Eighth International IBPSA Conference
Eindhoven, Netherlands
August 11-14, 2003

OPTIMIZING DAYLIGHT SIMULATION FOR SPEED AND ACCURACY

Benjamin Geebelen and Herman Neuckermans
Department of Architecture, Urban Design and Regional Planning
K.U.Leuven, Leuven, Belgium

ABSTRACT

A study by Reinhart and Herkel showed that long-term predictions of daylight availability in architectural spaces should take the daylight conditions of all individual time steps into account. However, most contemporary simulation tools and algorithms require such long computation times that it is impractical or even unrealistic to perform a simulation of the daylight penetration for each time step. This paper discusses two adaptations of a known algorithm, radiosity, that bring down the required computation time for an annual prediction from the order of days to the order of seconds.

INTRODUCTION

The growing interest in the potential of daylighting in buildings for saving energy and cutting carbon-dioxide emissions calls for a re-investigation of the prediction techniques and tools that we have at our disposal to assess the performance of a design.

Since the 1980s lighting simulation software has evolved considerably. Powerful tools, such as Radiance or LightScape[1], allow expert users to produce both amazing images and a numeric output of reasonable accuracy, but are they really suited for daylighting analysis?

This paper discusses the results of a study of contemporary lighting simulation algorithms. The objectives were twofold:

1. to identify among the leading paradigms the most efficient algorithm for daylighting studies;

2. to adapt this algorithm so that it would be suited for long-term prediction of daylight availability. This entails assuring accuracy and maintaining practical computation times (within the order of minutes).

[1] LightScape is no longer in development. Some of its functionality is currently incorporated into AutoDesk VIZ4.

CUMULATIVE DAYLIGHT AVAILABILITY PREDICTIONS

If we want to fully assess the quality of the daylighting concept in a particular design alternative, we are faced with the task of performing predictions over long time spans. Although the simulation of daylight penetration for a single instant may be quite interesting, an estimate on a seasonal or an annual basis, e.g. the daylight autonomy, the annual relative usable light exposure or the impact on the seasonal and annual energy consumptions, is far more essential.

Since the sky is a highly variable light source – its luminance distribution changes from one minute to the next – long-term predictions may present a challenge. The overall light distribution in a proposed room changes with the sky-luminance distribution. This implies that a single prediction of the indoor light distribution may not suffice for judging the performance of a daylighting concept over an entire year. However, given the long computation times that are still common among contemporary lighting simulation tools, running a simulation for each minute of a typical year is a sheer impossibility. Even with a time step of one hour a complete-year run requires a computation time that is only justifiable for academic purposes.

Different approaches have been tried and tested to overcome this problem by limiting the number of sky conditions to be simulated. The daylight-factor method assumes a single constant sky-luminance distribution and thus bases a whole-year estimate on a single simulation. Some tools, such as ADELINE (Erhorn et al., 1998) and DOE2 (Winkelmann et al. 1985), use interpolation schemes between sets of overcast and clear skies. Herkel and Pasquay proposed the reduction of simulations by categorizing daylight conditions according to sun position and direct and diffuse illuminance into some 450 classes (Herkel et al. 1997). Each of these solutions effectively brings down the number of simulations to be performed by at least an order of magnitude. However, a study by Reinhart and Herkel showed that considerable discrepancies can occur between the results of any of these approaches and a prediction that considers the daylight condition of each time step individually (Reinhart et al. 2000).

The following sections describe two adaptations of a classical radiosity algorithm that enable cumulative predictions within practical computation times without jeopardizing accuracy.

RADIOSITY VERSUS RAY TRACING

Daylighting software can be subdivided into two categories: tools that are mere digital translations of known simplified calculation methods for the direct and average internal reflected components of the daylight factor, and tools that try to model light propagation as physically accurate as possible. Most tools in the latter group adhere to one or a combination of the two key paradigms in physically based rendering, i.e. ray tracing and radiosity.

