Chapter

3 *Validation I: Preparation*

"Entirely by ironic coincidence, on the day I wrote this I received a postcard from an art gallery which somehow got my name on its mailing list. It was to announce a new one-man show entitled ***Primus corpus: DNA License****. A Sample of the artist's work was included. It consisted of a small plastic bag stamped 'Universal Notice - Only One - Original Human' containing some hair trimmings. Anyone struck dumb by this achievement is unlikely, I daresay, to be among the admirers of this chapter."*

NORMAN LEVITT (IN *PROMETHEUS BEDEVILLED*)

Recent advances in computer graphics techniques allow, in principle, the modelling of realistic architectural scenes for visualisation and illuminance prediction [Sillon 94, Ward 94]. Validation studies of these new programs have, to date, been of restricted value, one reason being that comparison against scale models measured in artificial skies are made using necessarily idealised sky brightness distributions [Selkowitz 82]. Also, where illuminance predictions have been compared with measurements taken in real rooms under real sky conditions [Bellia 94], the sky brightness distribution used by the program was based on a theoretical sky model generated from bulk values e.g. global and diffuse horizontal illuminance. Differences between the real sky luminance distribution and that used in

the program are not known. It is therefore impossible to determine where the errors arise; in the basic algorithms or the representation of the sky.

This chapter describes the preparation for a validation study that offers a considerable advance on previous work. In this study, illuminance predictions were compared with measurements taken in full size office spaces under real sky conditions. The simulation program used model sky luminance patterns that were based directly on measured sky brightness distributions. The uncertainties in model representation, that had limited the findings of earlier studies, were greatly reduced for this validation. It was possible therefore, to make a reliable evaluation of the absolute accuracy of the program under naturally occurring daylight conditions.

Section 3.1 describes the composition of the validation dataset. That section also discusses to what degree the validation dataset is representative of the full range of naturally occurring sky conditions for the UK. The lighting simulation models for both the luminous environment (sun and sky) and the office space are described in Section 3.2. The processing of the sky luminance measurements to a form compatible with the simulation program is also described. The section ends with a hypothesis concerning potentially unreliable sky-photocell combinations. In Section 3.3 the lighting simulation itself is described. That section includes a methodology for the setting and optimization of the ambient calculation parameters. The results of the validation are presented and analysed in Chapter 4.

3.1 The validation dataset

The first steps towards constructing a definitive world atlas of daylight availability were made when the Commission Internationale de l'Eclairage (CIE) organised the International Daylight Measurement Programme (IDMP). A major objective of the programme was to collect long-duration time-series data for a range of daylight parameters, including, at the stations designated 'research class', measurement of the actual sky brightness distribution together with integrated quantities. The IDMP has

coordinated the activities of 15 such 'research class' monitoring stations around the globe, the majority of which attempted to achieve continuous monitoring over a period of a year or more. One of the two UK 'research class' stations was at the Building Research Establishment (BRE), Garston, UK.

In conjunction with the sky monitoring programme, the BRE conducted an evaluation study of the light redistribution properties of five innovative glazing systems against standard clear glazing. The sky monitoring apparatus were positioned on the roof directly above the experimental rooms. Room illuminance and sky luminance measurements were recorded within seconds of each other. From matched samples of data from these two measurement programmes, a database for the validation of lighting simulation programs was constructed. This is referred to herein as the BRE-IDMP validation dataset.

3.1.1 Measured quantities and site details

The site details for the BRE station were as follows[1].

Station Location:	Latitude: 51˚43' N Longitude: 0˚22' W
Height above sea level:	80m
Operation:	Started[2] on July 16, 1992, ended on July 1, 1993.

The external quantities monitored and the measuring instruments used were as follows:

- Illuminances

Global horizontal:	LMT BAP 30 FCT
Diffuse horizontal:	LMT BAP 30 FCT
North vertical:	LMT BAP 30 FCT
East vertical:	LMT BAP 30 FCT
South vertical:	LMT BAP 30 FCT
West vertical:	LMT BAP 30 FCT
Direct normal:	Eppley, Solar Tracker

1. Information obtained from IDMP web-server **http://idmp.entpe.fr/**
2. Measurements were made available from earlier in 1992.

- Irradiances

Global horizontal:	Kipp & Zonen CM11
Diffuse horizontal:	Kipp & Zonen CM11
North vertical:	Kipp & Zonen CM 5
East vertical:	Kipp & Zonen CM 5
South vertical:	Kipp & Zonen CM 5
West vertical:	Kipp & Zonen CM 5
Direct normal:	Eppley, Solar Tracker

- Others

Sky luminance:	PRC Krochmann Sky Scanner (15 mn)
Dry Bulb Temperature:	Vaisala HMP 132Y
Relative Humidity:	Vaisala HMP 132Y

The instrumentation layout on the roof of the BRE office block (Building 9) and the obstructions to the view above the horizon are shown in Figure 3-1.

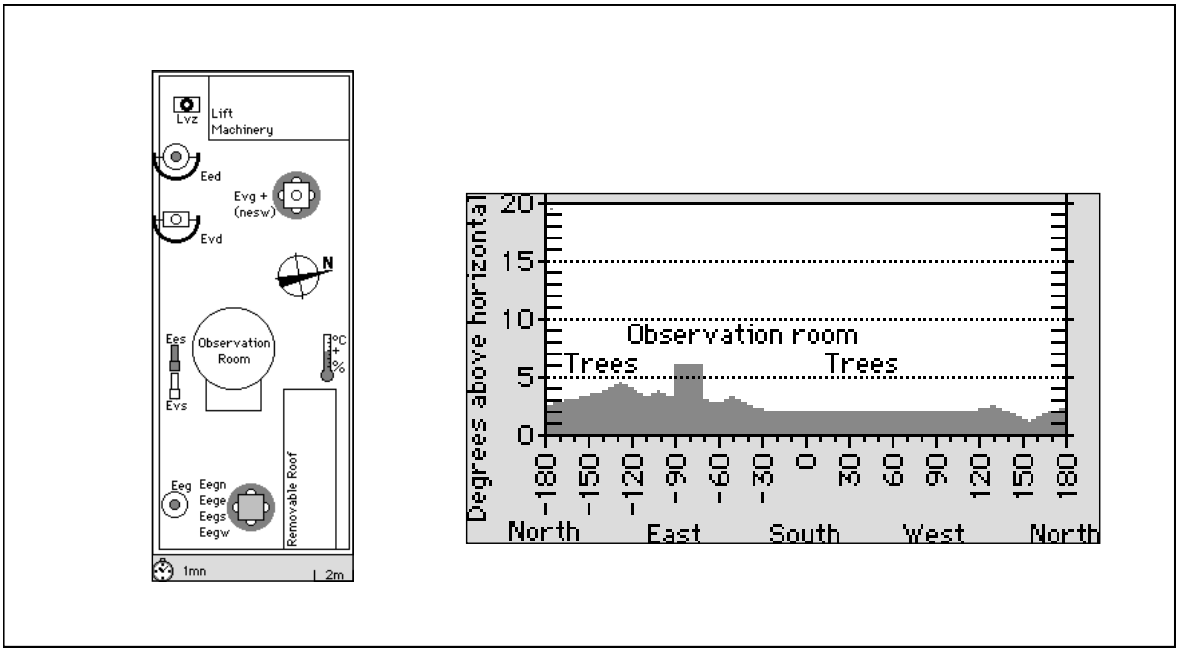

Figure 3-1. Instrumentation layout and obstructions to view above horizon

3.1.2 Internal conditions: illuminance measurements

Two full-size mock offices, adjacent to each other and with south-facing glazing were set up by the BRE on the top storey of Building 9, Figure 3-2.

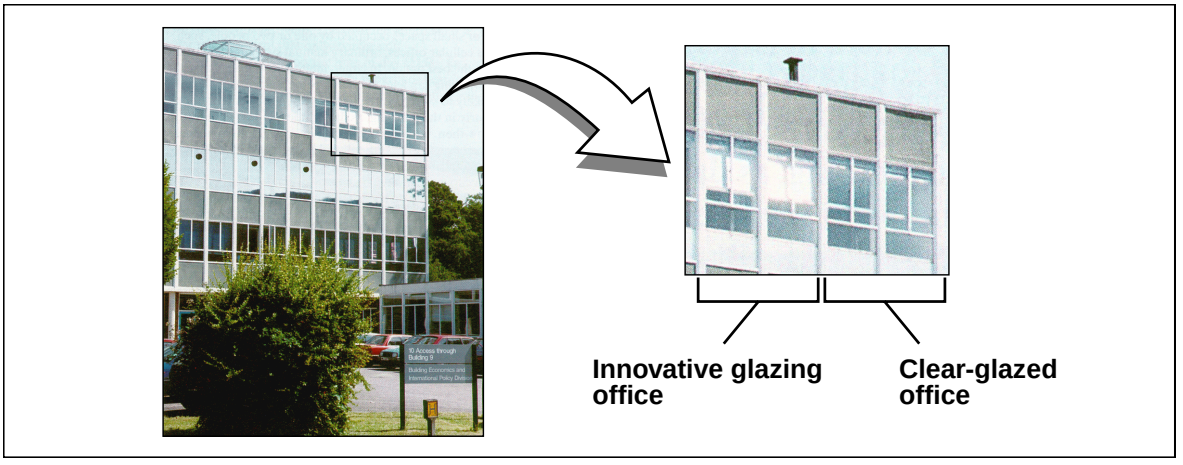

Figure 3-2. Building 9 with inset showing test offices

Room dimensions were almost identical; 9 metres deep, 3 metres wide and 2.7 metres high. The rooms were left unfurnished, though the surface reflectances were chosen to correspond to a typical office. The window of one office was adapted so that an innovative daylighting system could be installed, the other had conventional single glazing, Figure 3-3.

