector was modelled, which was positioned only 0.1 mm away from the first one. This configuration was intended to test numerical rounding tolerances (see chap. 5.2.2.3).

Finally, the representation of the artificial light sources should be as close to reality as possible. The integration of luminaires should be based on the acceptance of photometric distribution data published by lighting manufacturers. Not only the ease of integrating the photometric distribution data but also the appearance of the light fixtures themselves was evaluated.

The drawing unit was chosen to be "cm". The "bounding box"⁵ for the complete model has the dimensions 2630 cm x 3170 cm x 677 cm.

⁵The "bounding box" is the smallest cube enclosing the scene.

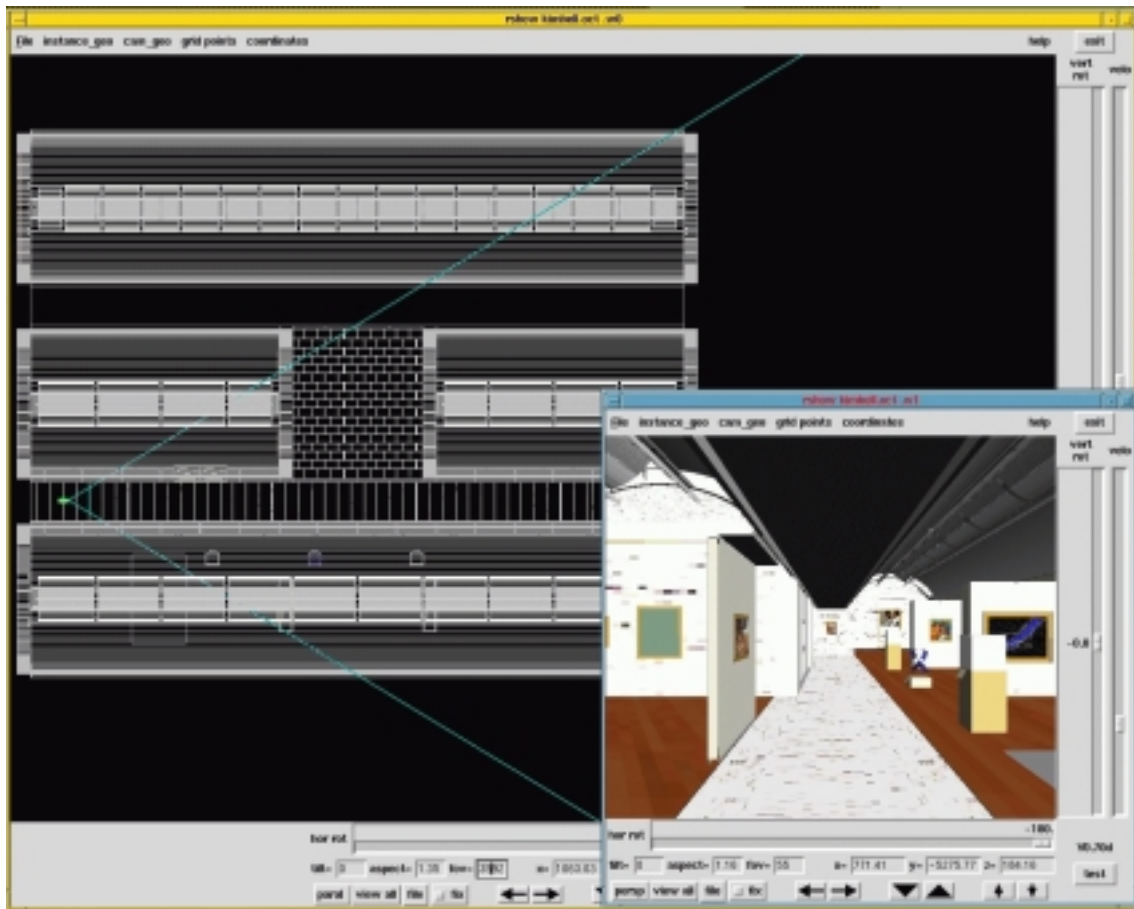


Figure 1.7: Roof view of the modified CAD model with the camera position in `rshow`.




Figure 1.8: Example for replicating an image over an area (texturing). Left: The image fills the entire area. Centre: The image is repeated twice in both directions. Right: The image is repeated 8 times in each direction. Underlying patterns in the texture image become more evident when it is repeated.

Chapter 2

Importing the CAD data

Which properties of the simulation programs should already be taken into account during modelling in the CAD program ? How easily can the CAD model be imported into the simulation programs?

2.1 Model structuring

To increase efficiency when extending, further processing or modifying CAD models, these are usually logically structured. The structuring approach is usually left by CAD systems to the user, who is thus not prevented from letting them grow at random. This leads to long computation times if the CAD data are used in simulation programs ¹. The influence of the specific needs of simulation programs on the initial structure of the work with the CAD program should not be underestimated.

During work with AutoCAD, a "reference object" (called a "block") is defined for frequently occurring elements, which subsequently only needs to be referenced. These reference objects themselves usually consist of references for the construction of a hierarchical data structure. Different colours, materials or surface properties are controlled via separate drawing layers (also called levels in other CAD programs).

Reference objects (blocks) and drawing layers thus form the basic structure for the model of the Kimbell Art Museum.

An important criterion for the exchange of CAD data is the acceptance of the surface orientation. In AutoCAD, the direction of the surface normals can already be influenced during object generation. The creator of a drawing is thus responsible for possible later corrections to the surface orientation.

In Lightscape and 3D-Studio MAX, the direction of the surface normals can easily be reversed interactively. In Radiance, this is still relatively complicated, even if it only has to be taken into account for luminous surfaces ("light", "illum") and solid glass bodies ("interface", "dielectric").

2.2 Overview

The import or translation of the CAD data into each of the investigated programs is the first criterion for their user friendliness.

Lightscape allows direct import via the standard file format from AutoCAD: DWG. With 3D-Studio MAX, the direct import of DWG files is via a plug-in. As the import functions in Lightscape and 3D-Studio MAX are not published, the user does not have much freedom to influence the import result.

¹In brief: GIGO = garbage in, garbage out.

A translator program within AutoCAD is needed to transfer CAD data from AutoCAD to a format which conforms with Radiance. For this investigation, a revised and extended version of the `torad` program was used, which was originally developed in 1993 by Georg Mischler and the Faculty for Building Science at ETH, Zurich. Programming knowledge is needed to modify the program, but this approach means that the CAD data can be processed without any restrictions.

Turning to each program individually:

2.3 LIGHTSCAPE

The imported model structure is essentially determined by the specifications from AutoCAD. All blocks (even nested ones) and all layers are correctly imported and can also be referenced according to these concepts. This does not apply to externally referenced blocks². These are ignored by the import function of Lightscape.

The colours which are defined in the model via layers are imported as material characteristics.

It is possible to open DWG files individually or to incorporate them in an existing model, which means that individual partial models generated in AutoCAD can be combined and inserted within Lightscape.

In addition, the user can decide during the import process how geometries should be treated, which were only "outlined":

Two different types of entities are meant here. The first consists of closed outlines, which do not have a "height" dimension. If the "Cap Closed Entities" option is chosen during import, the enclosed area is subdivided into triangular or quadrilateral polygons. Otherwise, these objects are ignored. The second type includes the upper and lower ends of outlines with a height attribute: A circle can thus be imported as an open pipe or a closed cylinder.

Curved outlines are converted into segmented polylines. The resolution of arcs and circles can be set.

Unfortunately, these options can be applied only globally to the file to be imported: A cylinder with an open lower end, e.g. the housing for a spotlight, cannot be imported in a single step, if it was constructed using a circle with a defined height and a circle as the cap³. If in doubt, the user should choose the closed option and edit the end surfaces subsequently by hand.

A further option during import is to generate so-called "smoothing groups": In Lightscape, all geometrical objects are constructed of flat facets⁴. A sphere, for example, is administered as a set of flat facets, although it appears round on the monitor or during rendering due to Phong-related interpolation of the surface normals. The representation of edges can be influenced by selecting this option, including entry of a threshold value for the angle formed between the surface normals of adjacent surfaces. Again, this option can only be applied globally.

2.4 3D-Studio MAX

3D-Studio MAX, with its powerful materials editor and numerous plug-ins, offers manifold possibilities for visualisation and animation. However, with the latter possibility, the program has a bias which confronts the inexperienced 3D-Studio MAX user with a plethora of functionalities which hinder rather than help basic understanding, even if they do not actually prevent desirably efficient modelling and subsequent lighting simulation.

²In principle, every drawing can be understood as a block. If a file is referenced without importing the blocks, the included blocks can only be globally scaled, rotated or translated, but their content cannot be manipulated. As the block does not become a component of the drawing, it is called an externally referenced block.

³Another solution to this problem is to model cylinders closed at one end in AutoCAD as an extruded polygon plus a closing surface.

⁴This is a direct consequence of the fundamental algorithm for the radiosity solution.

According to a distributor's information, the 3D-Studio VIZ program, where the DWG interface is stated to be more precisely adapted to AutoCAD, is better suited for importing 3D CAD data and subsequent visualisation for architectural purposes. 3D-Studio VIZ has similar basic functions to 3D-Studio MAX and is advertised as having a simple and intuitive user interface. The integrated modeller is supposed to allow easy positioning and processing of luminaire geometries and photometric data.

3D-Studio VIZ does not support any indirect calculations: Any light reflected off a surface will not illuminate other surfaces. As a workaround the users might add "invisible" light sources which than mysteriously illuminate the hithero dark areas of the room ⁵

Another way consists of a hybrid use of 3D-Studio VIZ and Lightscape, which uses the latter to solely calculate any indirect contribution. This calculations is based on material definitions in Lightscape, which are not implicitly linked to materials definitions used for final rendering 3D-Studio VIZ. Lightscape is than instructed to save only the indirect light to file. After importing this data to 3D-Studio VIZ, the final material gets defined and direct calculations are done. Since 3D-Studio VIZ material definitions are more powerful than Lightscape's this offers greater flexibility to the user.

This hybrid system is disadvantageous:

- Changing light sources or scene geometry requires redoing the procedure.
- Material definitions used in indirect calculations correlate not necessarily with the visible material definitions for direct calculations.
- Any hybrid approach implying manually following a stepwise procedure is prone to error under time pressure. (Murphy's Law).

For a quantitative and robust estimate of light distribution, apart from rendering a "nice" image, the use of 3D-Studio VIZ remains very doubtful. Failing to match the predefined aims of this study, neither 3D-Studio MAX nor 3D-Studio VIZ were investigated further.

2.5 RADIANCE

Data are exported from AutoCAD to Radiance via an auxiliary module, which runs within AutoCAD with direct access to the internal structure of the AutoCAD data. This is possible in this case, as the Radiance file format has been published completely, and its structure is relatively transparent.

This way of exporting CAD data is fruitful in general, whenever the following criteria are met:

- The CAD program allows direct access to internal data formats, and this interface is documented. The export procedure can be achieved technically with a script language such as AutoLISP, which is available to AutoCAD. An alternative is to "dock in" user programs to the CAD program (linking at runtime).
- The target format for the export is published and documented.

These criteria are met for the combination of AutoCAD and Radiance.

Further, Radiance offers the possibility to process instances ⁶ of objects which have been defined once. Thus, it is necessary only to map the block structure of AutoCAD onto Radiance in order to reap the benefits of an efficient model structure. This was a primary aim in the latest revision of the `torad` converter.

⁵These sources will not cast any shadows either. A feature we would very much like to see in the real world as well.

⁶An instance is a previously defined object, which can be referenced repeatedly. Each replication is stored only as a new transformation matrix. This means that a large number of such objects can be modelled with efficient use of memory space

Consistent definition of the geometry is more significant for export into Radiance than into Lightscape: As e.g. the inversion of surface normals after export into Radiance is not yet possible interactively, the amount of processing work increases for incorrectly defined geometry, if a Radiance material is used, for which the surface normal is relevant (e.g. for solid glass bodies and special materials).

Chapter 3

Paths of Simulation

Which data are added to the CAD geometry before simulation begins?

Elements are typically added from two libraries defined in the simulation program:

- Material libraries are collections of parameters for material surfaces, which can be referenced by a material name. The number of parameters and their physical relevance depend on the simulation program.
- Block libraries contain geometric models which are defined in detail and can be introduced into modelled scenes. These include e.g. luminaires or furnishings. These blocks can be used to replace roughly modelled dummies in the CAD geometry.