In accordance with their name, ray-tracing algorithms follow the paths of individual light rays through a scene in order to determine the luminance of the scene's surfaces at a limited set of singular points in the direction of a pre-set viewpoint. Radiosity algorithms, on the other hand, apply a finite-element approach to find a light distribution within the scene that comes close to an energetic equilibrium. The two counterparts are almost each other's opposite: Ray tracing performs best with specular materials and point light sources, allows any geometry, and is view-dependent, whereas radiosity performs best with perfectly diffuse materials and large light sources, requires faceted geometry, and is view-independent.

For daylighting analysis, one is interested in predicting accurate luminance and illuminance values in the scene. It would seem that radiosity is the more efficient algorithm to this end, since estimating the overall light distribution is its corner stone, whereas ray tracing only observes distinct points and light rays[2]. It is quite complex to assign appropriate light fluxes to the individual rays so as to get an idea of the overall light distribution. Especially diffuse inter-reflections between surfaces are hard to reproduce with ray tracing. Nevertheless radiosity is often dismissed for not being able to handle specular reflections and for requiring large amounts of memory. In the context of daylighting analysis, where a visual output is irrelevant, we find this reasoning unjustified:

1. In the simulation of the built environment, the overall light distribution is generally determined by inter-reflections that come closer to diffuse than to specular behavior. Only for special mirroring systems are specular reflections prevalent;

2. It is not impossible to introduce specular behavior into radiosity. The simple technique of using ray casting to determine extended form factors can already take specular reflections into account, thus facilitating the simulation of mirroring systems. This will not suffice for creating images, but performs adequately for producing illuminance values;

3. Most solutions to introduce diffuse inter-reflections into ray tracing involve a proliferation of rays to be traced, thus badly affecting its efficiency;

4. In a time of persistently growing computer memory the argument of radiosity's excessive memory demands has become obsolete.

One could argue that, because of their complementary natures, a combination of both algorithms – e.g. a two-pass procedure – would be the ultimate solution. However, based on the reasoning above, it is our belief that the ray tracing step is only necessary to generate the view-dependent effects when producing images, and that, when focusing on light distribution and numeric output, the bulk of our attention should go to improving radiosity algorithms.

In order to facilitate long-term predictions that consider the daylight conditions of all time steps, we will concentrate on two goals: the reduction of the required computation time for a single simulation, and the reduction of the number of simulations to be run.

LIMITED PATCH SUBDIVISION

Radiosity algorithms introduce a mesh of nodes across the surfaces of the scene, i.e. the geometrical model in which the light propagation is to be simulated. Simultaneously the surfaces are subdivided into patches. All light transfer is assumed to be concentrated in the finite number of light paths between the nodes. Then a radiosity[3] value is found for each node, so that the corresponding fluxes along the light paths constitute a state of energetic equilibrium. Finally, the radiosity and corresponding luminance is obtained for each point of the surface patches by interpolating between the radiosity values of the adjacent nodes (Glassner 1995). The quality of the simulation increases with the number of nodes; the quality of the interpolation increases with the number of nodes per patch. The remainder of this

[2] This assumption seems to be confirmed by an investigation of the computation methods used in Radiance, one of the most well-known and respected lighting simulation programs. Though essentially a ray tracer, Radiance uses a grid of nodes in which irradiance is cached in order to account for diffuse inter-reflections (Ward et al. 1998). The resemblance with radiosity is striking.

[3] *Radiosity* is a synonym for *radiant exitance*, which is defined as the power radiated per unit area of a surface.

paper will relate to classical radiosity, in which the luminance values of all points of a surface patch are constant and equal to the luminance of a single node located in the patch's center of gravity. In other words, the number of nodes is equal to the number of patches.

The Importance of Grid Density

Since the criterion for the solution is the energetic balance of the light transfer between the nodes, finding a solution will consist in solving a system of as many equations as there are nodes. For a scene that contains *n* nodes, this will result in a system of *n* equations with $n \times n$ coefficients. Each coefficient corresponds to one of the light paths between two nodes. The two most time-consuming tasks in any radiosity simulation consist of determining these coefficients, which corresponds with the computation of form factors, and solving the system. Due to geometrical complexity, the former task is generally avoided by using some fast approximating procedure, e.g. ray casting or treating patches as differential elements. Due to the mostly huge numbers of equations, the latter task is generally performed iteratively, e.g. by using Jacobi, Gauss-Seidel or Southwell iteration. Important, though, is that the number of nodes is critical for both tasks, and that, whatever the simplification strategy, computation time will always increase more-than-linearly with the number of nodes. It is therefore logical that this parameter constitutes an ideal starting point for reducing computation times.