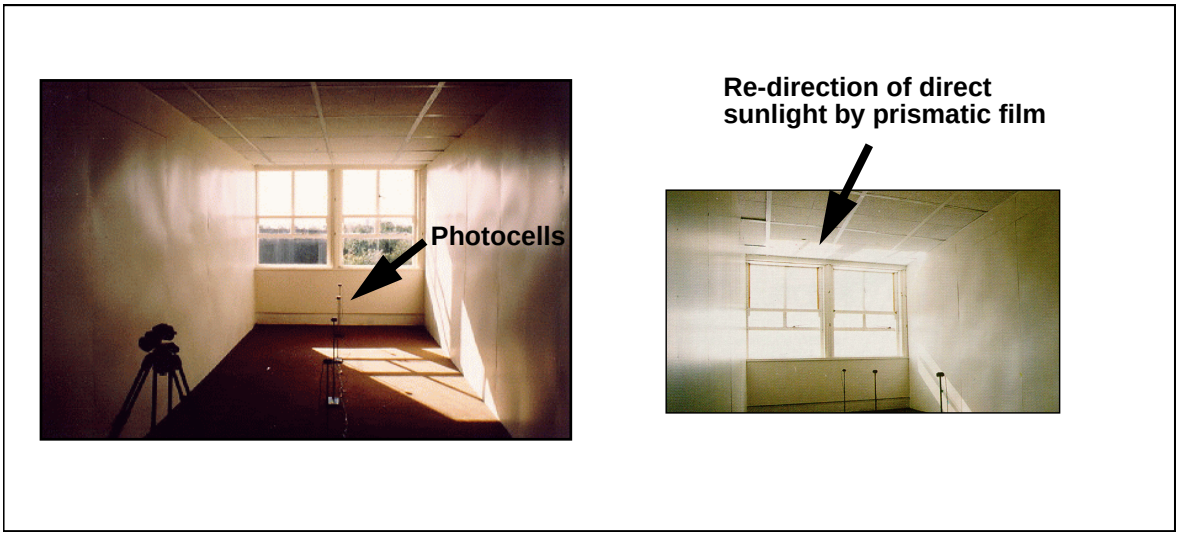

Figure 3-3. Photographs of the BRE office rooms (a) single glazing and (b) innovative glazing

Six illuminance cells positioned at work plane height (0.7m), regularly spaced along the centre line of each room, were used to monitor the illuminance distribution in the room, Figure 3-4.

The innovative glazing systems used in the BRE study were: diffuse and mirror finish light shelf, 3M prismatic film, Siemens prismatic glazing and Okalux mirrored louvre. Each system was evaluated for a period of about six weeks close to an equinox (23 September, 21 March) and again for two shorter periods during summer and winter months [Aizlewood 93]. The innovative glazing systems were installed in turn in one of the two office spaces. The other mock office had conventional single glazing throughout the entire monitoring period.

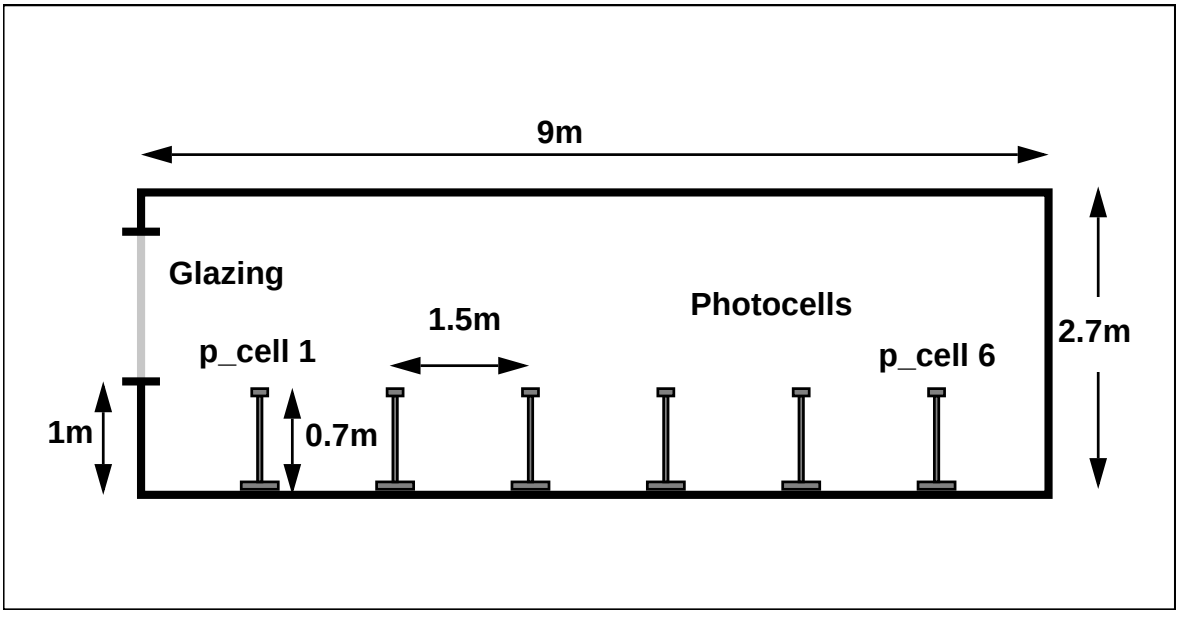

Figure 3-4. The BRE test cell

File formats

The internal illuminance data were obtained as ascii files, one for each day of monitoring. The illuminance measurements for the two mock offices were supplied as 5 minute averages of 1 minute data. For these files, each record contained 16 entries: time, 6 illuminance measurements (innovative glazing

office), 6 illuminance measurements (single glazing office), and sun position. A sample is shown in Table 3-1.

Hour	Min.	Illuminance measurements for innovative glazing office P_cell 1 - 6						Illuminance measurements for single-glazed office P_cell 1 - 6						Alt.	Azi.
12	45	21259	19495	3514	2259	1724	1592	19374	20454	18787	3814	1751	1709	18.7	195.3
12	50	20904	18689	3561	2367	1733	1594	17990	20335	18441	9499	1749	1698	18.4	196.5
12	55	20725	18803	3677	2547	1785	1620	19332	20843	18675	5240	1784	1726	18.2	197.8
13	0	21238	17485	3716	3585	1780	1604	19102	20056	12508	19291	1796	1727	18.0	199.0
13	5	21006	15742	3756	11969	1776	1598	18842	19560	7461	18953	1816	1733	17.7	200.2
13	10	20644	16832	3854	15513	1773	1579	18805	19219	18513	18569	1826	1728	17.4	201.4
13	15	20780	15221	3923	15623	1758	1566	18682	18852	18325	18213	1818	1711	17.1	202.6

Table 3-1. Sample from file for day 318_92.csv

3.1.3 External conditions: monitoring the sky and sun

The instrument used to measure the sky brightness distribution was a PRC Krochmann sky scanner, Figure 3-5. The sky scanner measured the sky luminance distribution every 15 minutes during daylight hours. Each scan consisted of 150 readings according to the pattern recommended by the CIE [Perez 91] and took approximately 25 seconds to complete. Of the 150 measurements taken, 145 were for unique positions on the sky vault (the zenith luminance was recorded 6 times during each scan). The scanner acceptance angle was 11˚ giving a sky coverage of ~68% [Tregenza 87], Figure 3-6. The scanner did not measure the sky luminance at the position closest to the sun, and a scan could contain one or more occurrences of 'out of range' measurements.

File formats

The global quantities file contained the basic illuminance and irradiance data together with a few environmental parameters. These quantities were five minute averages of one minute data. Each record in the file contained 22 entries, Table 3-2. An example few lines from a global quantities file is given in Table 3-3.