3.1 Lightscape

3.1.1 Geometric refinement

After the geometry has been imported, Lightscape is in the so-called *preparation mode* (Fig. 3.1), in which material and block libraries can be integrated:

In Lightscape, luminaires are treated like all other blocks, i.e. they can be loaded from an (external) library.

In addition to allowing the user to create individual libraries, Lightscape offers a range of libraries which can be installed. If these libraries are used consistently, even complex geometry can be handled well by the introduction of dummies.

The procedure for materials is similar.

3.1.2 Blocks

All blocks can be selected individually after import and processed further. The program changes to an editing mode, during which the main model is not displayed (fig. 3.2). However, all structuring elements (layers, materials, blocks) remain accessible. In particular, this mode is very helpful when materials are to be assigned or the surface normals need to be inverted. After terminating the editing mode, the changes are applied to all of the referenced objects.



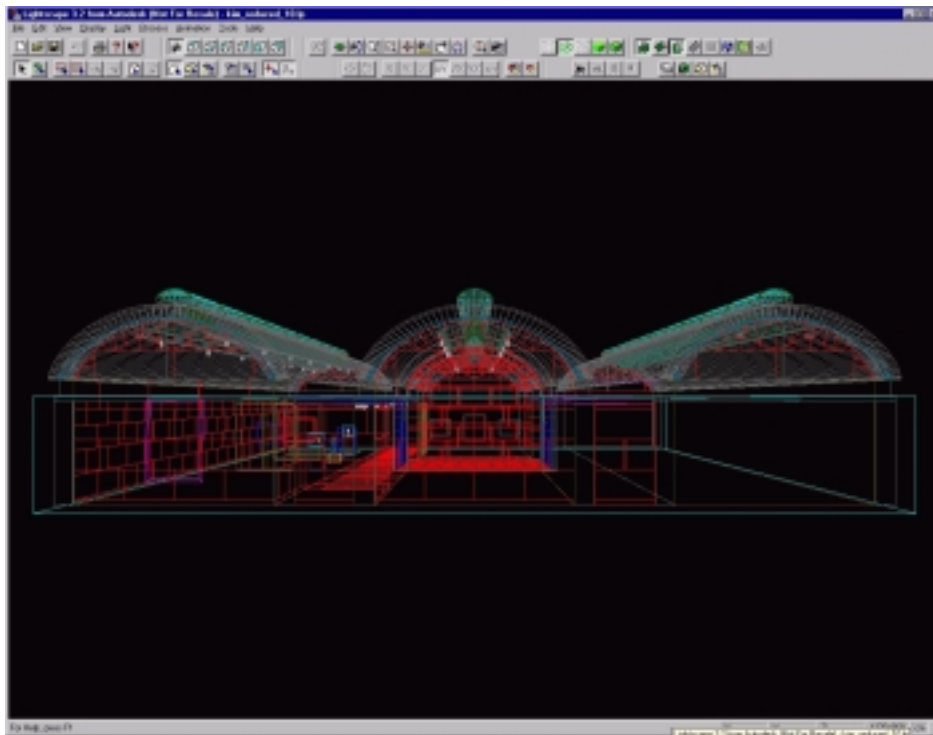


Figure 3.1: Importing the geometrical data to Lightscape

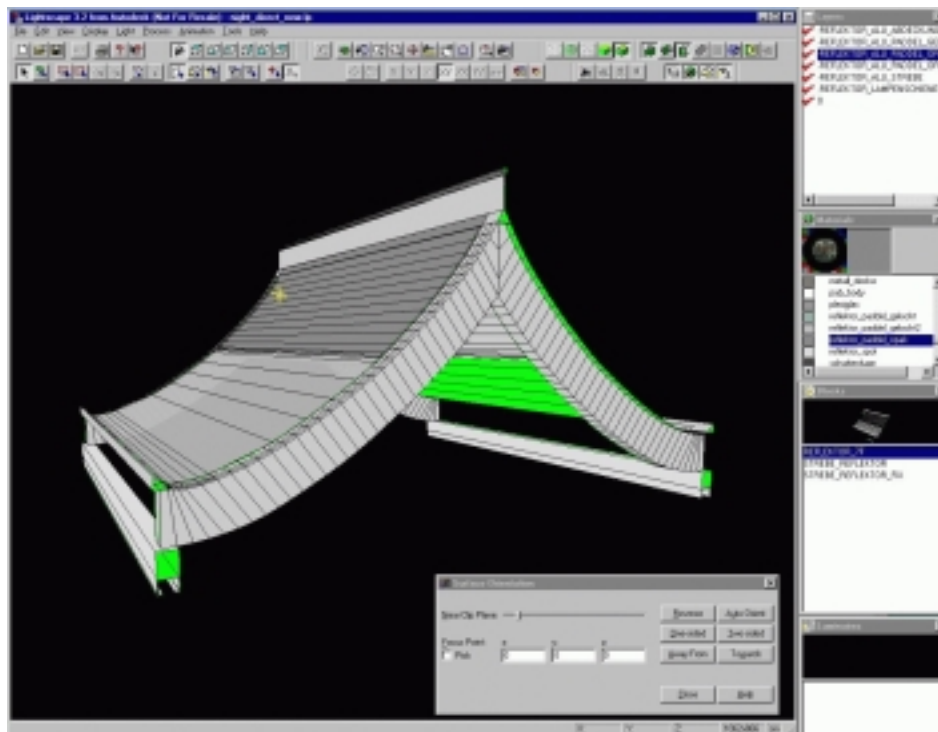


Figure 3.2: Block reference for the reflector The insert is the menu to invert the surface normals. The back surfaces are coloured green.

3.1.3 Light sources

The editing mode is similarly useful when photometric or candlepower distribution curves are to be assigned. To integrate a luminaire into the model, two approaches are possible. The user can resource libraries specific to certain manufacturers, which already contain not only the geometry and materials of the luminaire but also the appropriate light source. Alternatively, a lamp with the corresponding photometric data can be assigned to a modelled light fitting (fig. 3.3). The photometric distribution is recorded in an IES file.

The lamp itself is not visible: Lightscape makes an artificial distinction between the luminance of a light source, from which the energy is distributed onto other surfaces, and the visible luminance when this light source is viewed directly. A luminous material must be explicitly assigned to a surface in a luminaire block which is described by photometric data, e.g. the fluorescent tube of a neon light, if it is to appear as a light emitter when viewed directly (see fig. 3.4 and [Aut99, p 114]).

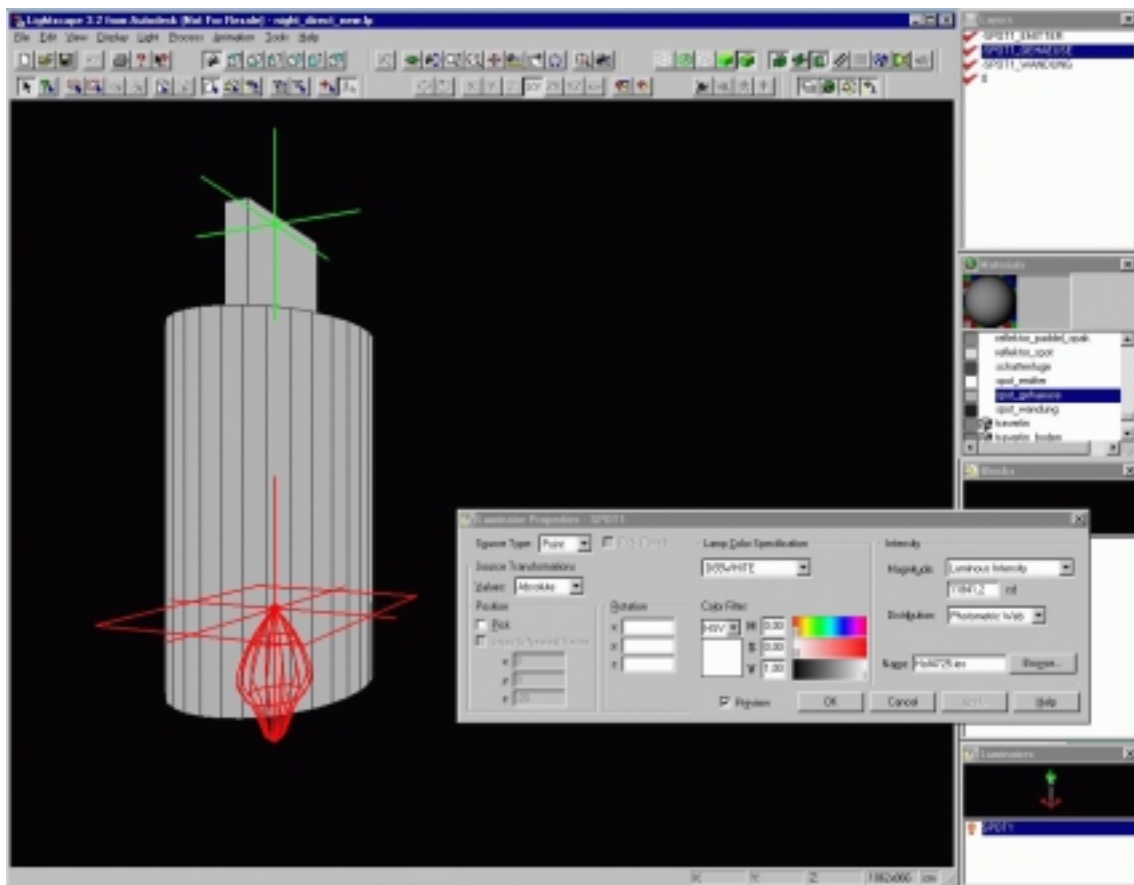


Figure 3.3: Block reference for a luminaire. The properties of the lamp can be specified via the menu.

However, a luminous material does not have any physical effect on the lighting distribution, but serves only the visual representation. This opens new possibilities for inconsistency, particularly in glare investigations, in which a fixed correlation is needed between illumination of the surroundings by a light source and its brightness when viewed directly.



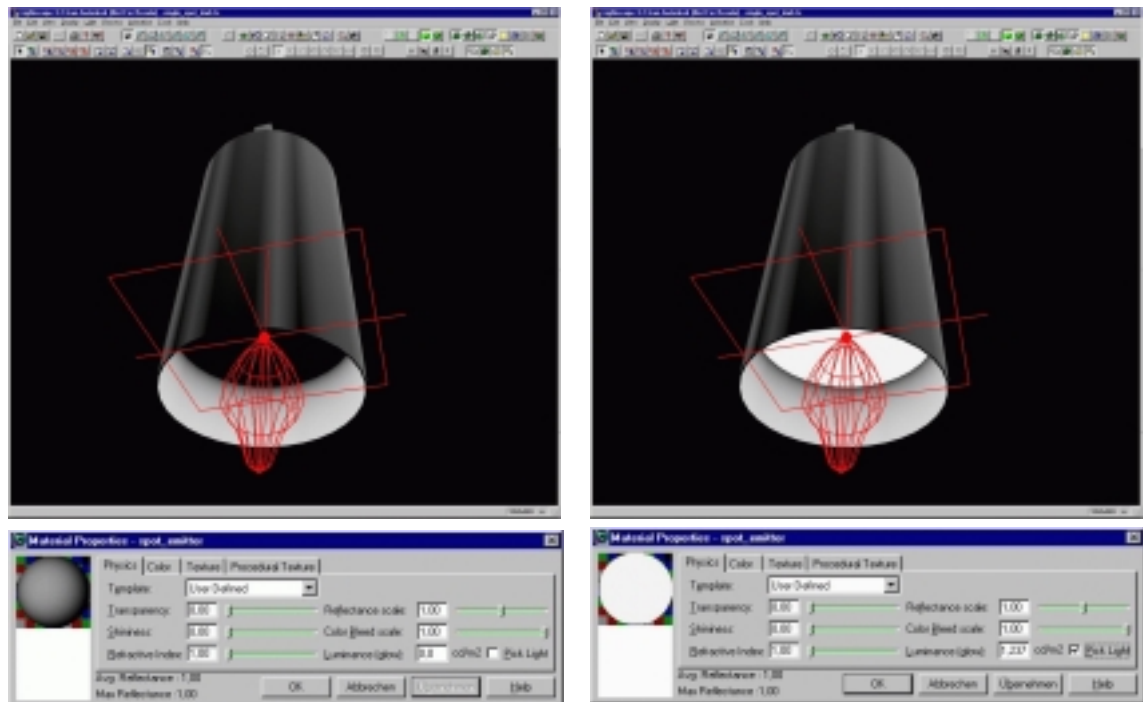


Figure 3.4: A photometric distribution (shown in red) has been assigned to both luminaires, which has its origin just below the lampshade. As usual, the light source itself is not rendered. A luminous material has been assigned to the emitter surface in the right hand example.