Traditionally radiosity algorithms are applied in the field of digital image synthesis, where the ambition is to produce highly realistic images without artifacts. For this purpose high-density node grids are required. Over the years, a great deal of research has therefore been devoted to techniques that can cope with great numbers of nodes. However, for daylighting design and analysis, visual output is irrelevant. We do not need to know every nuance across a scene's surfaces; we merely need the resulting illuminance values in a predefined set of reference points. We should therefore investigate whether we can keep the number of nodes deliberately low without jeopardizing the accuracy of the numerical result.

Experimental Setup

In order to isolate the effect of node density, which is directly linked to surface-patch size, on accuracy, we chose to simulate an ideal case and compare results with the analytical solution. In a case without any noteworthy luminance gradients across the surfaces, the effect of surface-patch size would most likely be minimal, rendering the experiment useless. We therefore opted for an infinitely long room with a longitudinal roof light under a CIE Standard Overcast Sky. All surfaces were either perfectly diffuse or perfectly absorbent. The infinite length of the room was mimicked by using perfectly specular mirrors at both ends. We believe this model to suit our needs, since it exhibits a light distribution that is comparable to many real rooms. The dimensions and material properties can be found in Figure 1. Along the central vertical axis we positioned 23 measure points (one every 0.25 m) in which we computed the daylight factor on a horizontal plane.

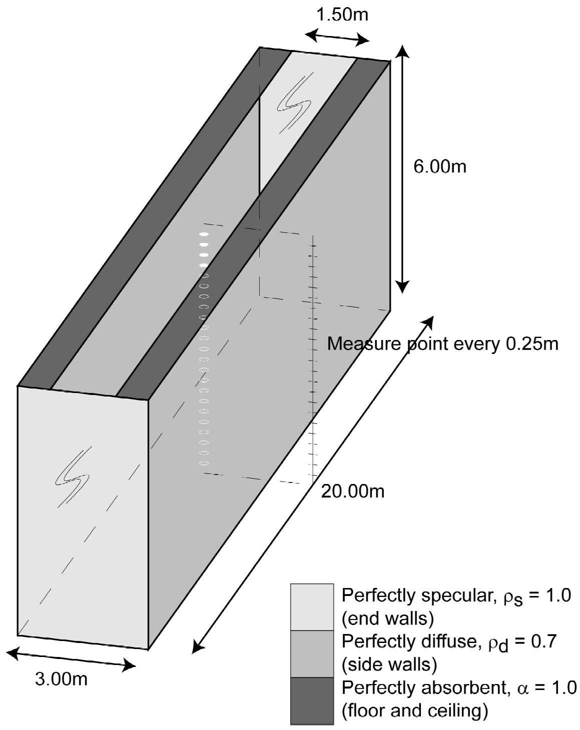

Figure 1. The experimental model to assess the impact of surface patch size on numerical accuracy.

For reference the luminance distributions across the side walls was computed mathematically. To this end the vertical surfaces were subdivided into one-centimeter horizontal strips of constant luminance. The sky components in the measure points and on the surfaces were computed analytically. So was the light transfer between any two strips of the side walls. The inter-reflections were then approximated iteratively until the daylight factor in the measure points converged to a precision of 4 decimal places.

As simulation software we used IDEA-L, a prototype developed at K.U.Leuven that uses classical radiosity to simulate daylight penetration into architectural spaces (Geebelen 2003). It uses ray casting to avoid the computation of form factors and Southwell iteration to solve the system of equations. The ray casting also takes specular material behavior into account. The program uses neither progressive, importance-driven or hierarchical patch-subdivision procedures, nor optimization techniques, so as to clearly see the effects of surface patch size. We studied seven different patch sizes: 0.5×0.5 m² (case

1), 1.0×1.0 m² (case 2), 1.5×1.5 m² (case 3), 2.0×2.0 m² (case 4), 3.0×3.0 m² (case 5), 4.5×4.5 m² (case 6) and 6.0×6.0 m² (case 7).