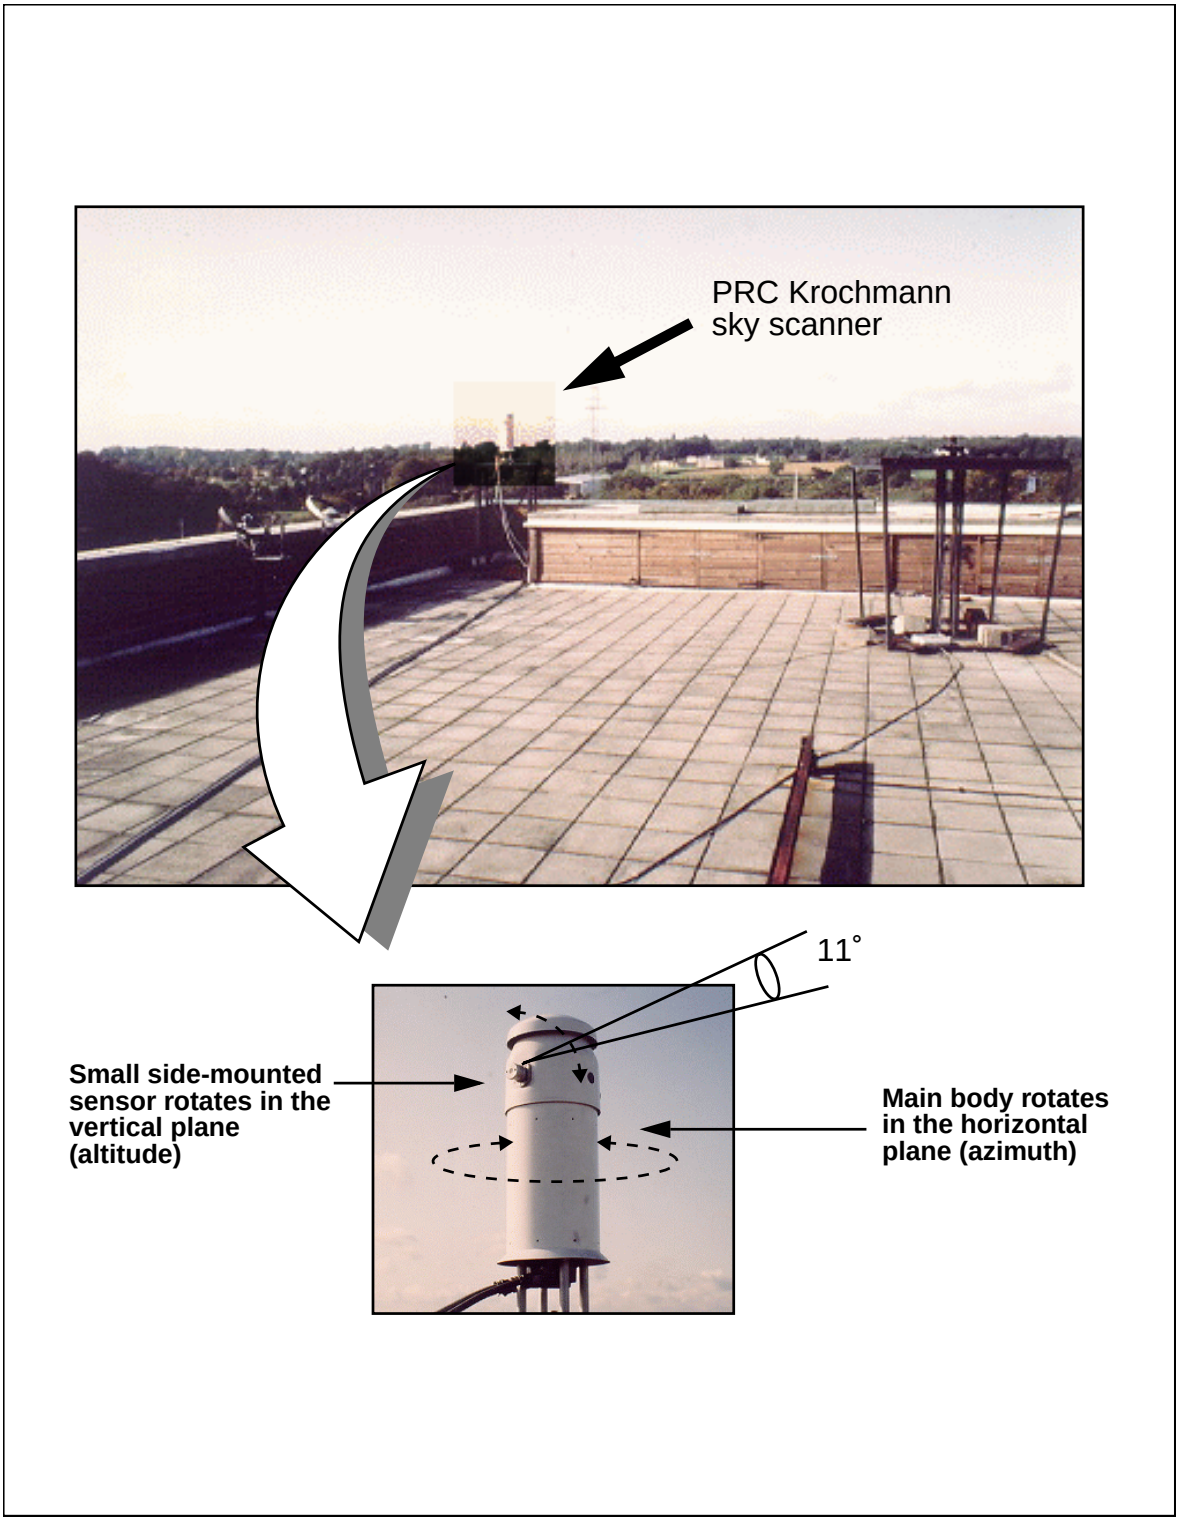

Figure 3-5. Krochmann PRC sky scanner positioned on the roof of the BRE lighting laboratory building and detail

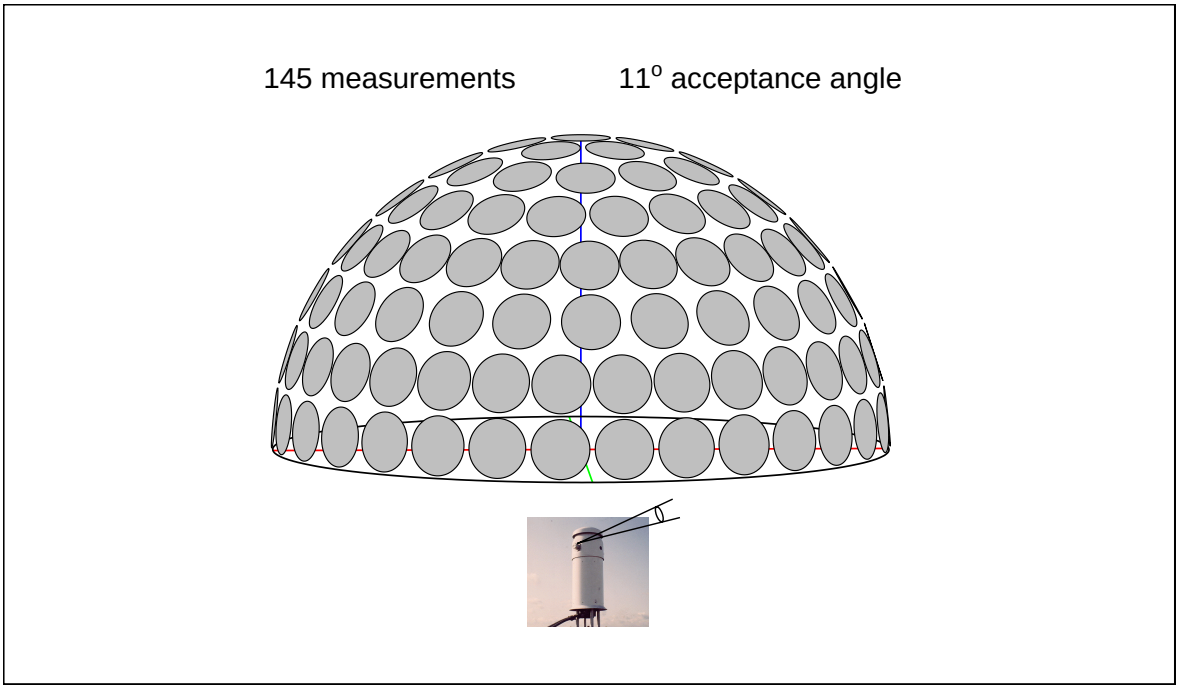

Figure 3-6. Sky scanner measurement pattern

Quantity	Units
Error code	-
GMT	hr., min.
Horizontal global illuminance	lux
Horizontal diffuse illuminance	lux
Vertical total illuminances (north, east, south, west)	lux
Direct normal solar illuminance	lux
Zenith luminance	cd/m^2
Dry bulb temperature	˚C
Relative humidity	%
Horizontal global irradiance	W/m^2
Horizontal diffuse irradiance	W/m^2
Vertical total irradiances (north, east, south, west)	W/m^2
Direct normal solar irradiances	W/m^2
Solar altitude (above horizon)	deg. (˚)
Solar azimuth (clockwise from due north)	deg. (˚)

Table 3-2. Global quantities file

SkyLog: V 3.0 (c) Copyright Cambridge Consultants (SE) Ltd. 1991

SL318_92.ILR

1992.

318.

	Time		Horiz	Diff	VertN	VertE	VertS	VertW	Direct	Zenth	T/Al	Hu/Az
	-----	-------	-------	------	-------	-------	-------	-------	-----	----	-----	
0	10	10	26564	10750	5070	31229	62109	5914	48233	1592	4.1	100.0
			254.27	84.56	37.78	298.31	624.20	44.54	495.19		17.0	156.6
0	10	15	27179	10863	5107	30366	63406	6019	49067	1617	4.2	100.0
			260.58	85.63	38.01	289.57	635.44	45.37	501.16		17.3	157.8

Table 3-3. Sample from global quantities file sl319_92.alr

The measured sky luminance distribution data were instantaneous values recorded at 15 minute intervals. Each record contained 155 entries: an error flag, the time of observation, sun position, 144 measurements of the sky luminance distribution, and six measurements of the zenith luminance (see Figure 3-6 for a graphic of the measurement pattern). A sample file showing the measurements taken at 13h00 for day 102 in year 1992 is given in Table 3-4.

32713										Flag
102	92	13	00	45.5	201.4					Day, year, hr., min., altitude, azimuth
24444	22900	23438	19435	14154	12554	9607	8512	7313	6258	150 measurements of the sky luminance
5912	5393	5395	5345	5108	5237	5248	4942	5077	5241	
5541	6138	6887	8088	10016	11831	15468	18043	21404	26991	
30099	27793	30696	24041	16433	14122	10613	9126	7689	6594	
5819	5665	5459	5368	5343	5241	5267	4948	5081	5226	
5619	6254	7098	8484	10755	12614	17844	20959	25620	36961	
42601	41862	30333	21104	14797	10499	8020	6461	5522	5001	
5083	4834	4775	4889	4494	4866	5433	6307	7851	10320	
13690	20758	31728	41523	99999	40227	33260	21457	15445	10945	No measurement at the sun position
8079	7073	5684	5492	4682	5119	4469	4653	4669	4781	
5283	6267	8413	10468	14618	21543	31644	51516	99998	33055	Out-of-range
25636	15314	10124	7482	6419	5091	4665	4830	5828	6108	
5566	8138	10582	15487	21929	31264	36910	27183	14230	9113	
7699	5319	5984	6809	6482	10328	15479	19513	20161	12472	
9634	8309	8060	14219	12571	9242	12078	10854	10238	13932	Six zenith measurements

Table 3-4. Measured data for case 102_92_13h00

3.1.4 Comparison of the validation dataset composition with the Kew TRY

Data files comprising 27 days monitoring from the year 1992 were provided by the BRE. The days supplied were pseudo-randomly scattered throughout the year, Figure 3-7. Due to the presence of a large tree east of the site -

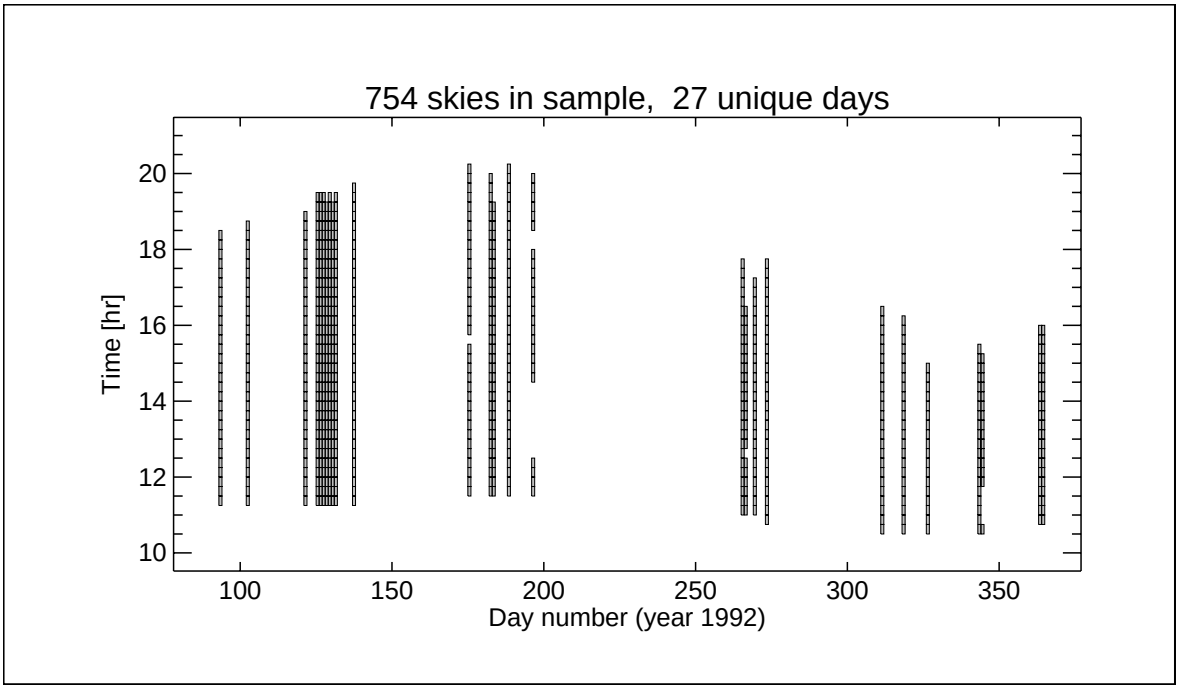

Figure 3-7. Distribution of validation dataset samples from the year 1992

which can cause shadowing on the windows of the mock-office (Figure 3-1) - all illuminance levels measured with the solar azimuth at less than 160˚ were removed from data by the BRE as part of their quality assurance procedures. The distribution in sun azimuth and altitude for the validation dataset is given as a two dimensional frequency histogram, Figure 3-8. The bin width for the altitude and the azimuth angles was 5˚ because this was roughly commensurate with the 15 minute timestep for the scanner measurements. For comparison, the distribution in sun position that would occur for an entire year (at 15 minute intervals) at the validation site is given also. The absolute numbers are of course very different: for one year (at 15