3.1.4 Sky models and daylight

The sun is modelled in Lightscape as a light source which emits parallel light. The sky is represented as a dome of infinite radius centred on the modelled scene. The brightness of the sky dome is correlated to the sun's position. If daylight is to be taken into account in the simulation, various parameters help to specify regional and local differences. In addition to parameters describing the sun's position, defined from the location, date and local time, both the colour of the sun and of the sky can be adjusted. The default settings agree with the IES standards¹. The degree of cloud cover can be chosen from three options: clear, half overcast, overcast. (See also the discussion of other sky models in section 3.1.4).

3.1.5 Definition of materials

Materials can be edited via four parametrisation masks.

Two preview images are displayed at the left edge of the materials editor, which show the material as applied to a spherical and a flat surface. When the parameter values are changed, the previews are updated immediately. The diameter of the preview sphere can be set to 1, 10, 100 or 1000 drawing units. This is particularly advantageous for textured materials. In addition, it is possible to switch off reflection or the coloured background to the previews. However, the previews only give a very rough impression of the later material appearance (fig. 3.5).²

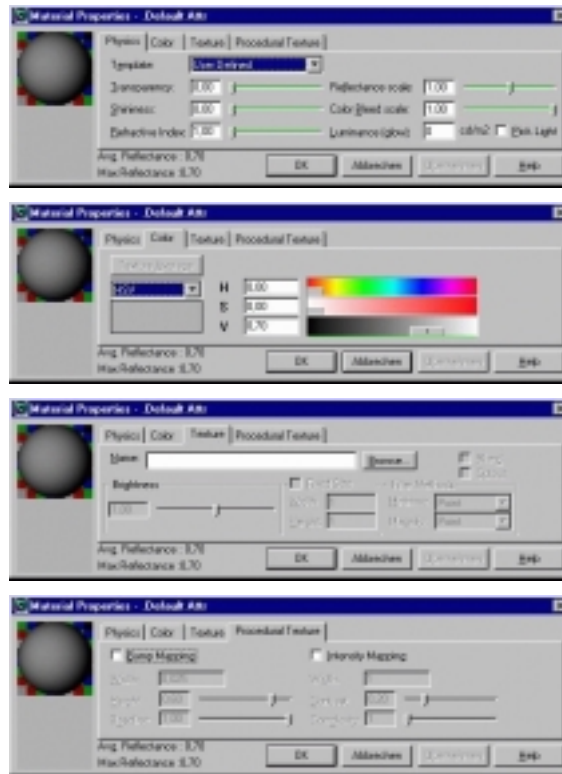


Figure 3.5: The menus to set the parameters for user-defined materials: Physical colour and texture attributes.

As an aid to defining materials, Lightscape provides several templates describing basic materials. These include default attributes for tiles, cloth, glass, paper, paint, metal, stone, water and wood. If one of

¹You can find more information on IES e.g. at <http://www.iesna.org/>

²This nomenclature seems to be doubtful: hat is the physical meaning of *shininess* ?




Figure 3.6: The restriction of allowable parameter values to reasonable ranges, depending on the material: At the top, a metal material, for which an error message appears if the reflectance is too high (centre). The bottom section shows the parameters for glass.

these templates is chosen, some parameters are automatically restricted to ranges which appear to be reasonable for the selected material (fig. 3.6).

For some materials, the restrictions are even greater. If "metal" is chosen, for example, the parameters for transparency and the refractive index are inactivated (fig. 3.6). If the user nevertheless wishes to change these parameters, this is only possible with a user-defined template, where all parameter values are allowed.

The parameters, *Avg. Reflectance* and *Max. Reflectance*, can be found at the lower edge of the materials editor. They specify the average and maximum values for the proportion of diffuse light reflected from the selected material into the rendered scene. If the material is generated using a template, the value for *Avg. Reflectance* is shown in red (fig. 3.6)), if it is beyond the range defined as reasonable.

Lightscape strongly recommends that the parameters be kept within the pre-defined limits when materials templates are used, if correct simulation results are to be obtained.

Two effects can be generated in the material by using the *Procedural Texture* input mask. If the *Bump mapping* option is chosen, the surface appears uneven. The *Intensity mapping* option creates random patches which are lighter or darker than their immediate surroundings. However, these effects only become visible once the ray tracer has been applied, and they do not have any effect on the light distribution between the surfaces.

If comprehensive libraries are not accessible, the materials editor allows simple materials to be generated quickly. In comparison to a materials editor such as that provided e.g. by 3D-Studio MAX, however, its offer is very modest. It lacks extensive possibilities for manipulation, such as the combination of materials or the selection of transparency according to illustrated examples. If the user does not wish to use standard materials, the limits for material description are quickly reached.

3.1.6 Simulation parameters

The global parameters for the accuracy and speed of the radiosity procedure, as well as the associated demand for working memory, are among the most sensitive settings for Lightscape. A "wizard" provides help³ by pre-setting defaults for all parameters, depending on the desired quality (1 = low, 5 = high) and the lighting conditions (with or without daylight, indoors and/or outdoors).

This method is strongly recommended to Lightscape novices, as the meaning and purpose of the parameters only becomes evident after some experience in using Lightscape.

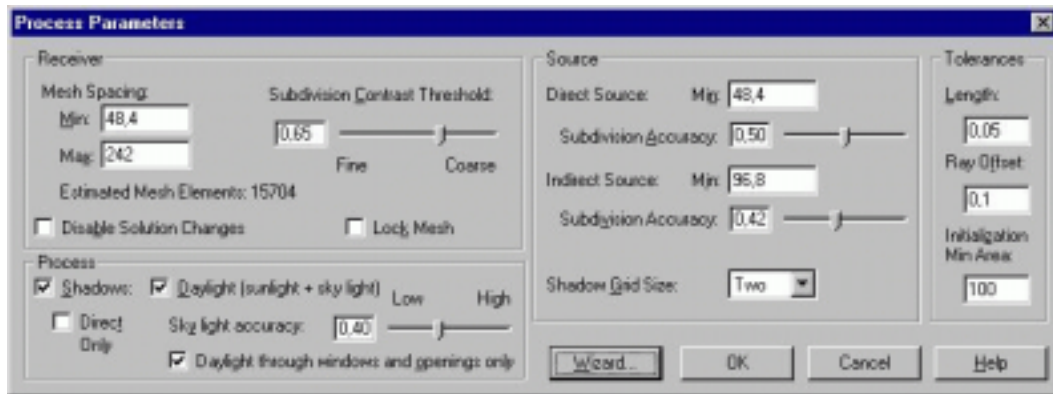


Figure 3.7: The basic parameters of the Lightscape simulation.

3.1.7 Surface properties

The response of surfaces with respect to lighting can be influenced by assigning surface attributes. The following definitions are possible:

Occluding : If this setting is activated, incident light is influenced according to the material properties. The surface casts a shadow. If this setting is deactivated, the light penetrates the surface without being affected by it.

Receiving : When this option is activated, incident light is included in the radiosity procedure for a surface. This option should be deactivated for luminous materials.

Reflecting : This option controls whether incident light on a surface is reflected or not. If the user wishes to know the illuminance for any surface within a model, the surface must be defined as "non-occluding" and "non-reflecting". The incident light on this surface can thus be monitored without the surrounding surfaces being affected.

Window : If an area is defined as a window, it is treated as a source during natural lighting computations. The emitted energy is obtained from the amount of transmitted light. A transparent material should thus be assigned to areas defined as a "window".

Opening : Openings behave similarly to windows. However, they are not part of the scene and do not receive or reflect any light. Instead, they are dummies, through which natural lighting can enter. Openings are not displayed in the rendered image.

In addition, the mesh spacing during the radiosity procedure can be defined for each surface separately. This reduces the computation time.

³The user is asked for his/her preferences regarding a subsequent procedure. Only one option can be chosen from the selection of suggestions.

3.1.8 Finishing touches to preparation mode

If all the settings have been entered, the simulation can start.

Beforehand, Lightscape requests the user to save the current status. Preparation files (with the default ending ".lp") are pure ASCII files. Their format is very similar to DXF files, but is not published. As the point co-ordinates are saved with at least 15 insignificant digits after the decimal point ⁴, the files quickly take on megabyte dimensions. The number of digits after the decimal point is independent of the drawing unit. This is even more annoying, considering that the tolerance suggested for the radiosity procedure does not correlate with the accuracy with which the 3D co-ordinates are processed.

Once the preparation mode has been concluded, the option for making geometric modifications to the model is also closed. This affects the positioning of surfaces, blocks and light sources. Only material or light properties can be modified in the following solution mode. For example, if the user later wishes to introduce a luminaire block, the associated preparation file must first be loaded and the changes made there. Afterwards, the simulation must be started again from the beginning.

3.2 Radiance

Radiance is available, free of charge, from Lawrence Berkeley National Laboratory (LBNL):

<http://radsite.lbl.gov/radiance/HOME.html>

The distributed version includes all calculation programs which are controlled by command lines. ⁵

The core of the simulation is the `rpict` program, with which the image-defining calculations of the luminance, the irradiance and the illuminance are started as background processes (fig. 3.8).

In addition, `rttrace` calculates the numerical lighting values for individual points, which e.g. can be further processed to give isolux contours.

The interactive `rview` program can be used for initial checking of parameters, aspects and lighting conditions (described in more detail in section 3.2.2).

All of these programs have the same command line options, which control the parameters for the simulation. Results, e.g. for indirect illumination, can be stored in files between consecutive starts of the same or different programs, and then used further, which drastically cuts down computation time.

We draw the reader's attention to the book [WLS98] published in 1998, which answers many questions on Radiance and lighting simulation for beginners and experts, and includes practical examples. The following discussion is a necessarily brief summary of the author's experience with Radiance over the past 10 years.

As of this writing (Spring 2001) further support and discussions lists will likely be installed at

<http://www.radiance-online.org>

For users with Radiance experience, the modular structure is very powerful, as each program has a dedicated, easily comprehensible function. However, beginners often would be glad of a graphical user interface (GUI), which would make it easier to use the program initially. Figure 3.9 shows the main menu of the `trad` program as distributed, which does not necessarily meet all expectations concerning user friendliness.

At the time of writing, several GUI's are available from third parties. The most widespread are:

⁴This corresponds to stating the distance from San Francisco to Los Angeles to an accuracy equal to the diameter of an atomic nucleus.

⁵A detailed discussion of auxiliary programs such as `mkillum` or `mktis` is beyond the scope of this report.

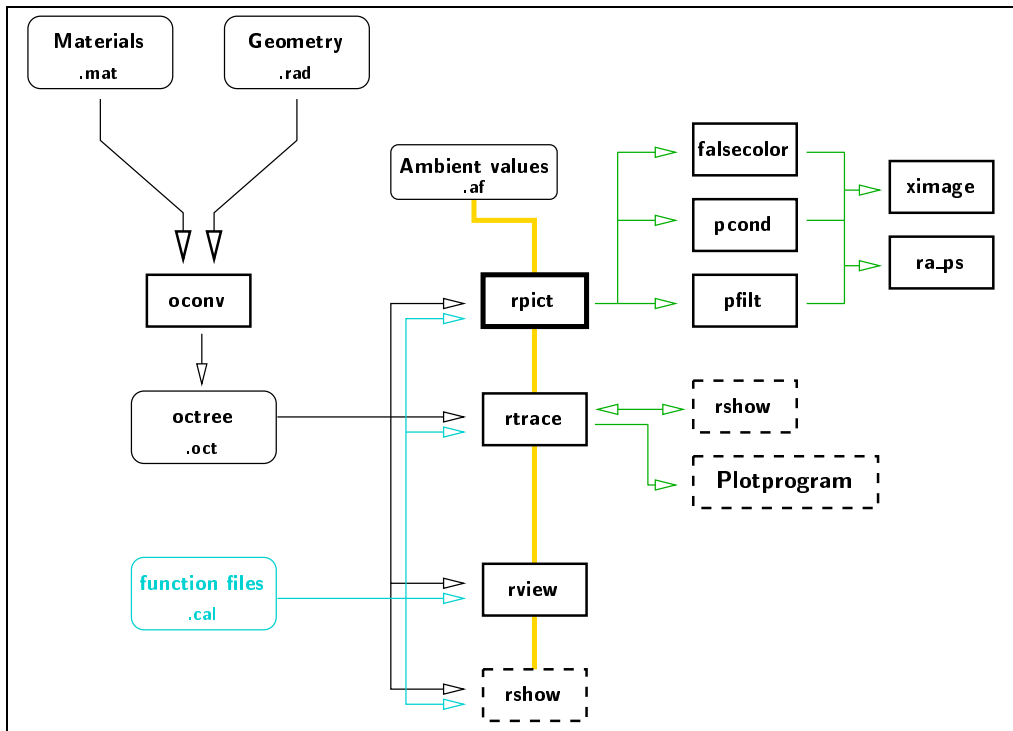


Figure 3.8: Overview of the modular programs in the Radiance system. Files are shown in the boxes with rounded corners, programs in the rectangles. Dashed outlines indicate the parts which do not belong to the core system.