For comparison the daylight factors in the reference points were also computed by using the *rtrace* program from the Radiance suite. Particular care went into the definition of the daylight opening, so as to get good accuracy. The program was invoked with the following parameters: `-av 0 0 0 -ar 120 -aa 0.05 -ab 5 -ds 0.05 -dj 0.05 -dp 0 –I.`

The experiments were run on a dual-processor Intel Pentium III 1000 MHz with 256 MB RAM and the Windows 2000 operating system.

Effect on Numerical Accuracy

Surface patch size only affects the internal reflected component (IRC). However, since that component approaches zero in the top reference points, relative errors on the IRC come out too distorted to be informative. We therefore prefer to use the relative root mean square error (RMSE) of the total daylight factor according to:

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}\left(\frac{DF_{i,sim} - DF_i}{DF_i}\right)^2}{n}} \qquad (1)$$

where

DF_i = the DF in measure point *i*,

$DF_{i,sim}$ = the predicted value of the DF in measure point *i* and

n = the number of measure points.

The test results are indicated in Figure 2. Although the surfaces are coarsely subdivided into patches, the overall light distribution in the scene is still a fairly good match for the theoretical case. This reflects in the error values. Although the IRC accounts for a portion of the total DF that ranges between 0% and 35%, cases 1 through 5 have very low RMSE values: roughly up to 1% of the total DF. Cases 6 and 7 show a considerable increase in error. Similar tests with different room dimensions have indicated that the point at which this sudden escalation occurs is not related to the size of the window opening, but largely depends on the dimensions of the room. We can roughly say that the steep part of the error curve encompasses those subdivisions that have mesh widths greater than the smallest dimension of the room. The main effect of window size is that errors in the flat part decrease with increasing size.

Effect on Computation Time

As could be expected computation time exhibits an almost quadratic relationship with the number of surface patches in the scene. Times range from four seconds for case 7, of which three went into input and output routines, to 25 minutes for case 1.