minute intervals) there were 17,635 occurrences of the sun altitude greater than 0˚, as opposed to 754 entries in the validation dataset. However, to reveal any bias that may exist for the sun positions in the validation database, each frequency map was normalised to maximum = 1 (see legend Figure 3-8). The distribution plots show that most of the actually occurring sun positions (for sun azimuth ≥ 160˚) were represented to a greater or lesser degree in the validation database. There is an arc of empty bins in the distribution that was due to the lack of samples from around either day 70 or day 290. At this stage, there is no reason to believe that this deficiency will have any significant bearing on the outcome of the validation.

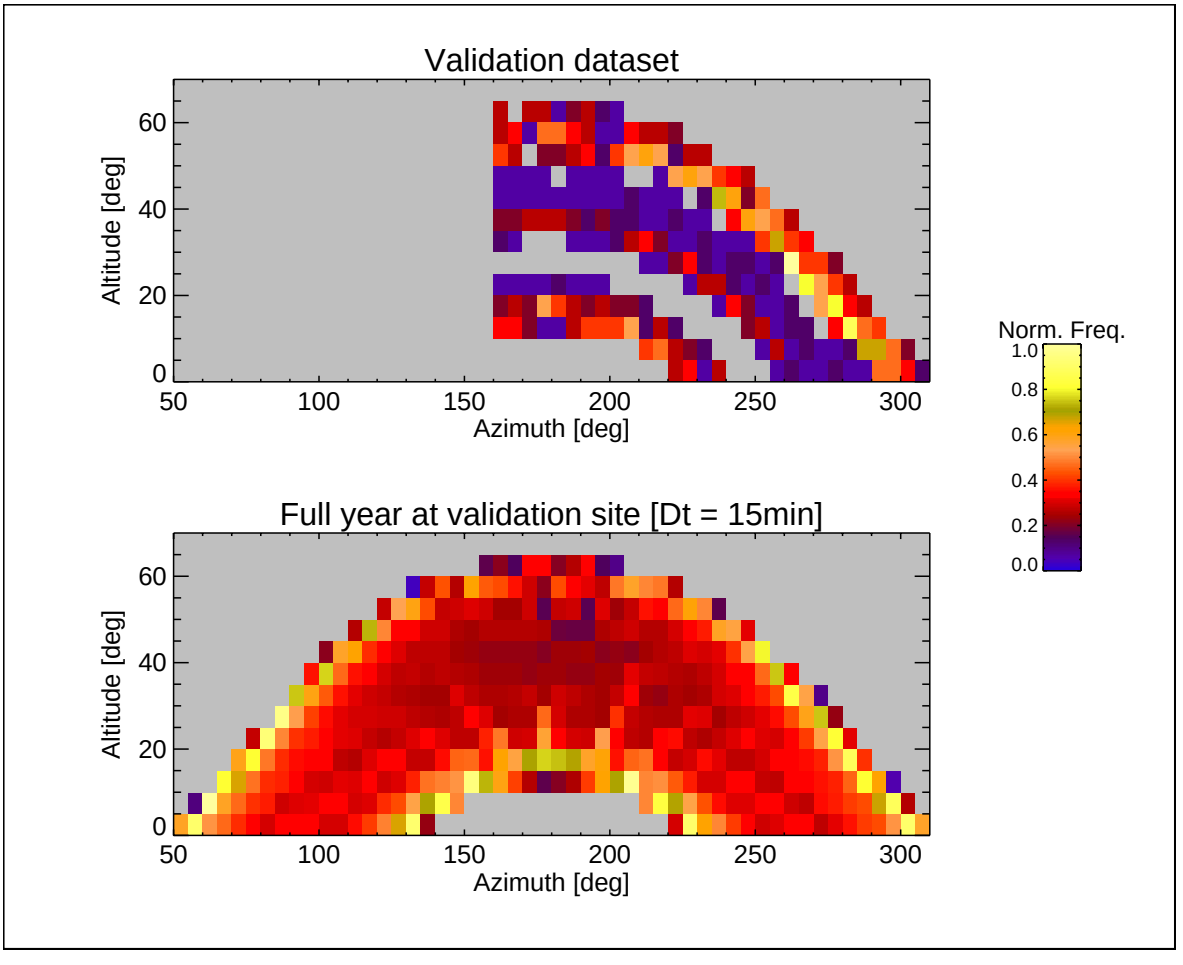

Figure 3-8. Distribution in azimuth and altitude for validation database and entire year

Just how representative these 754 skies were of the full range of naturally occurring sky conditions (clear, overcast etc.) in the UK can be judged from Figure 3-9. Here, the distribution in the sky clearness index for the

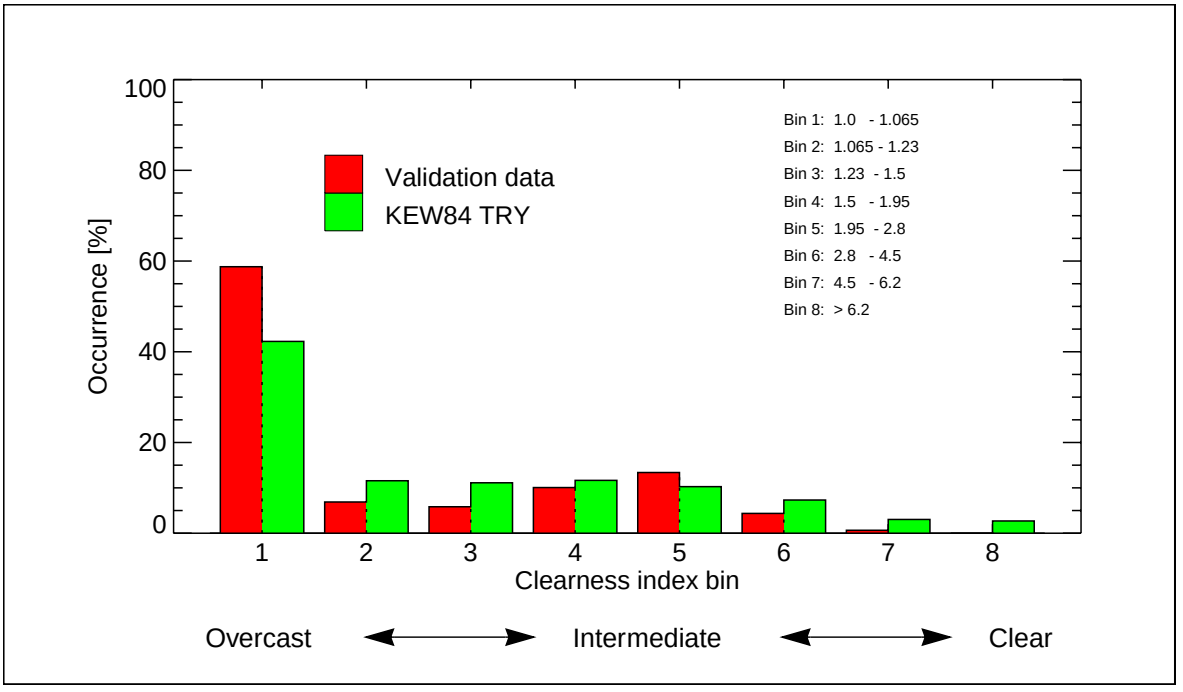

Figure 3-9. Distribution in clearness index compared to TRY

validation dataset and for a standard test reference year (TRY) are compared. The TRY data were recorded at Kew which is close to the validation site. The TRY time-series contains hourly measurements of the diffuse sky irradiance and the direct normal solar irradiance[3] for one year. The distribution in sky types for the validation dataset was broadly similar to that for the TRY. In the validation data, heavily overcast skies (bin 1) were somewhat over-represented whilst the very clearest skies were under-represented. The clearness index, ε, is given by [Perez 90]:

3. Irradiances were converted to illuminances using a constant value for luminous efficacy.

$$\varepsilon = \frac{\left[\left(\frac{E_{dh} + E_{dn}}{E_{dn}}\right) + 1.041Z^3\right]}{1 + 1.041Z^3} \quad \textbf{(3-1)}$$

where E_{dh}, E_{dn} and Z are the diffuse horizontal illuminance, the direct normal illuminance and solar zenith angle respectively.

3.1.5 Scope of the validation

The fixtures in the innovative glazings room were cycled throughout the monitoring period. So it was the clear glazing office that was exposed to the largest number of skies. Accordingly, the all-skies (754) validation exercise was carried out using this window type. Note that clear glazing is used for the majority of existing and new buildings in the UK. In a limited study, the diffuse and specular light shelves were modelled (see below). The other three innovative glazings - Okalux mirrored louvre, Siemens prismatic glazing and 3M prismatic film - were excluded from the validation because the optical transmission properties of these materials had not been measured. The *Radiance* program has the capability to model in detail the bi-directional reflection transmission distribution function (BRTDF) of a material, and it would be possible to extend the validation to include these materials if and when the BRTDF data becomes available.