Figure 3.9: Main menu of the trad program.

Rayfront : A powerful program to embed the scene geometry into the simulation with Radiance. See <http://www.schorsch.com> for further information (commercial software).

Desktop Radiance : Integrates Radiance into AutoCAD and uses its GUI (semi commercial software) <http://www.radsite.lbl.gov>.

rshow : Interactive representation of the geometry using Open-GL. The current version includes many useful practical features, but does not have a materials editor. See <http://www.pab-opto.de/progs> for details (free binaries).

A comparison of these programs is not the topic of this report, which will concentrate on the advantages and disadvantages of the Radiance core in the following. Where necessary, the advantages of the GUI's will be mentioned.

Working with command lines can be somewhat exacting for users who come from the GUI world, in which every action with a pop-up menu is ascertained with questions of the type, "Do you really not want to save?". There is not the same degree of double precautions against critical actions, such as unintentional overwriting of important files.

3.2.1 File organisation

After the CAD geometry has been exported, this is accessible in Radiance in so-called .rad files: These contain the scene geometry, structured according to the geometrical elements in Radiance: polygons, spheres, cones and cylinders. These files also include the materials, which can describe a very comprehensive range of properties, from simple standard surfaces, through anisotropic surfaces (e.g. brushed aluminium), to user-specified surfaces. Inclusion of measured values for the scattering function is possible. Finally, the .rad files also include the light sources and candlepower distribution data.

To accelerate the computations, the space in the entire scene, as originally described by the geometry in the .rad files, is recursively subdivided in three dimensions into cubes and subcubes. This is done by the `oconv` program, which is started once before the actual lighting simulation, and generates an .oct octree file from the material and geometric files.

```
oconv material.mat geometry.rad > scene.oct
```

If the scene geometry is changed, `oconv` must be started again. It is thus possible to later change the materials properties, without running the risk of introducing inconsistency.⁶ In practice, the computation time for `oconv` is negligible, being only several seconds, depending on the complexity of the scene. It is essential for speeding up the lighting simulations, as the later computation time depends sublinearly on the number of scene elements, i.e. Radiance processes complex scenes. A very long computation time for `oconv`, enormous octree files or errors during the subdivision always indicate a faulty geometry, mostly by congruent geometric elements (e.g. giving the same polygon multiple times).

For more complex scenes, in which different aspects are to be treated, it is worth distributing the various .rad files among different sub-directories, and creating a `make` file. The latter is a control file for the widespread UNIX program, `make`, which allows relationships between files to be handled elegantly and automatically.

3.2.2 Interactive geometric representation

The first steps of a Radiance simulation serve to check the object geometry and materials definitions. This is done either with the `rview` program, which is included in the standard packet, or with the

⁶The `-f` option generates *frozen octrees*, which also include the material properties. This "frozen" configuration of the scene is useful for comparison of different variations.

`rshow` program, which is available as an alternative.

`rview` operates adaptively, i.e. the image zones with high contrast are refined first. In addition, a rectangular image segment can be defined interactively (fig. 3.10). This allows a section of the complete image to be checked specifically. The computations can be made with the full range of options, i.e. the result is identical with that from `rpict`.





Figure 3.10: Interactive, adaptive image generation in `rview`. The more finely resolved segment was selected by the user with a mouse. Computation time is about 15 seconds.

`rshow` usually allows the view to be selected more quickly and simply ⁷. It uses the Open-GL standard to present the geometry of the scene (fig. 3.11). The image is similar to that from a CAD program. Interactive, real-time navigation through the scene is possible, but this advantage is inevitably gained at the cost of a crude rendering of the lighting conditions. In the language of computer graphics, this image is a Phong model, which is not physically correct. ⁸

⁷At present, `rshow` can only not be recommended if the relationship of scene complexity to computer power is unfavourable. To give an indication, the geometry presented here was handled without any difficulties with a Pentium II 400 MHz machine. According to the authors, this restriction should be relaxed in new versions.

⁸Attempts have been made to combine physically correct rendering of a scene with interactive, real-time navigation. At present, not even the power of a Cray supercomputer is adequate. This will certainly change in future with the increase in computer power.

Figure 3.11: Main menu of the `rshow` program.

3.2.3 Definition of materials

The standard Radiance system does not offer interactive selection of materials, luminaires or textures. The route is via tedious editing of ASCII text files, in which this is specified. This can well appear difficult to use for the beginner. A graphical user interface with guided selection of materials is urgently needed to reduce the initial threshold to Radiance.

Almost any conceivable properties can be assigned to the surfaces, without restrictions being imposed by the Radiance models. There are three common methods to select a surface:

- The "plastic" standard material, which describes the optical attributes of many more materials than just "plastics" (practically all non-metallic, opaque materials): It is defined by "color", "specularity" and "roughness". The three values for "color" specify the red-green-blue (RGB) colour values (see also section 3.2.5), the other two parameters control the gloss of the surface. This material model is the standard for simple models and the starting point for more sophisticated ones. It corresponds roughly to the possibilities of the Lightscape standard material.
- More complex material surfaces can be described by combining according to the building blocks principle. For example, a picture would be combined with a white area by:

```
void colorpict genpic_pic_8177
15 clip_r clip_g clip_b Trav_boden.pic picture.cal pic_v pic_u
-ry -90 -s 0.570033 -t 0 -0.5 -0.570033
0
0

genpic_pic_8177 plastic genpic_picp_8177
0
0
5 1 1 1 0 0
```

The first part defines the assignment of a picture to a surface, the second part assigns the "plastic" standard material to it. The two parts are simply linked sequentially: The result is a material with the physical properties of "plastic", with the reflectivity values as defined by a picture. This makes a complicated and laborious impression, but it contains neither extraneous nor inapplicable information. "Bump maps" and related effects can be specified similarly: waves, ripples, etc. in a material.

- The third approach to material modelling is via measurements: A sample of the materials to be modelled is measured optically, i.e. its bi-directional reflectance-transmittance function (BRDF) is determined, the data set is modelled and integrated into Radiance. Under certain circumstances, this can save a lot of time in materials modelling, as the usual procedure of trial and error is replaced by a more systematic approach. If the aim is to gain accurate information, e.g. on glare due to reflections, the effort of making measurements is certainly justified. (see e.g. <http://www.ise.fhg.de/radiance/gonio-photometer/intro.html>).

Some comments on colour rendering are given in section 3.2.5.

3.2.4 Sky models and daylight

The sky models in Radiance are more flexible than in Lightscape: The "gensky" program, which generates a CIE sky, is included in the standard Radiance packet. This distinguishes between an overcast, clear and intermediate sky, which can be further refined with a "turbulence factor".

If this is not accurate enough for some users, the "gendaylit" program can be applied to generate a sky according to the Perez model, in which the luminance distribution is specified by two parameters and corresponds better to typical measurements [Del95].

If measurements of the angular distribution of skylight are available (e.g. from so-called sky scanners), these can be incorporated into Radiance as function files.

Photos of the surroundings can be combined with the sky, and then do not only affect the appearance of the horizon when this section of the sky is viewed, but also automatically influence the computation of the lighting distribution.

3.2.5 Colour rendering

Radiance calculates in three colour channels. The exact definition of the spectral ranges does not affect the actual computations; corresponding assumptions are not made until the data are analysed with the `ximage`, `falsecolor` or `pcond` programs (see following chapter). This then also determines the spectral ranges on which the RGB colour values for the materials are based.

Two strategies are possible for colour information in Radiance:

The "conventional" approach, in which the RGB colour values are estimated for the material, and the final image is reproduced with an output device according to experience and visual impression. This is also the procedure applied with other programs. This is completely adequate for calculating the values needed for lighting analysis, as long as the spectrally averaged reflectance and transmittance values for the materials are correct.



If more exact information on colour is needed, e.g. to predict a "lighting mood", the two steps must be quantified with many more measured values: The RGB values for reflectance are then calculated from the reflectance spectrum for a specified sample. This can be measured with stationary spectrometers (e.g. Perkin-Elmer Lambda-900) or portable instruments (e.g. Minolta CM-525i). Radiance also supports the calibration of image scanners by analysing the image of a Macbeth ColorChecker reference colour card. However, the image output device is also decisive for faithful colour reproduction (see section 4.2.2). A detailed description is beyond the scope of this report.

3.2.6 Fundamentals and simulation parameters

Good selection of the parameters is needed for a simulation which is to provide the required results in the available time. Transparent and well-documented control mechanisms are open to the inveterate expert via the command line options. As an alternative, the graphical user interface offers an abstraction level, which converts general concepts (e.g. "Scene complexity: high, medium, low") into reasonable suggestions for the Radiance parameters, similarly to the Lightscape "wizard".

The Radiance system already includes the `trad` program, which allows simple administration of a simulation project and parameter selection (fig. 3.9).

The most important parameters are those for so-called ambient calculations. This is the Radiance method to calculate the interreflections between surfaces. Any form of indirect lighting is an example of this. Whereas the radiosity procedure used by Lightscape assumes the mathematical ideal of diffusely reflecting surfaces, the Radiance procedure is more general, but also more complicated.

In the following, this procedure will be described in more detail, and the most important parameters will be briefly explained: Radiance applies backward ray tracing, in which the value for a pixel is calculated by sending a view ray from the observer's eye through the imaginary image plane onto a surface of the scene (fig. 3.13) - the luminance found there determines the brightness of the pixel. The question which now needs to be answered is, "How bright is the surface?"