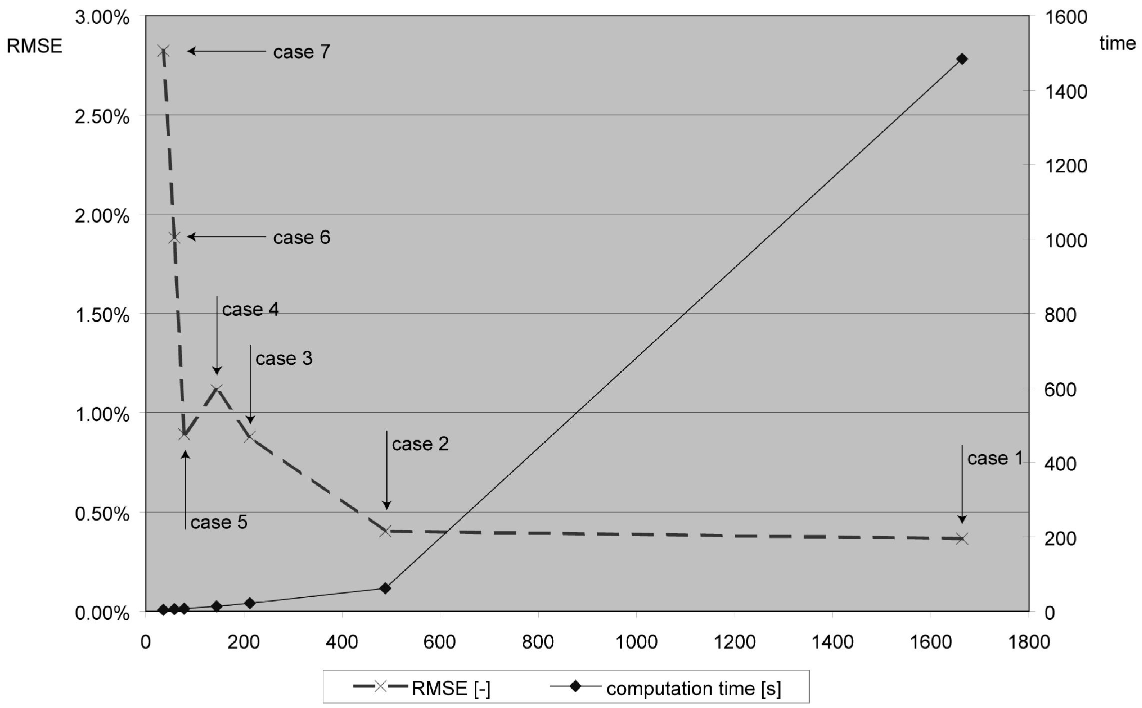

Figure 2. The RMSE of the total daylight factor and the computation time as functions of the number of patches in the scene. RMSE is in percent of the total DF; computation time is in seconds.

Discussion

This test shows that our goal, an accurate simulation within a practical computation time, can be met by manipulating surface patch size. Clearly a good equilibrium must be found between computation time and accuracy. Extensive testing has shown that for rooms with common office proportions a mesh width of one third to half of the smallest room dimension generally performs well on both accounts.

COMBINING TIME STEPS

The first adaptation of radiosity, i.e. a limit on patch subdivision, allows us to drastically reduce computation times while still maintaining good accuracy. However, this may not suffice if we want to consider the daylight conditions of each individual time step of a long-term prediction. After all, with one-hour time steps, an annual prediction involves about 4.700 instants; with one-minute time steps that number is raised to about 280.000! Even if a simulation can be performed in a minute or less, we cannot possibly perform so many simulations for a single prediction. An effective way to lower the required number of simulations is the use of daylight coefficients.

Daylight Coefficients

Tregenza first proposed daylight coefficients as an efficient way of coping with non-uniform sky-luminance distributions (Tregenza et al. 1983). The daylight coefficient for a measure point and a sky element defines the sensitivity of the illuminance in that measure point to the changes in the luminance of the sky element according to

$$\Delta E_{i,\vartheta,\varphi} = D_{i,\vartheta,\varphi}.L_{\vartheta,\varphi}.\Delta\omega_{\vartheta,\varphi} \qquad (2)$$

where

$\Delta E_{i,\vartheta,\varphi}$ = the contribution to the illuminance in measure point i due to the sky element with azimuth φ and altitude ϑ,

$D_{i,\vartheta,\varphi}$ = the daylight coefficient for measure point i and the sky element with azimuth φ and altitude ϑ,

$L_{\vartheta,\varphi}$ = the luminance of the sky element with azimuth φ and altitude ϑ

and

$\Delta\omega_{\vartheta,\varphi}$ = the angular size of the sky element with azimuth φ and altitude ϑ.

If we subdivide the sky vault into discrete solid angles and regard them as having constant luminance, we can compute the total illuminance in a measure point as follows:

$$E_i = \sum_{j=1}^{n} D_{i,j}.L_j.\Delta\omega_j \qquad (3)$$

where

E_i = the illuminance in measure point i due to the entire sky,

$D_{i,j}$ = the daylight coefficient for measure point i and discrete sky element j,

L_j = the luminance of discrete sky element j

and

$\Delta\omega_j$ = the angular size of discrete sky element j.

A daylight coefficient solely depends on the local geometry and on the material characteristics of the objects reflecting and transmitting the light on its way to the measure point. Unless one of those changes – e.g. in the case of adjustable louvers – it is invariant across all time steps.

Once the daylight coefficients have been computed, illuminance can be found under any given sky conditions as a simple linear combination. The great advantage of this approach is that it allows us to reduce the number of simulations: instead of one per time step, we can perform one per sky element, typically 145.

Errors Due to the Use of Daylight Coefficients

The principle of using daylight coefficients introduces two kinds of errors:

1. The first error is due to the fact that the sky-luminance distribution, whether it is defined by a mathematical model or by interpolation between measured values, is approximated by a discontinuous, stepped function. This error is negligible, even for skies with great luminance gradients. To test this, we computed the horizontal illuminance under an unobstructed sunny sky. For this purpose we chose a Perez All-Weather model (Perez et al. 1993a, 1993b) based on the TMY2 data for July 27th at 1PM in Atlanta, GA. The relative deviation between an estimate with a sky that was subdivided into 9824 elements, which we believe to be highly accurate, and one with a sky subdivided into 168 elements amounted to less than 0.2%.