3.2 The lighting simulation models

3.2.1 The office model

Geometrically, the office model created for the simulations was a very close representation of the experimental office. The dimensions of the clear glazed office room were measured by hand to an accuracy of ~1cm, and the space was described in the model as a collection of rectangular polygons. Particular attention was paid to the window bars and glazing panes which were measured to an accuracy of ~0.2cm and modelled as discrete elements. The illuminance meters themselves were not modelled, rather the horizontal illuminance at that point was calculated. All opaque surfaces

were modelled in the first instance as achromatic diffuse reflectors, although it is apparent from the photograph of the room that the paint used for the walls has a small specular component (Figure 3-3). The reflectances used in the model were the average of the values measured at the beginning and end of the monitoring period: walls 0.83, ceiling 0.80 and carpet 0.095 [Aizlewood 93]. Window transmittance was that for standard single glazing. A glazing maintenance factor of 5% was incorporated into the transmittance.[4] A rectangular ground plane of size 9 x 10m and reflectivity 0.15, was placed at ground level in front of the glazed facade of the office. This was the only non-luminous external object in the model. A line drawing created directly from the *Radiance* scene description for the single-glazed office room is shown in Figure 3-10.

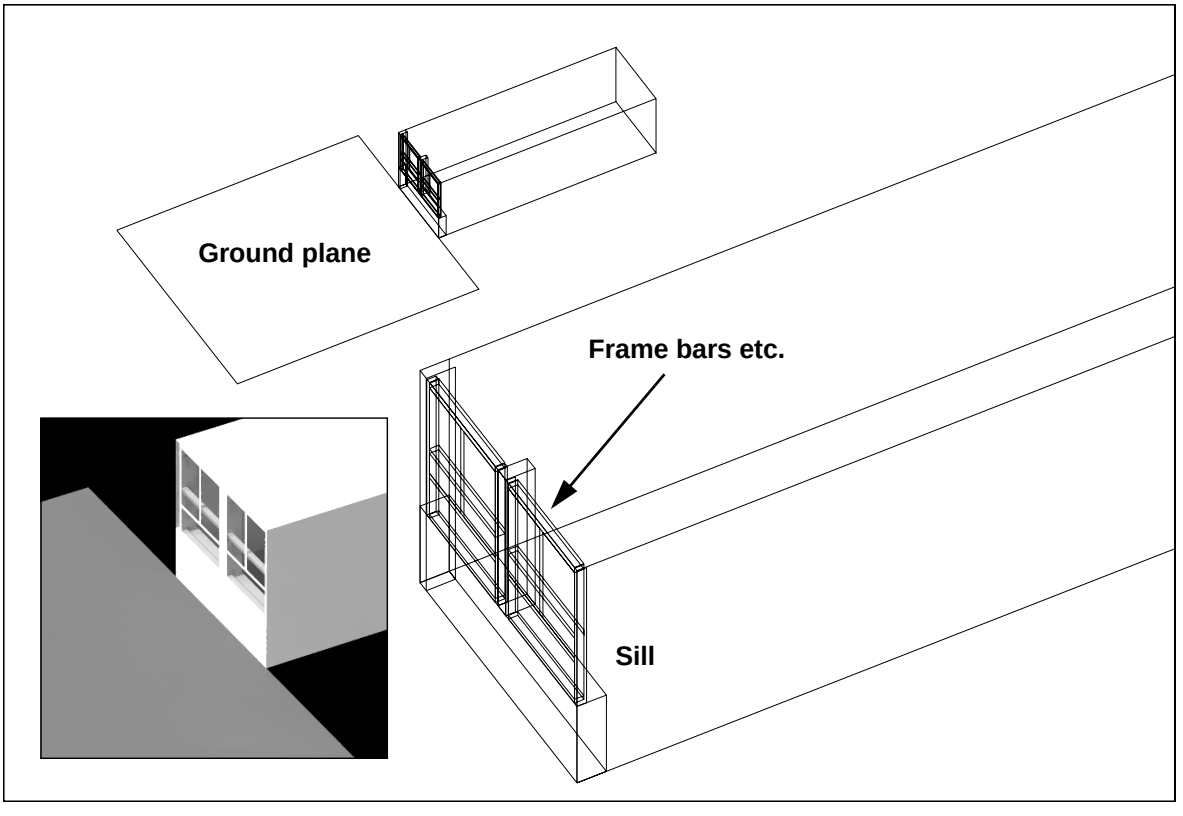

Figure 3-10. Line drawing and rendering of office scene description

4. Private communication - M. Aizlewood, BRE.

The experimental rooms were on the third storey of Building 9, whereas the ground plane in the model was placed at the same level as the office floor. The justification for this is as follows. The 3D model, necessarily, had to be an incomplete description of reality, and it was decided at an early stage that it should be as simple as possible.[5] For light transfers from the ground plane into the office, the 'view' of the ground plane from the ceiling just inside the office window is a major factor. In this respect, a small ground plane at the same level as the office floor functions in much the same way as a larger ground plane with the office placed above it, Figure 3-11.

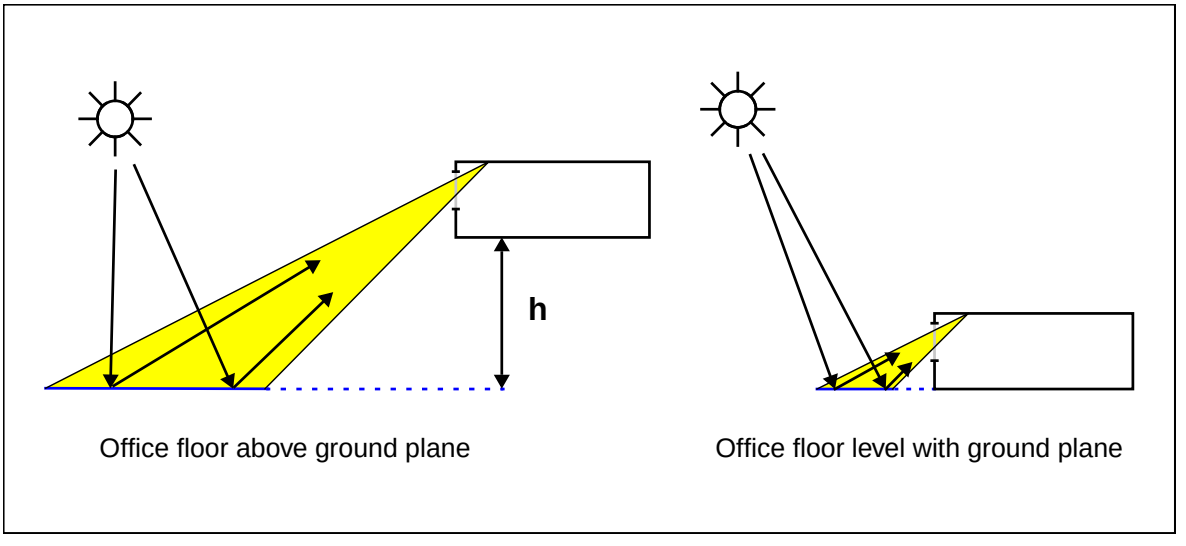

Figure 3-11. Simplified ground plane model

Furthermore, there are good reasons to prevent the maximum scene dimension from getting too large (see Section 2.5).

In the first published results of this work [Mardaljevic 95], a circular ground plane of radius 30 metres and centred on the room was used. With this ground plane, it was discovered that the inter-reflection calculation

5. Ockham's Razor, a principle urging the use of the most economical and least complex assumptions, is, in its original phrasing, particularly apt: "Entities should be not multiplied unnecessarily".

expended some effort in predicting the luminance of the external walls and adjacent ground plane. The luminance of the (external) side and rear walls however had negligible effect on the internal illuminance. Therefore the circular plane was replaced with a 'front-facing' rectangular ground plane.

In the limited study, two of the five innovative glazing fixtures were also modelled. These were internally mounted diffuse and specular (mirror) finish light shelves. Both shelves were the same size: full room width, 1.00 metre deep and fixed at a height of 2.08m. The diffuse finish light shelf was coated with a paint similar to that used on the ceiling and so was a assigned a reflectivity of 0.80. The upper surface of the specular shelf, in reality a polished aluminium sheet, was modelled as a mirror having a reflectivity of 0.90. Some uncertainty exists here: specular light shelf reflectivity was not directly measured and the value used in the model was based on typical value for this material.

3.2.2 The sun and sky models - generic form in the simulation

For lighting simulation, a model scene is constructed using various 'surface primitive' types (e.g. sphere, polygon, ring) and the illumination is provided by making one or more of these entities self-luminous. For the validation scene, there were two sources of (daylight) illumination - the sun and the sky. These were represented in the model using a special type of surface called **source**. A source is not really a surface, but a solid angle. And as such, a sun or sky described using source is effectively infinitely distant from the rest of the (finite) model scene. The source primitive has the basic format

```
mod source id
0
0
4 xdir ydir zdir angle
```

The arguments **xdir**, **ydir** and **zdir** give the direction to the centre of the source and **angle** is the number of degrees subtended by its 'disk'. A

schematic illustrating the extent and orientation of the source angles for the sky and sun is given in Figure 3-12.

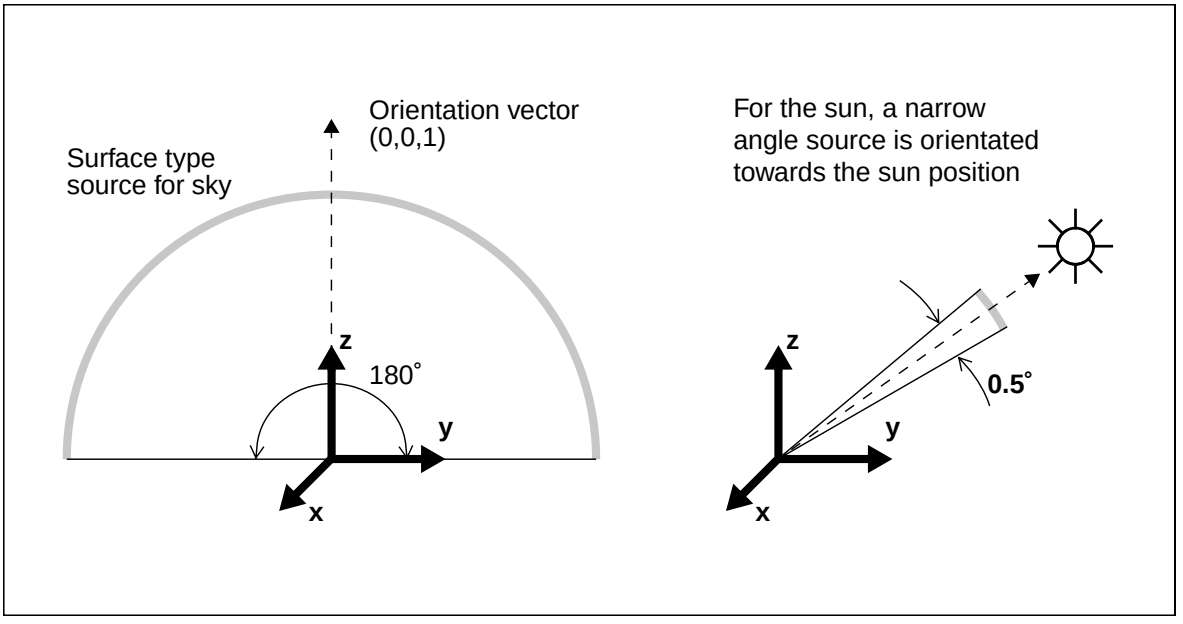

Figure 3-12. Sky and sun source geometry (not to scale)