The brightness is obtained as the sum of three components:

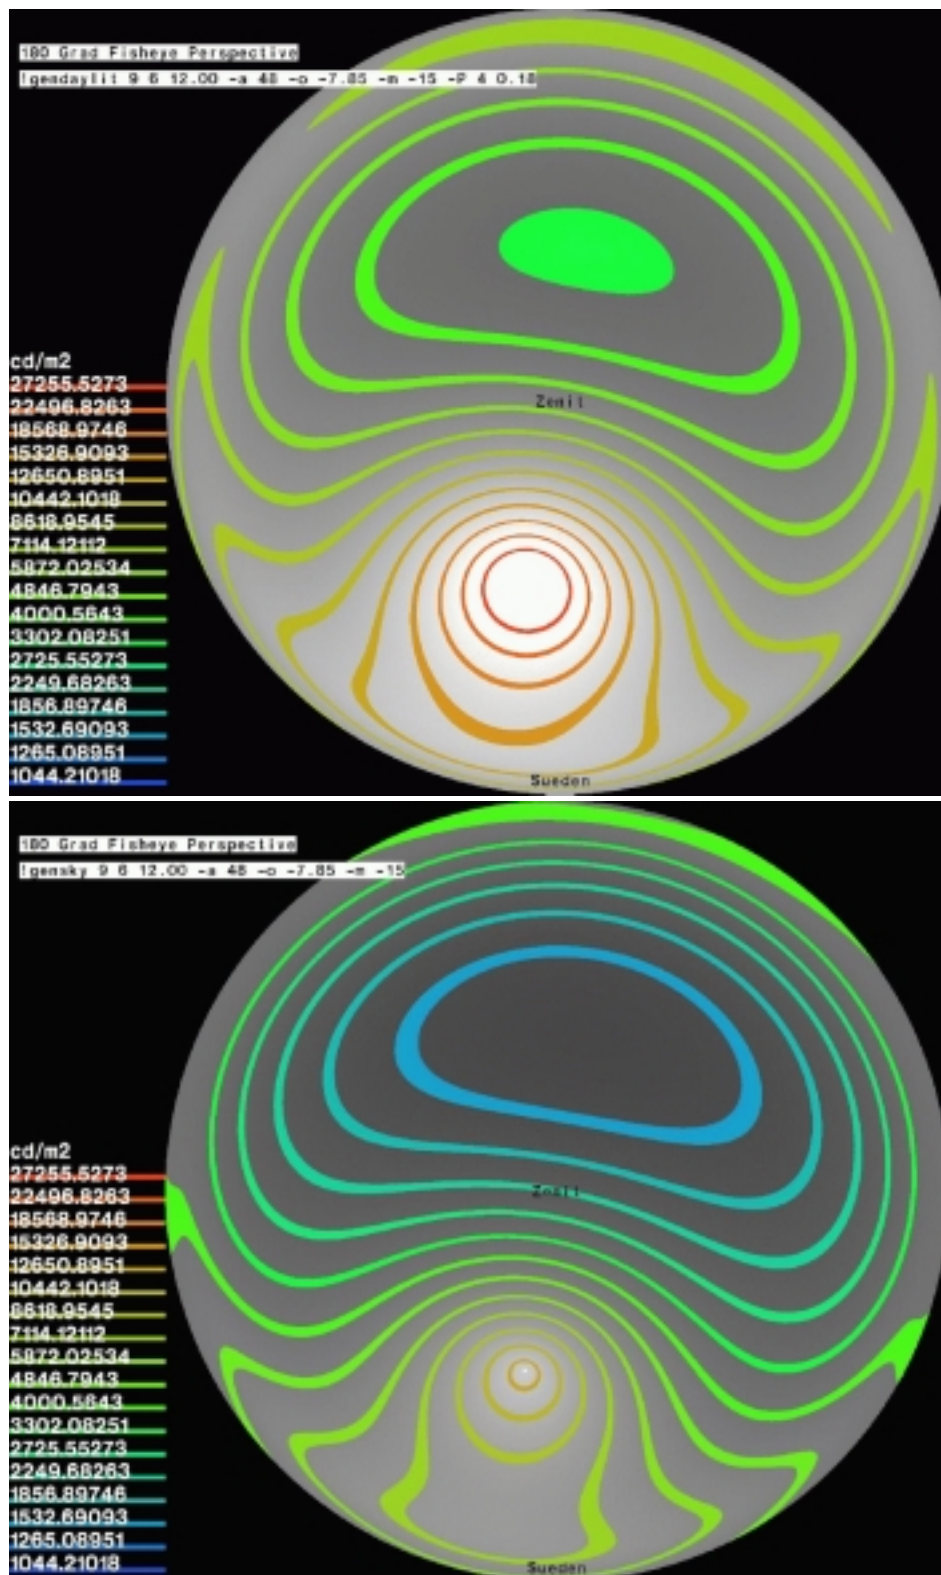


Figure 3.12: Angular distributions for skylight in Radiance: Obtained in the upper illustration from the Perez model with `gendaylit`, in the lower illustration with the CIE sky and `gensky`. The superimposed legend gives the following information: 6th September, 12 a.m. relative to 15°E (corresponding to the Central European time zone), local longitude 7.85°E, local latitude 48°N, Perez parameters $\epsilon = 4, \delta = 0.18$. For an explanation of the Perez parameters, see [Del95]. The higher level of circumsolar radiation in the Perez model should be noted.




Figure 3.13: Illustration of the *backward ray-tracing* procedure in Radiance: In the three-dimensionally modelled scene, a "synthetic" camera (represented here as a green box) generates an image in a given direction as seen from a certain view point. The image is constructed as a result of individual view rays being sent out from the camera into the scene (indicated by the red line), intersecting a surface in the scene, and the lighting conditions being calculated for the point of intersection. In doing so, the amount of light incident on this point from all directions is calculated. To this purpose, test rays are emitted recursively in all directions (yellow arrows).

direct illumination : All of the light which arrives directly from a light source on the imaged surface.

secondary direct illumination : As above, but reflected by specularly reflecting surfaces

indirect illumination : ambient calculations

The amount of light bounced onto the imaged surface from other surfaces in the scene must be determined to obtain the third component. This is calculated by sending out a bundle of radial rays (like hedgehog spines) into the hemisphere above the imaged surface. These rays then return the luminance which is incident on the imaged surface from this direction. The procedure is recursive.

Once the amount of light incident on the imaged surface is known, the sought, emitted luminance is obtained by weighting with the appropriate surface reflectivity.

The computations for this procedure can only be mastered by applying interpolation between the calculations for the "hedgehog". Otherwise, the amount of computing power needed for the recursion would be impracticable. It was Greg Ward who succeeded in implementing this concept in Radiance and thus making it practically useful [WRC88]. This interpolation confronts the user with two further parameters, for which the standard values are usually appropriate.

Thus, there are two very important and three important parameters, which play a central role in determining the image quality and the computation time: Typical values are listed in the table of fig. 3.14.

ab : Number of recursions, i.e. the number of indirect reflections (ambient bounces) considered.

ad : Number of "hedgehog" rays (ambient divisions)

as : Number of additionally emitted "hedgehog" rays (ambient supersampling)

aa : Accuracy of interpolation between the "hedgehogs" (ambient accuracy). A lower value corresponds to more computations.

ar : Maximum allowed spatial separation between two "hedgehogs", before a new one is calculated (ambient resolution)

parameter	minimal	reasonable maximum	typical	influence
ab	0	ca 12	3-7	increases computation time and quality
ad	500	ca 4000	1200	increases computation time and quality
as	0	500	100	100 - 300 is a good value
aa	0.5	0.05	0.2	if problems arise: set ad higher
ar	32	ca 1024	128	only needs setting for pathological cases

Figure 3.14: The most important parameters for ambient calculations.

Statistical fluctuations due to the Monte-Carlo approach applied are inherent to the algorithm. If the ad value is too small, this becomes apparent as "mould patches" in the image: Fig. *ambient calculations* shows images from three levels of ambient calculations and the corresponding computation times.

The cached ambient values from a calculation can be saved in a file, so that this time-consuming step only has to be taken once for different views of the same geometry. The ambient values are independent of the view point.

Intermediate saving in a file also speeds up the computation if an "overture" calculation is made before the actual simulation, using the same parameters but a lower image resolution. For instance, 200x200 pixels are adequate to "populate" the scene with ambient values and to accelerate the subsequent computations.

In contrast to Lightscape, the algorithm does not allow an input specifying how much of the incident light is already distributed, as there is no explicit definition of openings, windows or similar objects.

Also, the number of recursions (ambient bounces) specified by `ab` is not to be used incrementally, i.e. it does not make sense to try with `ab=1` and then to use the same ambient file with `ab=2`. In this case, Radiance (correctly) identifies the inconsistency between the two attempts.



If the ambient parameter values are saved in an ambient file, the geometry and materials parameters should not be changed before simulation.



Figure 3.15: Results of three different settings for ambient calculations in Radiance: From top to bottom: constant ambient value (i.e. no ambient calculation, $ab = 0$, computation time 0.135 h), one ambient bounce ($ab = 1$, $ad = 128$, 0.431 h) and seven ambient bounces ($ab = 7$, $ad = 16000$, 24 h). The "mould patches" on the ceilings have disappeared with the higher ad value.

Chapter 4

The simulation process

How is the actual simulation implemented in practice?

4.1 Lightscape

4.1.1 Initialisation

To prepare for simulation, an initialisation is made, in which the geometric data and the light sources are converted into a more efficient format for the subsequent radiosity procedure.

For large models, Lightscape advises against using layers which include a large number of surfaces. The initialisation is then supposed to proceed more quickly. Unfortunately, Lightscape's understanding of a large number is not specified.

Although the initialisation does not change the appearance or characteristics of the model, some modifications are made:

- All blocks are resolved fully. This makes sense, as blocks inserted at different positions will seldom be illuminated identically. The illuminance is assigned directly to the surfaces.
- Double-sided surfaces are converted into two separate surfaces with oppositely directed surface normals. As both surfaces are superimposed in the same plane, both are assigned the "non-reflecting" surface attribute. Lightscape advises the user to avoid such surfaces.
- The generated surfaces are collected into groups. Group members must lie in the same layer, they must have been assigned the same material and they must possess the same surface attributes. The surfaces must be coplanar and directly adjacent to group members. For Lightscape, a surface also lies in one plane if the distance of all vertices of the surface from the plane is smaller than a tolerance value. The value can be chosen without restrictions and influences not only the lighting distribution but also the "welding" of the vertices. In addition, so-called T-points¹ are eliminated when the surfaces are grouped.
- Lightscape generates an initial radiosity mesh for all surfaces, which depends on the simulation parameters. It connects the vertices of the input surfaces to form triangular and convex quadrilateral mesh elements. Each vertex is assigned an illumination value of 0.

Favourable surface forms, such as e.g. equilateral triangles or squares can be processed more effectively and do not cause visual artefacts as often.

¹T-points arise when a vertex of one surface is located on the edge of an adjacent surface. This can lead to discontinuities in the radiosity procedure. Lightscape automatically inserts a new vertex at the intersection of two edges.

Elements are removed from inactive layers.

After initialisation, the data are saved in a solution file (with the standard ending *ls*).

4.1.2 Iterations and radiosity procedure

During simulation, Lightscape uses an iterative procedure, continuously refining the distribution of light which originates from primary light sources (sun, sky, lamps), secondary light sources ("windows" and "openings") or the light reflected from surfaces.

At first, Lightscape selects the brightest light source (the one which emits the most energy) and calculates its contribution at all mesh vertices in the scene. Then, all light sources are treated in descending order of their light emission. If mesh elements obscure each other, a shadow is thrown.

Lightscape interpolates the values determined for the vertices to obtain the luminous energy incident on a mesh element. If the ratio of the lightest to the darkest vertex exceeds a certain threshold, the mesh element is subdivided into four similar elements. This process is continued iteratively until the contrast ratio is lower than the threshold or the minimum size for the mesh element has been reached. Both conditions can be pre-set with the simulation parameters.

This procedure means that the mesh elements are subdivided only at those positions relevant to the lighting analysis, such as along the edge of a shadow or the edge of a light cone.

Depending on the material properties, some of the incident light is absorbed, reflected or transmitted by the surface. All reflective surfaces are considered to be ideally diffusing (Lambertian), i.e. they reflect light identically in all directions.

After all the direct light sources have been treated, Lightscape determines the mesh element which reflects or, in the case of a window, transmits the largest amount of luminous energy. These now distribute energy in turn as secondary light sources. Then the distribution from all other mesh elements is calculated in descending order of luminous energy.

This process is continued until a state of equilibrium is reached between incident and absorbed energy, i.e. the simulation has reached a state of convergence.

Part of the total energy in the scene is distributed in each of these iterative processing steps. As the number of iterations increases, the additional amount of energy distributed in a new iteration step decreases (convergence of the solution). In practice, usually no further change in appearance can be perceived after only a fraction of the total iterations has been completed.

During the simulation, only that amount of light which has already been distributed is represented on the monitor, so the view of a scene is originally black when the initialisation process has been completed. However, as iteration proceeds, more and more light becomes visible. Progress in the light distribution process is indicated by a number giving the percentage of originally available light which has already been distributed.

The iteration process can be interrupted and recorded at any iteration step, and then continued later with the next iteration step.

4.1.3 Subsequent modifications

If the iteration process has been interrupted, material properties and the photometric properties of luminaires can be modified. Changes in the materials properties are adapted visibly immediately. However, if light had already been reflected from a subsequently modified surface, or if luminaire properties were modified, the program requires several further iterations to respond completely to the modifications. Alternatively, the whole simulation can be repeated, of course. If significant modifications were made, Lightscape advises taking the latter approach.

4.1.4 Analysis

After all iterations have been completed, or if the iteration process is interrupted, the distribution of diffuse light in the defined scene is available as the result (fig. 4.1).