2. The second error is due to the fact that all light originating from an entire solid angle is relocated to the solid angle's center. It is of a purely geometrical nature as it depends on the geometry of the scene and the location of the measure points. It is therefore difficult to assess. We can, however, assume it to be rather insignificant, since it will only be noticeable for cases where luminance peaks are relocated, thus causing shadows to be displaced. For

the direct component this problem can be avoided by super-sampling each sky element. For the internal reflected component the effect of slightly redistributing luminance values across the surfaces is negligible, as is demonstrated by the previous sections.

Parallel Computation

The idea to derive daylight coefficients from the simulation of an entire homogeneous sky was first proposed by Reinhart and Herkel and successfully demonstrated in their adapted version of the lighting simulation program Radiance, which uses backward ray tracing (Reinhart et al. 2000). All investigated rays are written to a file that subsequently serves as input for another program, which traces the origins of each ray back to one of the discrete solid angles in order to obtain each sky element's contribution to the indoor illuminance. The great advantage of this approach is that all daylight coefficients can be obtained from a single simulation of the inter-reflections in the scene, which can thus act as a pre-processing step to any long-term prediction.

This principle can be introduced into a radiosity algorithm even more easily. As mentioned above, one bottleneck in radiosity is the computation of how the light that is emitted by one patch will be distributed over all other patches in the scene, i.e. the calculation of form factors. Form factors are invariant to the origin and the distribution of light in the scene, and are thus identical for all daylight coefficients. It is therefore logical to perform this operation only once for all daylight coefficients.

We have adapted our radiosity program so that, instead of one luminance/illuminance pair per surface patch, each patch contains as many of these pairs as there are daylight coefficients to be computed. Each pair then holds the light originating from one sky element. The algorithm itself does not change. However, all light fluxes transferred between the surface patches of the scene, are treated as if consisting of several smaller fluxes. How these fluxes are to be redistributed needs to be worked out only once and only the very last step of performing the actual redistribution is repeated once for each sky element.