3.2.3 Modelling the sky and sun

The model sky and sun, when based on measured quantities, can have a representation that is subtly different from what one might expect, given the operational characteristics of the measuring instruments themselves. This difference is demonstrated in the following example in which a model sun is based on a measurement for the direct normal illuminance, E_{dn}. The model description requires a value for the brightness of the solar disc which is:

$$L_{sun} = \frac{E_{dn}}{\pi \sin^2\left(\frac{\Theta_{sun}}{2}\right)} \qquad \textbf{(3-2)}$$

The solar disc angle, Θ_{sun}, is usually taken to be 0.5˚ even though the acceptance angle of the measuring instrument was much larger: 6˚ for the Eppley solar tracker. In *Radiance*, the rationale for this is related in part to the program's hybrid deterministic/stochastic sampling approach. In this,

small concentrated sources of light (i.e. the sun) are sampled deterministically whereas large diffuse sources of light (i.e. the sky) are sampled stochastically [Ward 98]. The material type that is specified for a light source decides the domain in which its contribution to illuminance is calculated: type **`light`** in the deterministic domain and type **`glow`** in the stochastic domain, Figure 3-13. Note from this illustration that, (i) a single ray is used to sample the sun (solar penumbras are therefore not calculated), and (ii) any direct light source that is intercepted by an indirect ray returns zero.

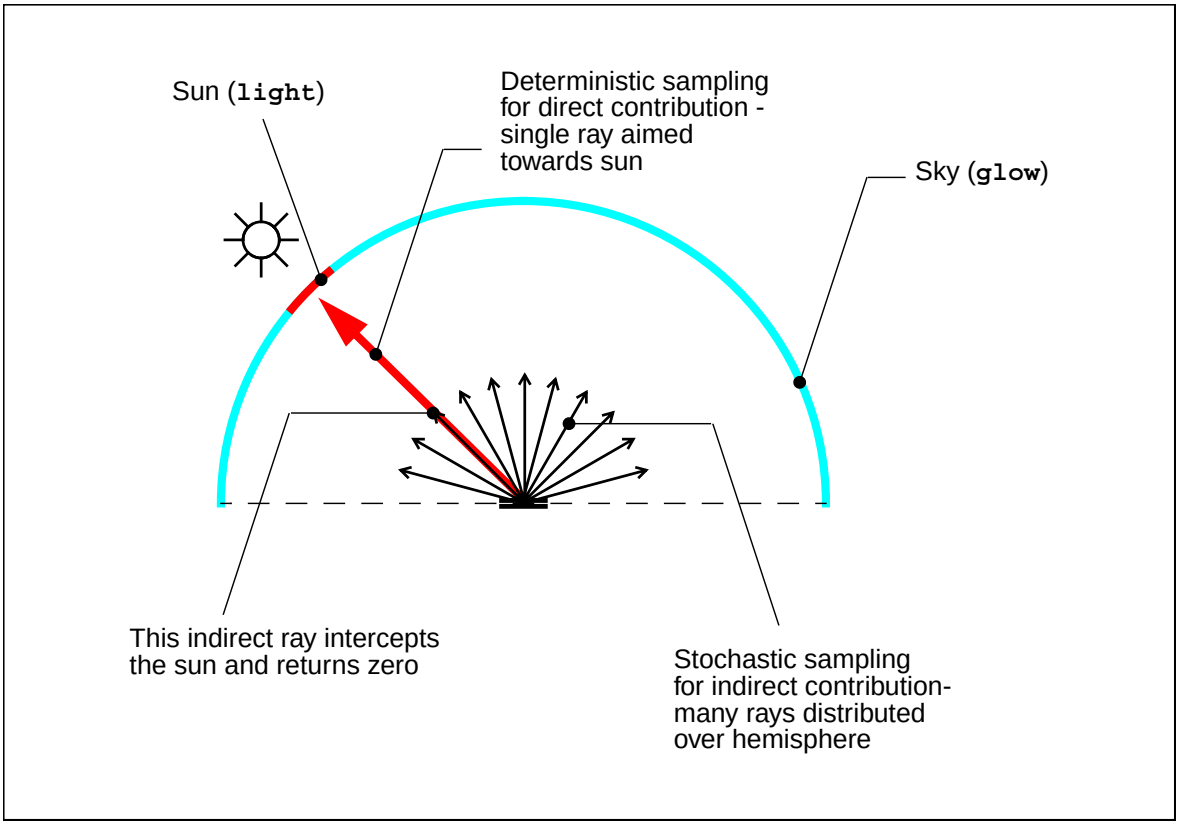

Figure 3-13. Hybrid deterministic/stochastic sampling of the light sources for the sun and the sky

Now, the sun luminance could be defined as a 0.5˚ or a 6˚ solar disc and, for either angle, the resulting direct normal illuminance will be the same. This is because a single ray is aimed towards the source centre regardless of source angle. The prediction of diffuse horizontal illuminance however, will

not be exactly the same for both cases: with a larger (direct) light source, there is an increased probability that indirect sampling rays will intercept it and return zero.[6] If this happens, the indirect illuminance will be underpredicted, albeit by a small amount. Given that the direct calculation is insensitive to the size of the source for the sun, it makes good sense to use a small solar disc. Although the sun source size could be arbitrarily small, convention has it that the actual size used is 0.5˚ - small enough to not interfere significantly with accuracy of the indirect calculation.

3.2.4 The brightdata format

The brightness of the sky source solid angle may, at its simplest be constant, it may take its form from a mathematical function or sky model (see Section 2.3.2), or it may be based on discrete data values - that is, measured sky luminance patterns. To use measured sky luminance data in a *Radiance* simulation, the data values need to be applied as a pattern modifier to a constant (e.g. unit) brightness sky. This can be done using either the colordata or brightdata pattern types.

The definition for the two pattern types is as follows[7]:

> Colordata uses an interpolated data map to modify a material's color. The map is n-dimensional, and is stored in three auxiliary files, one for each color. The coordinates used to look up and interpolate the data are defined in another auxiliary file. The interpolated data values are modified by functions of one or three variables. If the functions are of one variable, then they are passed the corresponding color component (red or green or blue). If the functions are of three variables, then they are passed the original red, green, and blue values as parameters.

```
mod colordata id
7+n+
      rfunc gfunc bfunc rdatafile gdatafile bdatafile
      funcfile x1 x2 .. xn transform
0
m A1 A2 .. Am
```

6. The probability is related to the source angle, for the 6˚ disc this is ~150 x that for the 0.5˚ disc.

7. Taken from the *Radiance* documentation for Version 3.1.

Brightdata is like colordata, except monochromatic.

```
mod brightdata id
3+n+
      func datafile
      funcfile x1 x2 .. xn transform
0
m A1 A2 .. Am
```

So there is just one auxiliary data file for the monochromatic form. The monochromatic brightdata pattern type was used to model the measured skies. The effect of the modifier is illustrated in Figure 3-14. To use

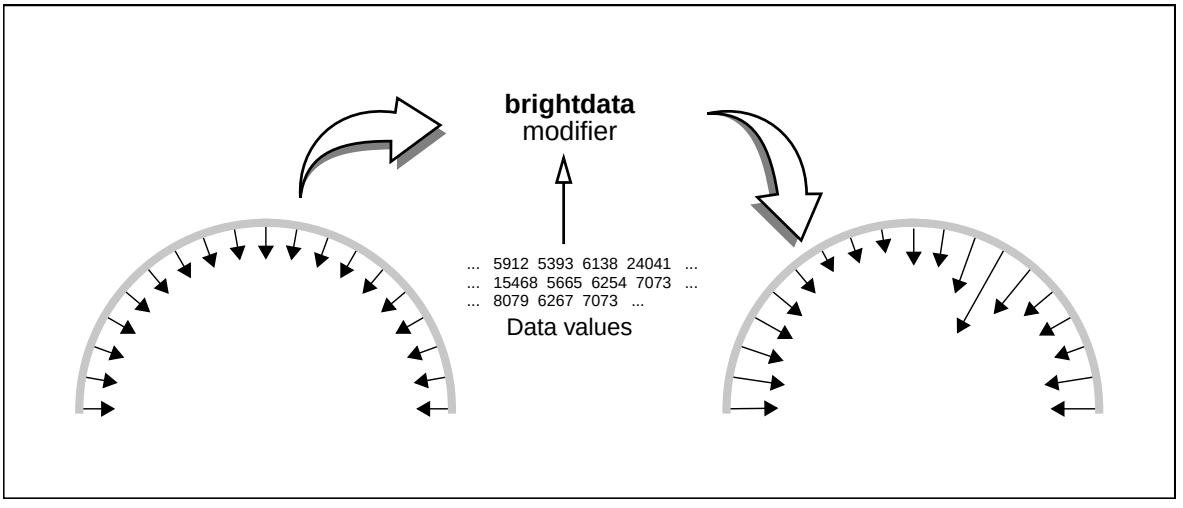

Figure 3-14. Application of brightdata pattern type

measured values with the brightdata modifier, the data must be in a regular grid form - regardless of the projection mapping of the data values, i.e. onto a plane or onto an arbitrary curved surface (e.g. hemisphere). This is so that the bi-linear interpolation scheme in *Radiance* can work effectively. The sky luminance measurements - which were (approximately) evenly distributed across the hemisphere - had therefore, to be mapped to a regular grid. In the regular grid, the spatial increment in either dimension is arbitrary, but it must be constant across the dimension. This means that there must be - for a hemisphere - the same number of azimuth data values at all altitudes. Thus the zenith region will be more 'crowded' with data values than the horizon. In fact, although the zenith is a point, it requires the same number

of azimuth data values as the row of minimum altitude. The minimum azimuth increment for the scanner measurements was 12˚ (for altitudes 6˚ and 18˚), whereas the altitude increment was 12˚ (i.e. constant) across the range. The regular array increments for both dimensions were therefore set to 12˚. This ensured that resolution of the regular grid was commensurate with the resolution offered by the irregular measurement grid. The mapping of the measurement grid to the regular grid is illustrated in Figure 3-15.