One advantage of the radiosity procedure is that the results for the light distribution are independent of the observer's viewpoint. The light distribution can be examined in more detail in regions of specific interest, without the need to restart the simulation. This feature distinguishes the radiosity procedure fundamentally from the ray-tracing procedure, in which a new simulation must be started for each viewpoint selected.

The distribution of the luminance or the illuminance can be presented graphically with false colours or a grey scale to allow analysis of the lighting conditions. Linear scaling of the false colours is the default setting. However, the colour conversion can also be scaled logarithmically for investigation of areas which are only very weakly illuminated in comparison to the maximum illumination of the scene.

The illuminance and luminance values can be obtained for every point in the scene which is accessible per mouse click.

A "measurement grid" with an unrestricted choice of grid spacing can be superimposed on any desired surface of the scene. Either the illuminance or the luminance values can be displayed at the grid intersection points. Unfortunately, the font size of the displayed values cannot be edited. In addition, the number of digits displayed before and after the decimal point can be selected only globally. This results in abstrusely presented values at times. If 8 significant digits have been selected, the values 1234.5678 lux and 12.3456 lux are presented as "1234.5678" and "12.34560" respectively. If 2 significant digits are selected, the values are shown as "1.2e+003" and "12" respectively. As the displayed values overlap if the grid spacing is small, the whole display becomes unintelligible and does not have any real value. A means of recording the determined values in a file was not found.

The possibilities for analysis are restricted entirely to the results of the radiosity procedure. However, the representation e.g. of the directly incident sunlight can be suppressed in order to accelerate the simulation. Then, the sun is used only to calculate the indirect illumination. This option is applied primarily when the user subsequently wishes to generate an image with the ray tracer. Among other effects, shadows are rendered more accurately with the ray-tracing procedure.

If direct illumination had been included in the rendering from the radiosity procedure, this direct contribution and its effects would first have to be removed again, before the ray tracer is activated, as otherwise the direct illumination would be counted twice in the result. The simulation would not provide correct results. If the direct illumination is included in the radiosity procedure, the computation time becomes significantly longer.

4.1.5 Ray-tracing

Ray-tracing in Lightscape is applied only after the actual lighting distribution has been calculated with the radiosity procedure, and serves more the visible improvement of the image rather than the physically correct calculation of lighting conditions. The main applications are to render reflections of light sources from materials and textures, and reflections of other objects (such as on the glass cases in the modelled exhibition gallery).

The user can decide whether each light source (lamps, sun, diffuse light from the sky) is to be included in the ray-tracing procedure or not.

4.1.6 Batch mode

The radiosity procedure and ray tracing can also be implemented via commands. This allows the Lightscape user to take advantage of scripts when dealing with complex problems. One effect is shorter computation times, as each iteration step no longer needs to be generated on the monitor. It



Figure 4.1: Results of Lightscape calculations for two different positions of the sun.




Figure 4.2: Results of Lightscape calculations in a false-colour representation

is also possible to work through variations of a scene or to generate a number of views or individual images for films with a corresponding script.

Most settings for both procedures can be read from LP or LS files. However, several options must be explicitly activated with command options.

4.1.7 Distributed computation

Parallel use of several computers within a network, as is typical for modern office configurations, can be very helpful in the solution of a simulation problem. According to the manual, Lightscape supports up to 1000 CPUs. However, the radiosity computation for a single scene cannot be distributed. What can be distributed is the ray tracing for an image, image variations, images for film sequences or several different radiosity calculations. All the machines involved are equipped with control software and require access to a defined working directory on a server. Any work station can act as the server. The jobs are then automatically distributed and worked through.

4.2 Radiance

4.2.1 Starting the simulation

After the octrees have been generated (section 3.2.1) and the simulation has been started, it is no longer possible to change the settings for the surfaces, in contrast to Lightscape.

Unlike Lightscape, it is not possible to make individual surfaces "transparent", so that they are visible but do not cast any shadows. It is practically impossible to start an inconsistent calculation: Everything which is visible in Radiance affects the light distribution. The disadvantage of this feature is that orientational geometric aids, such as points to mark co-ordinates, arrows and similar devices, cannot be simply incorporated directly into the Radiance image.

The `rpict` program is used to produce pictures, whereas `rttrace` is for individual points. In both of them, the user can choose whether the radiance of a point in a given direction or the irradiance on a surface element is to be calculated.

The photometric quantities, luminance and illuminance, are obtained from the radiometric quantities by the usual weighting with the standard photopic response function $V(\lambda)$.²

4.2.2 Analysis

The "classic" output for photorealistic images is reproduced in fig. 4.3: A perspective view of the luminance distribution from the camera viewpoint. In addition to perspective views, Radiance offers other image projections: fish-eye views, parallel projection and cylindrical panorama views (see fig. 4.4). In the sequence shown in fig. 3.8 after the picture has been calculated with `rpict`, the image is output to the monitor with `ximage` or as a PostScript file to the printer with `ra-ps`.



Figure 4.3: Calculated Radiance image for luminance distribution

These images can also be presented in false colours, in which case the radiance or luminance is converted into a linear or logarithmic colour scale with the `falsecolor` program (fig. 4.5 and fig. 3.8). This is appropriate for quantitative analysis, e.g. to assess glare problems, as otherwise the large dynamic range for radiance or luminance cannot be reproduced (see also section 4.2.3).

In addition to the radiance and luminance distribution, images of the irradiance or illuminance distribution can also be calculated. As both quantities are related via the material attributes of the surfaces,

²Skotopic weighting might be done by the user, but requires user given weights for the three channels.



Figure 4.4: Different types of projection of a Radiance image from the same viewpoint: (clockwise from the top: cylindrical, linear fish-eye, cosine fish-eye). A view in parallel projection is not shown here.




Figure 4.5: False-colour image of the results of fig. 4.3: With a linear scale (above), and with a logarithmic scale (below).

a second computation ³ is necessary. However, the ambient values from the radiance and luminance calculations can be used again. An example is shown in fig. 4.6.

Finally, the images for the luminance and the illuminance can be combined in one single image, in which isolux contours quantify the illuminance (fig. 4.6).

The calculation of illuminance values in planes, a classic feature of lighting technology programs, is offered via the `rtrace` program. It can accept any number of position co-ordinates as input and provides the illuminance values for them as output ⁴. The co-ordinates are defined e.g. within the CAD program and then exported, or are generated by the `rshow` program, which is available as a further option. The latter also allows the calculated results to be presented within the scene geometry (fig. 4.7).

4.2.3 Further processing for output devices

One aspect of simulation work which is often overlooked concerns the open questions related to outputting the results as a colour print, slide, video or image on a computer monitor. This is equivalent to the photographer's task of reproducing a scene, such as that in fig. 1.5, on photographic paper or a slide.

A real scene can easily span a dynamic range in luminance (contrast) of 1 : 10⁵: The dark sections of the image correspond to only a few candela/m², whereas directly illuminated areas or reflections of luminaires or the sun can have several ten thousand candela/m².

Radiance preserves this contrast range completely in the calculations. This is one of the decisive differences to other ray tracers in computer graphics: Radiance calculates and saves each value in a format which preserves the complete dynamic range.

However, currently available output devices are not able to reproduce this dynamic range: a slide has a range of about 1:100, i.e. the light parts are about 100 times lighter than the dark ones. The contrast is similar for computer monitors and video projectors (maximally 1:1000 for three-tube light valve projectors). All reproduction forms which are not light-emitting, such as photos and prints, have an even smaller dynamic range. Various recent methods attempt to compensate for this and are approaching the values for luminous media [Kir97].

How should the calculated values, which have a large dynamic range, be reproduced by output devices with a very restricted dynamic range, such that the subjective visual impression on viewing corresponds to reality?

Clearly, this problem has been solved more or less satisfactorily for photography, so that fig. 1.5 appears approximately "real". If we consider that such a photo was preceded by careful preparation, selection of the daylight conditions and photographic techniques, and that further post-processing occurred for the enlargement and printing, achievement of this reproduction quality for photography is not trivial.

In the terminology of computer graphics, this aspect is called *tone mapping*[Rus93], and is an area which has not yet been fully explored.

Radiance offers two programs for tone mapping: `pfilt` transforms the image in the same way as a real camera, where different apertures can be chosen for a given lens. Parts of the image which are too bright appear overexposed on the film, whereas dark parts of the image disappear into blackness (fig. 4.8). This corresponds to the so-called optical density function for a negative or slide film.

The more recent `pcond` program applies information about visual perception by the human eye: Bright parts of the image over-illuminate darker sections (scattering on the iris and in the vitreous humour of the eye), dark sections are seen scotopically (i.e. in black and white, and less well focussed), corresponding to the sensitivity of the rods in the retina (fig. 4.9). In addition, the contrast range of

³with the `-i` option of `rpict`

⁴`rtrace` can output any quantity from the Radiance system, but the irradiance or illuminance values are the standard application.




Figure 4.6: False-colour image of the illuminance with linear and logarithmic scales. As is to be expected, the logarithmic scale shows a greater dynamic range. The lowest illustration shows isolux contours superimposed on the image from fig. 4.3.

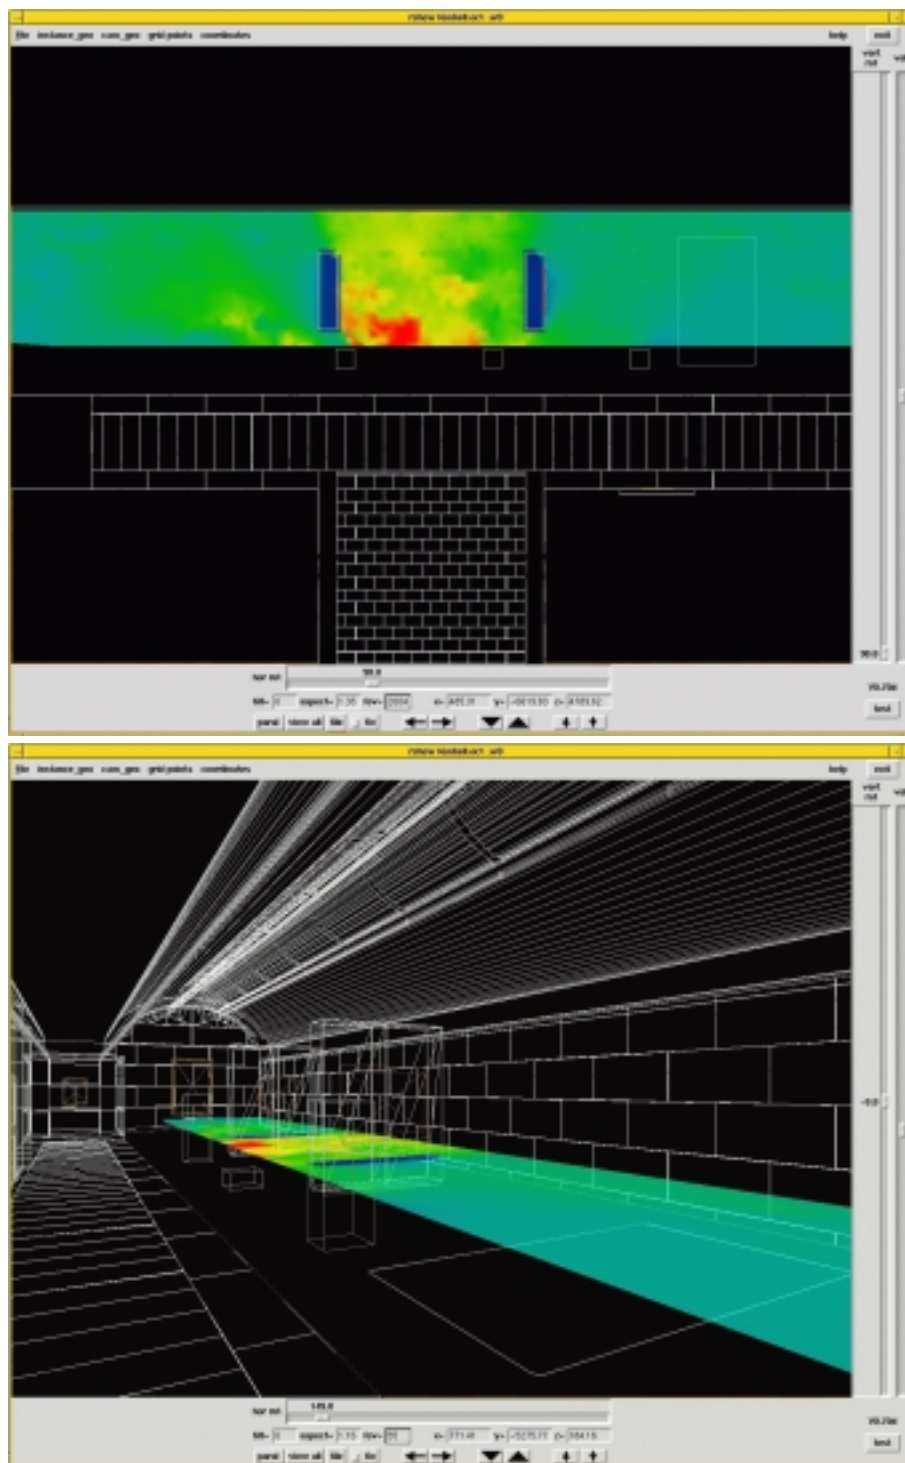


Figure 4.7: False-colour image of the illuminance with a linear scale, inserted into the geometric outline of the scene (`rshow`).