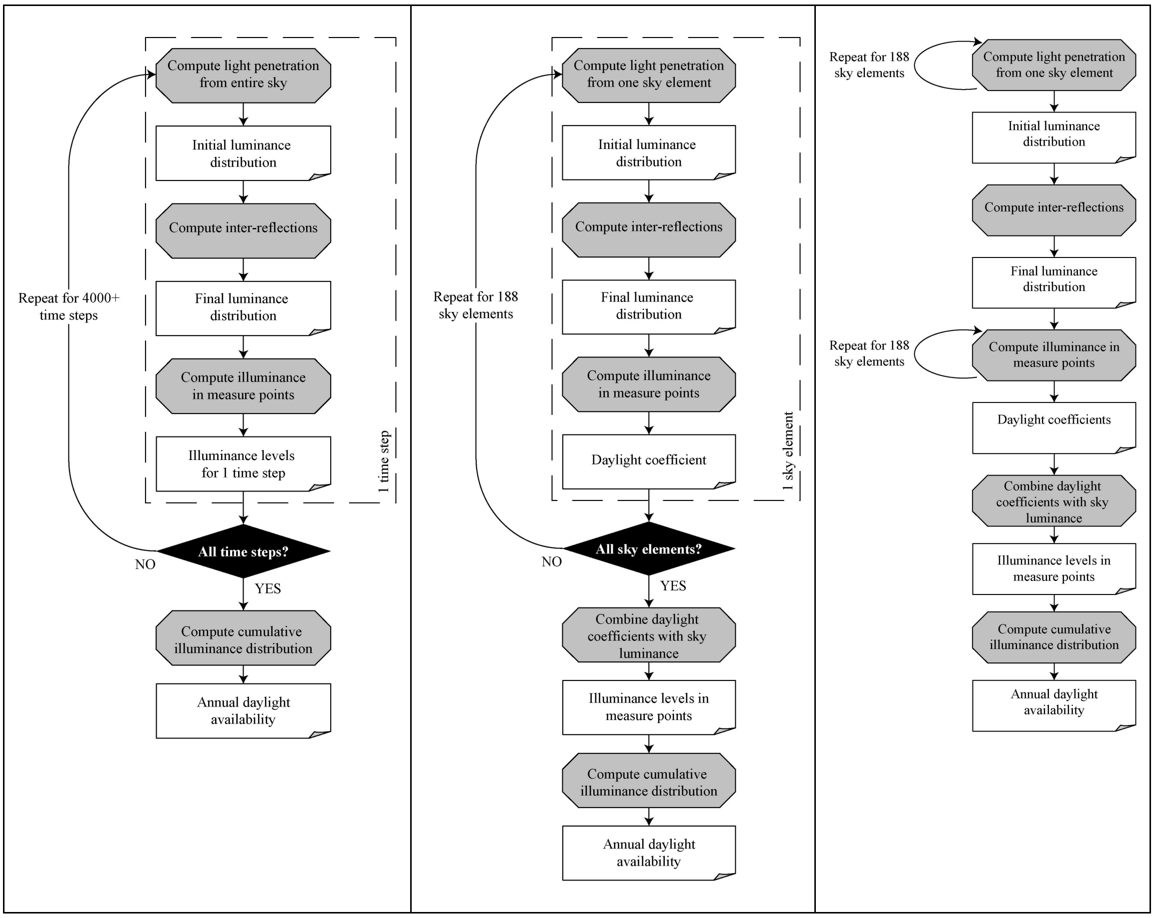

Figure 3. Flow charts representing three methods for performing long-term daylight-availability predictions while taking the daylight conditions of each individual time step into account: one simulation per time step (left), one simulation per daylight coefficient (middle), parallel computation of daylight coefficients (right).

Table 1
Comparison of the computation times needed for different simulation procedures.

	ONE SIMULATION PER TIME STEP	SERIAL COMPUTATION OF DAYLIGHT COEFFICIENTS	PARALLEL COMPUTATION OF DAYLIGHT COEFFICIENTS
Computer	Dual-processor Intel Pentium III 1000Mhz, 256Mb RAM, Windows 2000 (the algorithm discussed here has no built-in parallel-processing features)		
Number of patches in model	128		
Number of measure points	35		
Number of sky elements	N/A	188	188
Time for computing daylight coefficients [s]	N/A	306	11
Time for computing illuminance values [s]		8	8
Time for other tasks [s]		1	2
Total computation time [s]	**17696 (or 4 hr, 54 min and 56 s)**	**315 (or 5 min and 15 s)**	**21**

The only drawback of this procedure is the increased memory use. In our implementation this amounted to an additional 256 bytes per surface patch and per daylight coefficient. However, for most contemporary hardware this will not be a problem.

The parallel computation of daylight coefficients influences accuracy only marginally. Most radiosity algorithms solve the matrix equation in an iterative way – e.g. with Southwell relaxation – until a pre-defined threshold is met. If each daylight coefficient is computed individually, no more iteration steps will be performed than are necessary to meet the threshold for each separate daylight coefficient. When using parallel computation, the iterative process will be maintained until the threshold is met for all daylight coefficients. This implies that more iteration steps will be performed than are absolutely necessary for most daylight coefficients. If anything, this will improve accuracy.

The great advantage of this approach is of course the reduction of computation time. As an example we took the model of a typical rectangular office space with two windows and computed the daylight autonomy. The results can be found in Table 1. It is obvious that both the use of daylight coefficients and the parallel computation principle lead to great improvements.

CONCLUSION

Other authors have pointed out that the prediction of cumulative daylight availability should take the daylight conditions into account of all individual time steps. However, with most contemporary simulation programs this guideline leads to computation times in the order of days or even months, which is clearly highly impractical.

We have proposed two principles that can be applied to any radiosity algorithm. Both have only a minimal effect on the accuracy of the calculated result, but when combined they reduce the computation time for a cumulative prediction to the order of seconds. This makes annual and seasonal predictions far more accessible to the design process, so that the daylight availability of design alternatives can be more easily assessed and compared.

The studies discussed above have allowed us to draw the following conclusions.