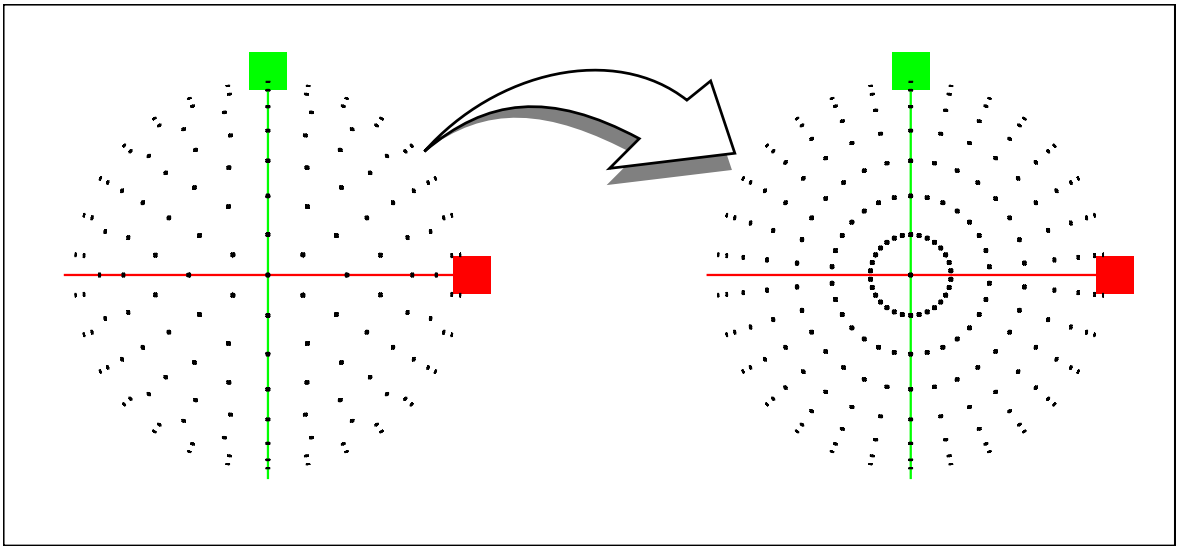

Figure 3-15. Sky luminance data - measured and brightdata-format grids

Example code showing how the regular grid was used to create a *Radiance* model sky that was based on measured values is given below, Table 3-5. To the right of the code is a brief description explaining the function of each block, except for the **`brightdata`** block which functions as follows:

- **`noneg`** - interpolation should not be allowed to give a negative result;
- **`102_92_13h00.dat`** - the file name for the (regular) array of sky brightness values;
- **`.`** - dot character signifies that additional function files are not needed;
- **`Asin(Dz)/DEGREE`** - effect transformation between z-direction vector and altitude (degrees); and,

`# Example measured sky/sun *Radiance* file`	Comment line

`void light solar`
`0`
`0`
`3  3.05528e+06  3.05528e+06  3.05528e+06`
	Declare material (light) for sun and sun R,G,B radiance values - void indicates no previous modifier

`solar source sun`
`0`
`0`
`4  -0.255746  -0.652586  0.713250 0.5`
	Apply modifier for sun material to a surface (source) and define surface orientation (sun position x,y,z vector) and opening angle (0.5˚)

`void brightdata skylumdat`
`5 noneg 102_92_13h00.dat . Asin(Dz)/DEGREE \`
`mod(atan2(Dx,Dy)/DEGREE-201.4,360)`
`0`
`0`
	See main text

`skylumdat glow sky_glow`
`0`
`0`
`4 1 1 1 0`
	Apply sky brightness data modifier to sky material (glow), and set (un-modified) sky radiance R,G,B to 1

`sky_glow source sky`
`0`
`0`
`4 0 0 1 180`
	Apply modifier for sky material to a surface (source) and define surface orientation (upwards) and opening angle (180˚)

Table 3-5. Code example sun/sky input file

- **`mod(atan2(Dx,Dy)/DEGREE-201.4,360)`** - effect transformation between x and y direction vectors and azimuth (degrees) and account for 'offset' angle of the data (each row begins at the sun azimuth, here 201.4˚).

The auxiliary data file for this example (**`102_92_13h00.dat`**) is given in Table 3-6. These data were processed from the original measured data given as an example in an earlier section (Table 3-4 on page 52). Note that there are 31 data values for each row of fixed altitude - the first at 0˚ and the last

2						n-dimensions
6 90 8						Altitude start, end and num. of increments
0 360 31						Azimuth start, end and num. of increments
136.559	150.788	119.575	100.799	86.4134	66.0950	6˚
55.9553	45.1844	38.4749	34.2905	30.9553	29.2793	
28.3631	27.6089	29.3184	29.2570	28.5363	29.8603	
30.1397	30.1285	33.0279	34.9609	40.8547	47.5531	
53.6704	70.1341	79.0726	108.575	130.939	127.933	
136.559						
168.151	206.486	143.128	117.089	99.6871	70.4693	18˚
60.0838	47.3966	39.6536	34.9385	31.3911	29.1955	
28.3855	27.6425	29.4246	29.2793	29.8492	29.9888	
30.4972	31.6480	32.5084	36.8380	42.9553	50.9832	
59.2905	78.8939	91.8045	134.307	171.486	155.268	
168.151						
237.994	236.086	203.288	151.139	106.084	76.4804	30˚
60.6619	48.9486	39.7924	34.0338	30.3520	27.7962	
25.4683	25.8452	27.4294	26.6760	26.7622	28.0178	
28.1867	28.2075	30.8492	34.8234	40.8592	49.4719	
62.5644	82.6648	109.953	145.591	198.390	238.857	
237.994						
282.588	297.349	223.550	146.622	112.400	81.6648	42˚
61.7808	51.2653	42.0207	33.3626	29.5140	27.1163	
26.1069	26.2172	25.6384	24.9665	28.2673	26.9186	
27.5888	31.0334	31.7542	37.9835	42.5023	49.9655	
65.8157	86.2849	110.192	160.183	202.698	235.063	
282.588						
195.356	187.484	165.442	132.729	105.939	86.5196	54˚
68.6423	55.6800	48.1407	38.8632	31.0950	31.9863	
34.5964	33.5086	30.1922	26.9832	25.9508	26.2149	
27.5053	31.2501	35.8603	38.9974	43.9113	52.8805	
65.1501	85.5531	119.297	153.836	178.786	192.466	
195.356						
206.201	174.311	127.900	97.2256	88.8061	86.4749	66˚
77.7416	64.4761	51.5409	41.7282	36.2123	35.7056	
37.4289	38.0035	36.3338	33.4302	30.3943	29.0543	
31.5555	37.3313	43.0112	46.2249	48.8764	53.8903	
63.4782	79.4972	103.670	134.485	170.005	200.370	
206.201						
112.631	111.850	106.802	98.8772	89.3515	79.4358	78˚
70.1500	61.8362	54.6998	48.9937	45.0279	43.0180	
42.5980	43.3012	44.7060	46.4190	48.1200	49.7323	
51.2366	52.6092	53.8212	54.9422	56.4584	58.9768	
63.1449	69.6760	78.9307	89.5837	99.9183	108.205	
112.631						
64.1667	64.1667	64.1667	64.1667	64.1667	64.1667	90˚ (zenith)
64.1667	64.1667	64.1667	64.1667	64.1667	64.1667	
64.1667	64.1667	64.1667	64.1667	64.1667	64.1667	
64.1667	64.1667	64.1667	64.1667	64.1667	64.1667	
64.1667	64.1667	64.1667	64.1667	64.1667	64.1667	
64.1667						

radiance values (31) in evenly spaced azimuth increments (12˚) for each row of constant altitude:

Table 3-6. Data file for sky 102_92_13h00 (radiance values)

at 360° are, of course, identical. This is to ensure continuity across the azimuth range for the *Radiance* bi-linear interpolation scheme. How the mapping from the measured to the regular grid was achieved is described below.

3.2.5 Pre-process of the sky luminance measurements

The PRC Krochmann scanner began each sky scan, and each subsequent row of fixed altitude measurements, at the solar azimuth position. The measurement pattern, though regular, possessed therefore a rotation offset about the zenith axis which was different for each scan. For each row, measurements were taken as the scanner rotated 'anti-clockwise', i.e. N → W → S → E, Figure 3-16(a). For the simulation however, the brightdata pattern type expects the data file to read 'clockwise', i.e. N → E → S → W, Figure 3-16(b). This was another factor that needed to be taken into account in preparing the measurements for use in the simulation.

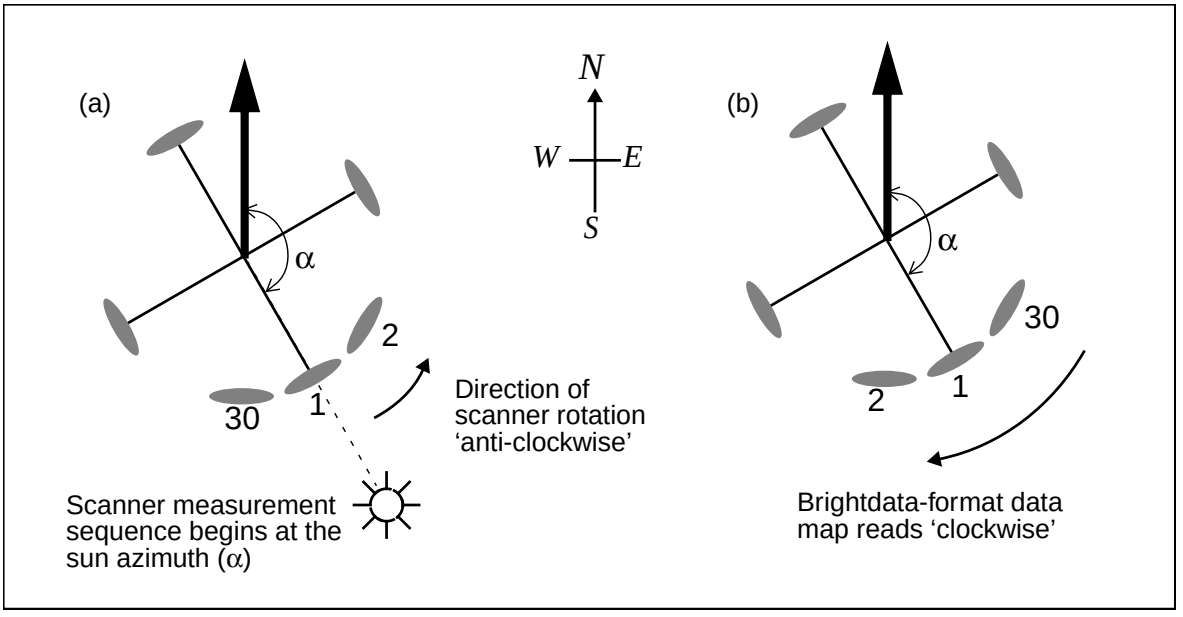

Figure 3-16. Comparison of the measurement pattern (a) with the brightdata format grid (b)

The processing of the scanner measurements for simulation involved the following procedures:

1. The azimuth order of the measured data was reversed.
2. The data were then interpolated to the regular grid pattern and normalised (this stage included the estimation of the out-of-range measurements).
3. The files containing the sky description and the auxiliary data (in *Radiance* format) were written to disk.