Figure 4.8: The artificial aperture in `pfilt` acts like the aperture of a camera: These three pictures are based on the same calculated data for the luminance. The pictures correspond to a so-called f-stop sequence in photography.

the output device and its colour reproduction can be specified.

With `pcond`, Radiance offers a good tool for photorealistic rendering, particularly in creating the same spatial impression for the viewer. In order to achieve this aim, however, the entire context of the presentation of a simulation must be taken into account, e.g. the viewing angle at which the image is seen by the observer. Such aspects are beyond the scope of this investigation.

4.2.3.1 Output to special devices

Figure 4.10 shows a so-called anaglyph, as an example for "unorthodox" reproduction of the calculated image: A superposition of a red and a green image, which generates a three-dimensional image when seen through special viewing spectacles. The spectacles simply consist of a red filter for the right eye and a green filter for the left eye.

An anaglyphic representation does not contain any colour information about the object itself and is mainly used as a simple method to achieve easily viewed 3D images, which are immediately perceived as three-dimensional.

Other stereoscopic representation forms which also reproduce colour include e.g. projections with crossed polariser planes (spectacles with two polarisation filters) or computer monitors with temporally multiplexed images (shutter spectacles), as well as reproduction with two monitors or slides.

These methods of reproduction are applied in order to convey an easily perceived impression of an unfamiliar three-dimensional geometry to all viewers in a wider audience.

4.2.4 Distributed computation

Parallel computation is supported by Radiance in two ways:

On the one hand, on multiprocessor UNIX machines: Whereas these were still confined to the expensive UNIX hardware from SUN, HP or SGI several years ago, now 2 or 4 processor machines based on Intel, with Linux as the operating system, are available for a moderate price. Here, Radiance is able to use all of the processors for computation.

In order to calculate two images with the same octrees from two different viewpoints, using a PC with two processors and Linux as the operating system, the `rpict` programmes are started simply with the additional `-PP` option:

```
rpict -PP lock.file ... rpict options ... -vf 1.vf -o 1.pic abc.oct
rpict -PP lock.file ... rpict options ... -vf 2.vf -o 2.pic abc.oct
```

The second program shares the memory space for the scene geometry with the first one. This means that complex, and thus memory-intensive, geometry only occupies space once in the main memory.

The `rpiece` program organises the computation of *one image* by several `rpicts`, so that a computation is accelerated in proportion to the number of processors or computers in the network.⁵

Parallel computation in a local area network (LAN) is also supported and only differs slightly in the program start from the first solution. This is transparent for the user. However, a completely automatic distribution of loads among several computers goes beyond the scope of lighting simulation alone. To this purpose, the simulation is typically coupled with a job distribution system (e.g. the `psub` system):

<http://www.ise.fhg.de/radiance/psub/psub.html>

⁵As of this writing (April 2001), support has been added into Radiance for PVM (Parallel Virtual Machine) and MPI (Message Passing Interface). Both support parallel computation on NOW (network of workstations) and multi-processor machines (e.g. Cray).



Figure 4.9: The same calculated values as in fig. 4.8, corrected with p_{cond} for reproduction on an output device with maximally 100 cd/m^2 and 1:32 contrast. The lower image is also weighted with visual perception parameters (scattering of bright points, scotopic vision).




Figure 4.10: Stereo image for red/green spectacles. These so-called anaglyphs are composed of two Radiance images for the view from the left and the right eye. The viewer needs spectacles with red and green filters in order to see the image in three dimensions. The optimal viewing distance in this case is about twice the image width.

Chapter 5

Conclusions

5.1 Degree of modelling detail in a scene

When a simulation is carried out, the question about the degree of detail needed in the CAD models is often raised. If the details are not standard components of an object library, such as tables, chairs or picture frames, increasing their number in a scene increases the modelling complexity and thus the costs.

All of our experience indicates that this extra effort is fully justified.

The visual impression is determined more by "minor details" than by large areas: grooves, edges, rounded corners and the like contribute essentially to the visual credibility of a scene. Naturally, this applies in particular to a zoomed image, but also to an extent which should not be underestimated, when the details form only a relatively small part of the overall image.

As an example, fig. 5.1 shows three levels of detail for the model of a glass case:

Firstly, the glass case was modelled as two simple cubes: A lower one of wood and an upper one of glass. The glass was modelled as a "single shell", i.e. an infinitesimally thin glass surface, with the reflection defined by the refractive index and absorption coefficient. This quickly prepared model provides physically correct results for windows flawlessly. However, in this case it does not look very real, as is immediately apparent.

In the second step, the glass of the case was modelled as 8 mm thick panes, with the reflection now determined by the refractive index, the absorption coefficient and the pane thickness. This extra work on the model adds essential features to the image, in that the edges of the panes are now shown correctly: They appear dark due to multiple reflections in the glass, or light due to total internal reflection (upper right edge). However, the show case still does not look completely convincing.

In a third step, a darker material (e.g. black velvet) is assigned to the lower surface of the show case, and the glass panes are positioned in a groove, 10 mm clear of the edge of the pedestal.

Consistent modelling with "attention to detail" results in better modulated images with a more convincing visual impression.

5.2 Lightscape

5.2.1 Manageability of Lightscape

The Lightscape lighting simulation program provided by Autodesk is an easily used product, which can be worked with effectively after a short introduction period, thanks to its structured graphical user




Figure 5.1: Models of the glass case with different levels of detail.

interface (GUI). However, as so often, it is the details which create the problems.

The import of geometric data is generally convenient and transparent. Only the capping function can be irritating.

We particularly liked the editing mode to process individual blocks. If the user's hardware includes a rapid graphics card, the orientation of surfaces and the assignment of materials can be changed quickly and easily.

If the user has set up material and block libraries over time in a strategically favourable configuration, the assignment of materials or substitution of dummies by already configured library elements is a matter of seconds.

The only restriction is in the possibilities for generating materials, which however is a considerable detraction for photorealistic rendering.

Filter functions help in the selection of objects to be processed.

However, the creation of perspective views proves to be very tedious at times. If the user chooses the interactive mode for generation, then all of the settings must be made by mouse clicks, without exception. Modifications are not implemented until the mouse has moved a certain minimum, pre-determined distance. Fine adjustments then become practically impossible. With large models, the display changes by jerks when the mouse is moved, which also makes fine adjustments difficult. As an alternative, the viewpoint and vanishing point for the perspective can be determined from an overhead view of the scene by mouse click, and then finely adjusted by keyboard entries. The settings can be checked using the perspective view. Thus, the user is continually switching between the overhead view and the perspective view until the final choice has been made. The image is regenerated each time, a time-consuming process for large models.

Altogether, adequate conformance between the hardware capabilities and the model size is urgently recommended. It is true that a large model can in principle be processed via individual components. However, alone the effort needed to co-ordinate possible interdependence between the individual models would not be in any relation to the benefit. It is also revealing that Lightscape does not support the import of external references from AutoCAD.

Assistance from the "wizard" in setting the simulation parameters is undoubtedly a great aid to Lightscape novices. The only disturbing feature is that a few parameters are not adapted by the wizard and lead to results which cannot be readily comprehended. Understanding of such "refinements" demands considerable experience with the program.

The large size of the files generated by the program is a negative point. It was not possible to determine how the CAD co-ordinates are stored. For instance, if the same geometry is imported from AutoCAD, once with "metres" as the reference unit, then with "millimetres", and the stored CAD co-ordinates are compared, the numbers are found to differ from about the fifth or eighth digit respectively after the decimal point. As rounding tolerances are obviously responsible, it is not clear why 14 or 17 digits after the decimal point should be saved.

The computation times increase rapidly, as soon as the quality is in the medium to high range. Repeated further processing of the surface attributes was needed in order to achieve results from a single computer which represented an acceptable relationship between the computation time required and the benefits obtained.

5.2.2 The test points

5.2.2.1 Model size, subsequent model processing

If the model size is larger than the corresponding hardware capacity, it is almost impossible to work effectively on the complete model. Every time the material or geometry is modified or an input window is closed, the scene is reconstructed on the monitor.

In order to process a model of this size despite this limitation, the following strategies were adopted:

Adapting level-of-detail for various parts of the model: However, this means that the views and the viewable parts must be selected early in the work.

Manipulating the surface attributes: The process of increasing the grid mesh resolution during the radiosity procedure was either greatly reduced or suppressed for all surfaces which were irrelevant for the overall result. The post-processing is time-consuming, even for a well-structured model. In addition, there is no possibility for globally displaying the level of detail for selected surfaces.

Restricting the amount of light distributed: The simulation was terminated after 90% of the initial energy had been distributed.

Restricting the quality: The "wizard" was used to set the simulation parameters, in order to approach normal working conditions. The quality was always set to "medium". A simulation with the "highest quality" setting was terminated after several days of calculation, when no appreciable progress in the distribution of the light could be observed.

If the radiosity procedure is to be applied efficiently, the subsequent distribution of the grid mesh must already be taken into account during the modelling phase. Unevenly proportioned surface forms result in "frayed edges" along a transition from light to dark areas, e.g. along shadow boundaries, in the radiosity procedure.

If the user does not work with a "pure" version of AutoCAD, but instead uses an additional application on AutoCAD for modelling or imports geometries from other CAD programs, it is seldom possible to influence the geometric erection of pre-processed elements, e.g. a wall with a window opening.

During conversion, the approach is usually to minimise the number of individual surfaces, which is generally not optimal for Lightscape.

5.2.2.2 Treatment of light

If daylight is to be taken into account during simulation, simulation parameters can be set to determine whether the entire sky should be used for the computations (for outdoor scenes) or only that portion of daylight which passes through windows or other openings (for indoor scenes). The latter approach results in better accuracy and higher efficiency, as only that sector of the sky corresponding to the apertures is taken into account. The disadvantage of this method is that all windows or openings have to be explicitly defined as such. As a result, "holes" in the model can hardly be noticed. For purely visual results, this situation is not critical. However, if accurate information on lighting values is needed, the "light tightness" of the model must be painstakingly checked initially.

5.2.2.3 Light distribution via the reflector plate

Ensuring that the incident daylight was re-directed onto the ceiling vault and from there into the gallery proved not to be completely straightforward.

The template for metal was used as the material for the upper, opaque zone of the reflector plate. The reflectance was set to a value of 90%. During the first test runs, it appeared as though the underside of the plate were shining brightly. In order to suppress this diaphanous effect, a second, opaque plate was modelled, with a matt grey painted surface, positioned 0.1 mm away from the original one. Despite these precautions, it still appeared as though the daylight entering from above shone through both surfaces. The cause of this phenomenon was an inappropriate value for the distance tolerance (see fig. 5.2).

Lightscape set this parameter by default to a value of 0.5 mm.

If two points are closer than this distance, apparently they are no longer processed separately. By reducing this value to 0.0005 mm, not only was the desired effect achieved, but the computation also became faster. This tolerance value is not treated when the simulation parameters are set using the "wizard". A far-reaching result besides the visual difference is, however, the distribution of light inside.

To illustrate the point, a section of the model for the ceiling was simulated (fig. 5.2). The settings for the two variations differ only in the tolerance value, but the results are very different.

The errors in the calculation arise mainly from the assumption of a diffuse material for the upper part of the reflector plate. Back-scattering of light into the PMMA cover leads to underestimation of the portion reflected by the ceiling into the room.