1. Radiosity has proven to be a highly efficient algorithm for estimating the light distribution in the scene. Even with very coarse meshes the resulting illuminance levels show relatively little error. Moreover, if one renounces the ambition of producing realistic images, computation times can be reduced considerably. A comparison with Radiance – a ray tracer that is regarded in the academic world as highly reliable – seems to corroborate the hypothesis that radiosity's fundamental principles make it more efficient for approximating the way in which light is spread throughout a scene than its counterpart ray tracing. The argument that only diffuse material properties can be simulated is obsolete and only applies to classical radiosity. Moreover, it must be noted that the described experiments were performed with a rather crude implementation of radiosity. The introduction of known optimization techniques would only reduce computation times even more, thus consolidating the conclusion.

2. The importance of the number of patches in the scene for computation time and its fairly small impact on the resulting illuminance levels emphasize the significance of sensible modeling. Scene

descriptions should be devoid of unnecessary detail. This is the shared responsibility of the user, who should be trained in the art of daylighting, and of the developers of modeling tools, which should guide the user in producing good models. Automatic or semi-automatic procedures for reducing detail in an intelligent way merit future research.

3. The tests scenes did not include any façade elements other than plain glazing. The introduction of systems such as Venetian blinds or anidolic mirrors may have a considerable impact on computation time and possibly render the effect of limited patch subdivision less significant. Future research will have to point out whether the addition of more elaborate transmitting, reflecting or refracting material behavior could allow more complex daylighting systems without raising the number of surface patches.

4. The use of daylight coefficients, especially when combined with parallel computation, offers an effective way of performing long-term simulations that consider the conditions of all individual time steps within practical computation times. Should future research indicate the importance of time steps shorter than 1 hour, simulation times will no longer be an obstacle.

As test bed we deliberately chose an algorithm without any optimization measures, so as to isolate the effects of the principles presented in this paper. Inevitably the choice of Southwell relaxation and ray casting affect the reported accuracies and simulation times. However, the discussed theories depend in no way on either technique, which implies that they are applicable to all radiosity variants. The introduction of techniques such as hierarchical or importance-driven meshing, spatial partitioning or fast computation of form factors may slightly attenuate the improvements due to limited patch subdivision, but there is no reason to assume that the conclusions of the tests will be affected.

ACKNOWLEDGMENT

This research was financed by project OT/00/18 of the K.U.Leuven Research Fund.

REFERENCE

Erhorn, H., de Boer, J., Dirksmöller, M. 1998. ADELINE – An Integrated Approach to Lighting Simulation. In: *Proceedings of Daylighting ´98*, Ottawa, Natural Resources Canada, 21-28.

Geebelen, B. 2003. *Daylight Availability Prediction in the Early Stages of the Building Design Process*, PhD thesis, K.U.Leuven, Leuven.

Glassner, A.S. 1995. *Principles of Digital Image Synthesis*, Morgan Kaufmann, San Francisco.

Herkel, S., Pasquay, T. 1997. Dynamic link of light and thermal simulation: on the way to integrated planning tools. In *Proceedings of the 5th International IBPSA Conference*, Prague, IBPSA, 307-312.

Perez, R., Seals, R., Michalsky, J. 1993a. All-Weather Model for Sky Luminance Distribution – Preliminary Configuration and Validation. In: *Solar Energy* **50**(3), 235-245.

Perez, R., Seals, R., Michalsky, J. 1993b. Erratum to All-Weather Model for Sky Luminance Distribution – Preliminary Configuration and Validation. In: *Solar Energy* **51**(5), 423.

Reinhart, C., Herkel, S. 2000. The simulation of annual daylight illuminance distributions – a state-of-the-art comparison of six Radiance-based methods. In: *Energy and Buildings* **32**, 167-187.

Tregenza, P., Waters, I. 1983. Daylight Coefficients. In: *Lighting Research & Technology* **15**(2), 65-71.

Ward Larson, G., Shakespeare, R. 1998. *Rendering with Radiance – The Art and Science of Lighting Visualization*, Morgan Kaufmann, San Francisco.

Winkelmann, F., Selkowitz, S. 1985. Daylighting simulation in DOE-2: theory, validation and applications. In: *Proceedings of the Building Energy Conference*, Seattle, WA, 326-336.