Figure 5.2: The light-distributing plate under the light slot in the cycloidal ceiling vault, as modelled in Lightscape: Above, with the default tolerance value of 0.5 mm. The light "shining through" the upper part of the plate can be clearly seen. Below, the effect of changing the tolerance value to 0.0005 mm is illustrated. The change in the luminance values is clearly visible.

5.2.2.4 Limits of the material definitions

Definition of the material for the lower, perforated part of the reflector plate exceeded the limits set by Lightscape. The material can not be defined as required. On the one hand, it is not possible to combine two materials, e.g. to control transparency via the brightness differences of pixel images. On the other hand, there is no way of generating a material with "mathematical" definitions, analogously to the function files in Radiance. The only possibility would be to imitate the geometry of the perforated metal sheet exactly in the model. This would mean that the grid mesh for the radiosity procedure needed to be extremely fine, which would automatically lead to longer computation times. In the model presented here, 5% transparency was assigned to the material for the lower part of the reflector plate, but the results were not convincing.

5.2.2.5 Computation times

Using a PC equipped with an AMD-K6-III 400 MHz CPU, it took about 50 hours until 90% of the initial light was distributed. The program had completed about 5000 iterations. The radiosity procedure had generated about 170,000 grid meshes with about 345,000 mesh vertices. About 58 MB of working memory was occupied toward the end of the computation. The directly incident light from the sun and the luminaires was not saved, as a ray-tracing image was subsequently generated. Only the daylight which entered via windows and other openings was taken into account.

Afterwards, it took about 12 hours to generate an image of 1834 x 1536 pixels with the ray tracer. The pixel colours were interpolated with an anti-aliasing factor of 4. The directly incident light and the shadow pattern were recalculated for the sun and all luminaires. The calculation option with "soft" shadow edges was activated for the sun. Otherwise, the default settings from the solution file were used.

5.2.2.6 Manageability of luminaires

In order to allow direct comparison between Radiance and Lightscape, the light fixtures were modelled in AutoCAD as blocks. The candlepower distribution curves could then be positioned using the editing mode in Lightscape without any difficulties. If the reference points for the luminaires are correctly defined, the orientation of a luminaire can be set automatically with mouse clicks ("aiming a luminaire instance").

5.3 Radiance

5.3.1 The test points

5.3.1.1 Model size

The degree of complexity of the model requires about 100MB main memory, which should not present any problems for most personal computers.

The `rshow` program to view the geometry was started with the `-qs 0.1` option, which means that all details of dimensions less than 10% of the overall scene are not portrayed. This made interactive selection of the views and camera positions feasible¹.

The model size did not present any problems in Radiance.

5.3.1.2 Light distribution

Daylight enters the gallery only indirectly via the reflector plate (see next section) and via the small inner courtyard, which is illuminated directly by the sun and thus is very bright. This diffuse source demanded high values for the `ad` and `ab` parameters in Radiance to avoid errors and thus "mould patches" in the image. The computation time was determined almost exclusively by these parameter settings.

5.3.1.3 Light distribution via the reflector plate

This incident light is reflected via the almost ideally reflecting plate and the diffusely reflecting ceiling downwards into the room. If sunlight is directly incident on the curved plate, as in this case, the light

¹with the same hardware, Matrox G400 graphics card




Figure 5.3: The light-reflecting plates under the light slot in the cycloidal ceiling vault, modelled in Radiance: Light transport via the curved reflector plate is modelled with perceptible statistical fluctuations, even with higher values of the ad parameters.




Figure 5.4: Example for replication of a material texture in Radiance: Left, the original image of a parquet texture, which was replicated four times (arrows). The irritating light/dark fluctuations are much less pronounced due to appropriate pre-processing in the texture shown at the right.

path is treated inadequately and at high computation cost in the current version of Radiance (fig. 5.3). An extension of the algorithms is needed here.

5.3.1.4 Material definitions

There are hardly any limits to the materials which can currently be modelled in Radiance. For example, the hole-to-material ratio of the finely perforated reflector plate was modelled as 15% transmitting, with the remainder reflective. The material for the flat ceiling between the vaults is defined as an anisotropically reflecting metal, with parameters which were modelled on the basis of original photographs, as material samples were not available. The photographs of reflections in the ceiling material indicate that it is rolled metal sheeting, probably anodised aluminium.

The results for replicated texture images, as discussed in section 1.3, are shown in fig. 5.4.

5.3.1.5 Computation times

The computation times for presentation quality ($ab=7$, $ad=16000$, 4000×4000 pixels) were about 24 hours for the first image and 2-3 hours for further views of the same geometry². No special "tricks"

²on a computer with Pentium-II, 400MHz processors

had been applied to accelerate the computations. A disadvantage of Radiance is that it lacks a check for convergence, e.g. indicating how much light has been distributed. This is difficult due to the nature of the currently used algorithm.

5.3.1.6 Manageability of luminaires

This is currently a major "minus point" for Radiance, as luminaires and candlepower distribution curves can not be entered interactively.³

The modularity of the Radiance system is its strength and its weakness simultaneously: It is extraordinarily flexible, and thus initially somewhat obscure. Good graphical interfaces are an essential requirement to overcome this situation.

5.4 Summary

The algorithms in Radiance (distributed ray-tracing) are superior to those in Lightscape (radiosity), as they calculate the light transport for surfaces which are not ideally diffusely reflective. If redirection of daylight or indirect lighting via the ceiling is considered, with real ceiling materials which usually deviate from an ideal Lambertian light-scattering characteristic, the radiosity solution is inadequate.

The disadvantage of the Radiance algorithm is the relationship between subjective image quality and computation effort: Whereas radiosity already provides visually appealing images after a few iterations, the Radiance images are unattractively patchy, if their calculation is based on fast, but inaccurate, parameter settings.

Computations must be made accurately and thus require appreciable time, even for relatively simple scenes.

However, it should not be forgotten that Lightscape requires similar computation times until all the light has been distributed, i.e. until quantitatively stable results are available. It just looks more attractive beforehand.

The function files in Radiance make its material modelling very much more powerful than Lightscape. This is primarily significant when new materials are to be modelled, with properties which fundamentally can not be modelled in Lightscape.

The Radiance core needs extension with regard to ambient calculations using user-defined BRDF materials and the calculation of caustics.⁴

The latter facility is a necessary extension for materials like the highly reflective reflector plates in the museum, as then the reflection of directly incident sunlight on curved reflectors would be treated correctly. Due to its modular structure and accessible source code, Radiance is a good platform for such extensions.

It is much easier to select and handle luminaires or materials, or incorporate candlepower distribution curves, in Lightscape than in Radiance. The current version of Radiance has clear weaknesses in this area.

Both programs demonstrate the wide range of possibilities available today for professional, photorealistic lighting simulation, which can be achieved in practice with average personal computers.

³This may be available with 3rd party packages like "rayfront" or "Desktop Radiance"

⁴As of this writing (March 2001), such development had been started (see <http://www.radiance-online.org>).

Bibliography

- [Aut99] Autodesk, Inc. *Lightscape User's Guide 3.2*, April 1999.
- [BL91] David B. Brownlee and David G. De Long. *Louis I. Kahn, In the Realm of Architecture*. Rizzoli International Publications, New York, 1991.
- [Bra92] Michael Brawne. *Louis I. Kahn, Kimbell Art Museum, Architecture in Detail*. Phaidon Press Ltd, London, 1992.
- [Del95] Jean-Jacques Delaunay. *Contribution à la Modelisation de la Lumiere Naturelle en vue de son Application à la simulation de l'Eclairage de Locaux*. PhD thesis, Université Louis Pasteur de Strasbourg I, 1995.
- [Fra99] Francesi. *Louis I. Kahn, The construction of the Kimbell Art Museum*. Skira editore, Milan, 1999.
- [Kir97] Kirchner. Novospace- neue drucktechnik (lichtakademie gr"undungstreffen). private communication, 1997.
- [Rus93] Holly E. Rushmeier. From solution to image. In *Making Radiosity Practica*. Siggraph Course Notes, Course 22, August 1993.
- [WLS98] Greg Ward Larson and Rob Shakespeare. *Rendering with Radiance*. Morgan Kaufmann, 1998.
- [WRC88] Gregory J. Ward, Francis M. Rubinstein, and Robert D. Clear. A ray tracing solution for diffuse interreflection. In John Dill, editor, *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 85–92, August 1988.

About the authors

Kurt Altmann studied architecture at the Technical University of Karlsruhe, the Royal Technical University of Stockholm and the University of Tampere. As a student, he gained extensive experience with CAD and lighting simulation. After graduating in 1997, further experience was gained in work for various employers (including Fraunhofer ISE, Freiburg, Germany). The main emphasis lay on the conversion and processing of CAD data, and visualisation of lighting effects and visualisation of lighting. He set up his own office in Karlsruhe in 1999 (<http://www.kurt-altmann.de>).

Peter Apian-Bennewitz studied physics in Freiburg and obtained his Ph.D. in 1995 with a thesis on the measurement and modelling of materials in optical simulation programs. Since 1991, he has worked at Fraunhofer ISE and in 1999 he established his own company, pab-opto (<http://www.pab-opto.de>), for technical visualisation.

The Fraunhofer Institute for Solar Energy Systems ISE, is one of 40 institutes in the Fraunhofer Society, which offers industrially relevant, applied research as a service. Further information on the Fraunhofer Society can be found under <http://www.fhg.de>. As well as other aspects of solar energy application, Fraunhofer ISE also concerns itself with daylighting, solar building, materials development and thermal/optical measurement of materials: <http://www.ise.fhg.de> and <http://www.ise.fhg.de/radiance>.

The Lichtakademie was established in 1998 for inter-disciplinary analysis of questions concerning light. The questions range from the aesthetics of lighting-induced atmospheres through photo-biological effects to fundamentals of simulation: <http://www.licht-akademie.de>

Many thanks to Mrs. Helen-Rose Wilson for providing the English translation of the initial German version.

This report was written using \LaTeX on an Intel Pentium-II PC running Linux.



The hardware and software

Hardware for Lightscape:

Gigabyte dual mother board with two times 200 MHz Pentium I, 128 MB working memory EDO-RAM, Gloria Synergy graphics card from Elsa with 8 MB RAM

Asus P5A mother board with AMD K6 III 400 MHz, 128 working memory SDRAM, Gloria Synergy II graphics card from Elsa with 32 MB RAM

Hardware for Radiance:

Asus mother board with two Pentium 400 MHz and 390 MB ECC-RAM, G400 graphics card

Operating system for Lightscape:

Windows NT Workstation, version 4.0 with Service Pack 4

Operating system for Radiance:

Linux 2.2.13 (Slackware 7 distribution)

Software used:

Modelling:

AutoCAD R14.01 (without additional options), distributor Autodesk

Lighting simulation:

Lightscape v3.2, distributor Autodesk

3D-Studio MAX v2.5 (plug-in for the import and export of DWG files), distributor Autodesk

Radiance 3R1P20, distributor <http://www.radsite.lbl.gov/radiance>

Index

— A —

Anaglyph, 45
Analysis
 Lightscape, 35
 Radiance, 38

— D —

distributed computation
 Lightscape, 37
 Radiance, 45

— I —

IES
 Lightscape, 15

— L —

Lightscape
 Analysis, 35
 batchmode, 35
 Blocks, 13
 distributed computation, 37
 handling, 49
 Light sources, 15
 Materials, 17
 limits of, 53
 Opening, 19
 optimizing strategies, 52
 Raytracing, 35
 Simulationsparameter, 19
 sky models, 17
 test points, 51
 Window, 19

— M —

Materials
 Lightscape, 17
 Radiance, 25

— P —

Perez sky model, 25

— R —

Radiance
 ambient calculations, 29
 analysis, 38
 color, 26
 distributed computation, 45
 materials, 25
 MPI, 45
 oconv, 22
 output devices, 41
 overview, 20
 parameter, 29
 pfilt, 41
 PVM, 45
 rshow, 23
 rview, 23
 Sky models, 25
 test points, 54
Raytracing
 backward, 26
 Lightscape, 35

— S —

Sky models
 Lightscape, 17
 Radiance, 25
Stereo image pairs, 45

— T —

test points
 Lightscape, 51
 Radiance, 54
tonemapping, 41