For this, a 1-dimensional interpolation scheme was applied across each reversed-order row (i.e. fixed altitude) of in-range scanner measurements.[8] To ensure continuity across the full 360˚ in azimuth, the row vector was concatenated with itself, and mapped to an extended range of azimuth values, Figure 3-17.

Following interpolation, the sky luminance distribution was normalized to the diffuse horizontal illuminance, E_{dh}, which was obtained from measurements of the global horizontal illuminance, E_{gh}, and the direct normal illuminance E_{dn}:

$$E_{dh} = E_{gh} - E_{dn}\sin\gamma_s \qquad \textbf{(3-3)}$$

where γ_s is the sun altitude. This derived value is considered more reliable than using the shadow-band corrected measurement for diffuse horizontal illuminance.[9]

8. A 2-dimensional interpolation in spherical co-ordinates (i.e. a surface fit) is, potentially, a more accurate technique for estimating missing values, because this fit takes into account all neighbouring in-range data. However, the additional complexity was not considered warranted for this application.

9. Private communication - P. Littlefair, BRE.

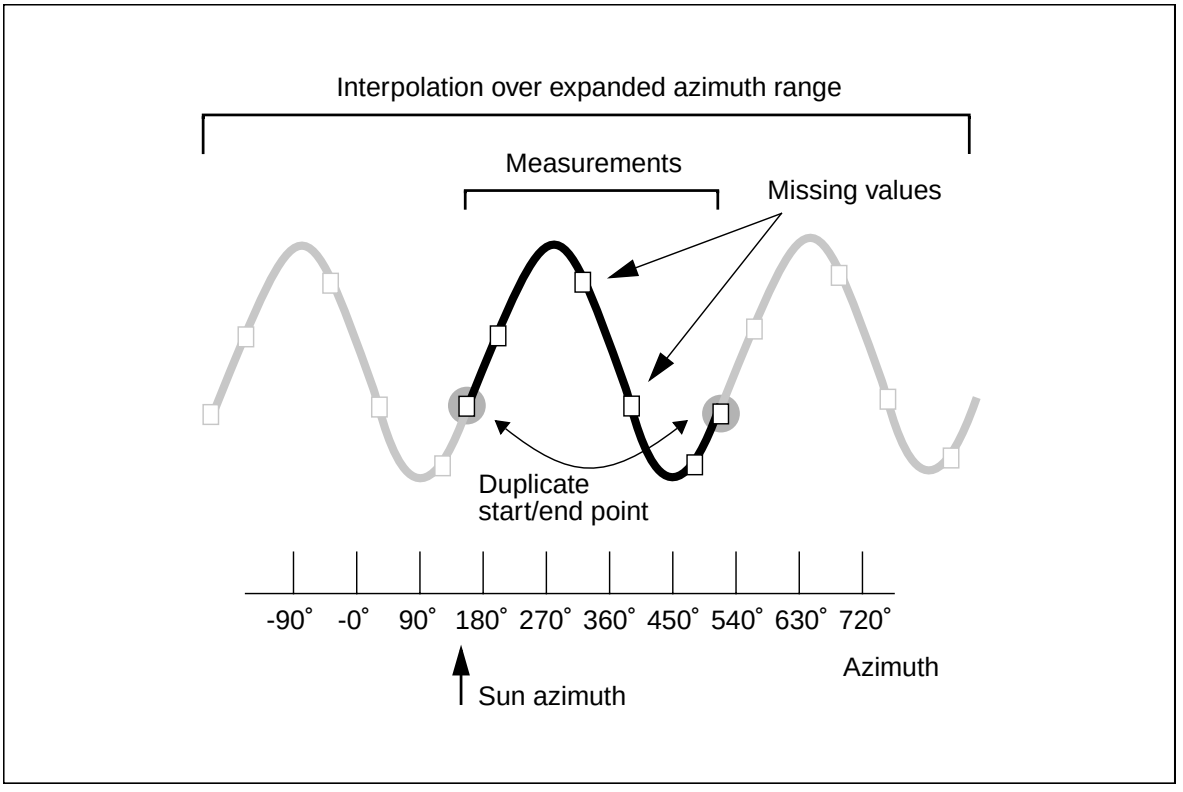

Figure 3-17. Interpolation across concatenated vector

Each (interpolated) sky luminance measurement L_j^u was then normalized to L_j^n using the normalization factor f_{norm}:

$$f_{norm} = \left(\frac{E_{dh}}{\sum_{i=1}^{145} L_i \Omega_i \sin\gamma_i} \right) \qquad \textbf{(3-4)}$$

$$L_j^n = L_j^u f_{norm}$$

Where L_i, Ω_i and γ_i are, respectively, for 'rectangular' patch *i*, the luminance, the solid angle and the altitude of the patch centre, Figure 3-18. The solid angle of the rectangular patch for each row in the measurement pattern is given in Table 3-7.

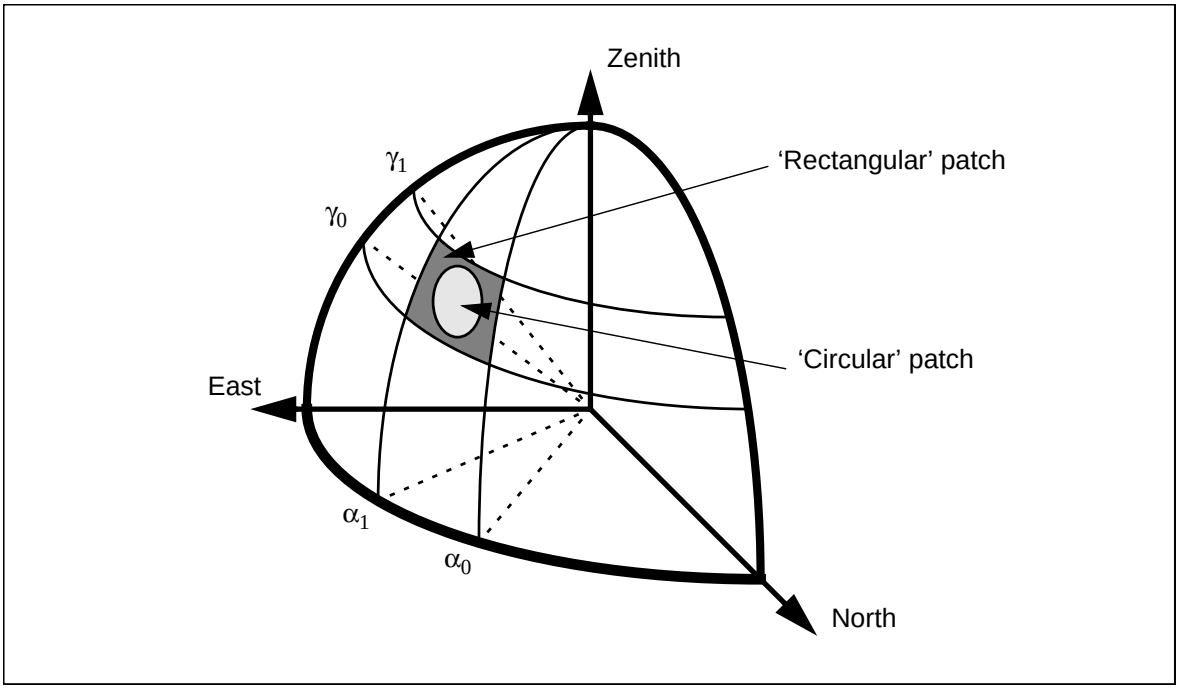

Figure 3-18. Rectangular and circular patch geometry

Row	Altitude	Number of patches per row	Azimuth increment	Solid angle subtended by 'rectangular' patch (sr)
1	6˚	30	12˚	0.0435
2	18˚	30	12˚	0.0416
3	30˚	24	15˚	0.0474
4	42˚	24	15˚	0.0407
5	54˚	18	20˚	0.0429
6	66˚	12	30˚	0.0445
7	78˚	6	60˚	0.0455
8	90˚	1	360˚	0.0344[a]

Table 3-7. Pattern of rectangular patches

a. 'Polar cap', not 'rectangle'.

A sky luminance interpolation/visualisation software tool was created to examine and display the processed luminance distribution for the skies in the validation database. The tool functions either interactively or in batch mode and it displays to either X-windows or PostScript devices. The display for a sky contains the following information:

- Seven plots of the luminance versus azimuth (at fixed altitude) for measured and interpolated sky luminance - both normalized to E_{dh}. The sun azimuth is marked with a dashed vertical line.
- A plot showing the scatter in the six zenith luminance measurements with a horizontal line to indicate the mean.
- A false-colour map of the array of interpolated-normalized sky luminance values. The sun position is at the intersection of the dashed lines.
- A projection of the false-colour map onto a hemisphere. This gives an 'external view' of the sky luminance distribution. The view direction is from the sun position to the hemisphere origin.
- A legend showing the mapping of colour to luminance.

Example output for three skies are shown in Figure 3-19 to Figure 3-21. For sky 093_92_13h15 (Figure 3-19), the estimate for the 'missing' scanner measurement at the sun position is likely to be reliable since this sky exhibits fairly overcast conditions. For clear and intermediate skies (e.g. 125_92_13h15, Figure 3-20), the estimate will be less reliable because it is impossible to accurately reconstruct potentially large luminance gradients when the highest luminance value is missing. The medium-tension cubic-spline algorithm used for the interpolation does allow for estimates greater than the peak measurement in a row (see plot for Scan alt. = 54˚, Figure 3-20). Whilst this may be more realistic than a linear interpolation - which can never exceed the neighbouring values - it cannot be regarded as a truly reliable estimate.

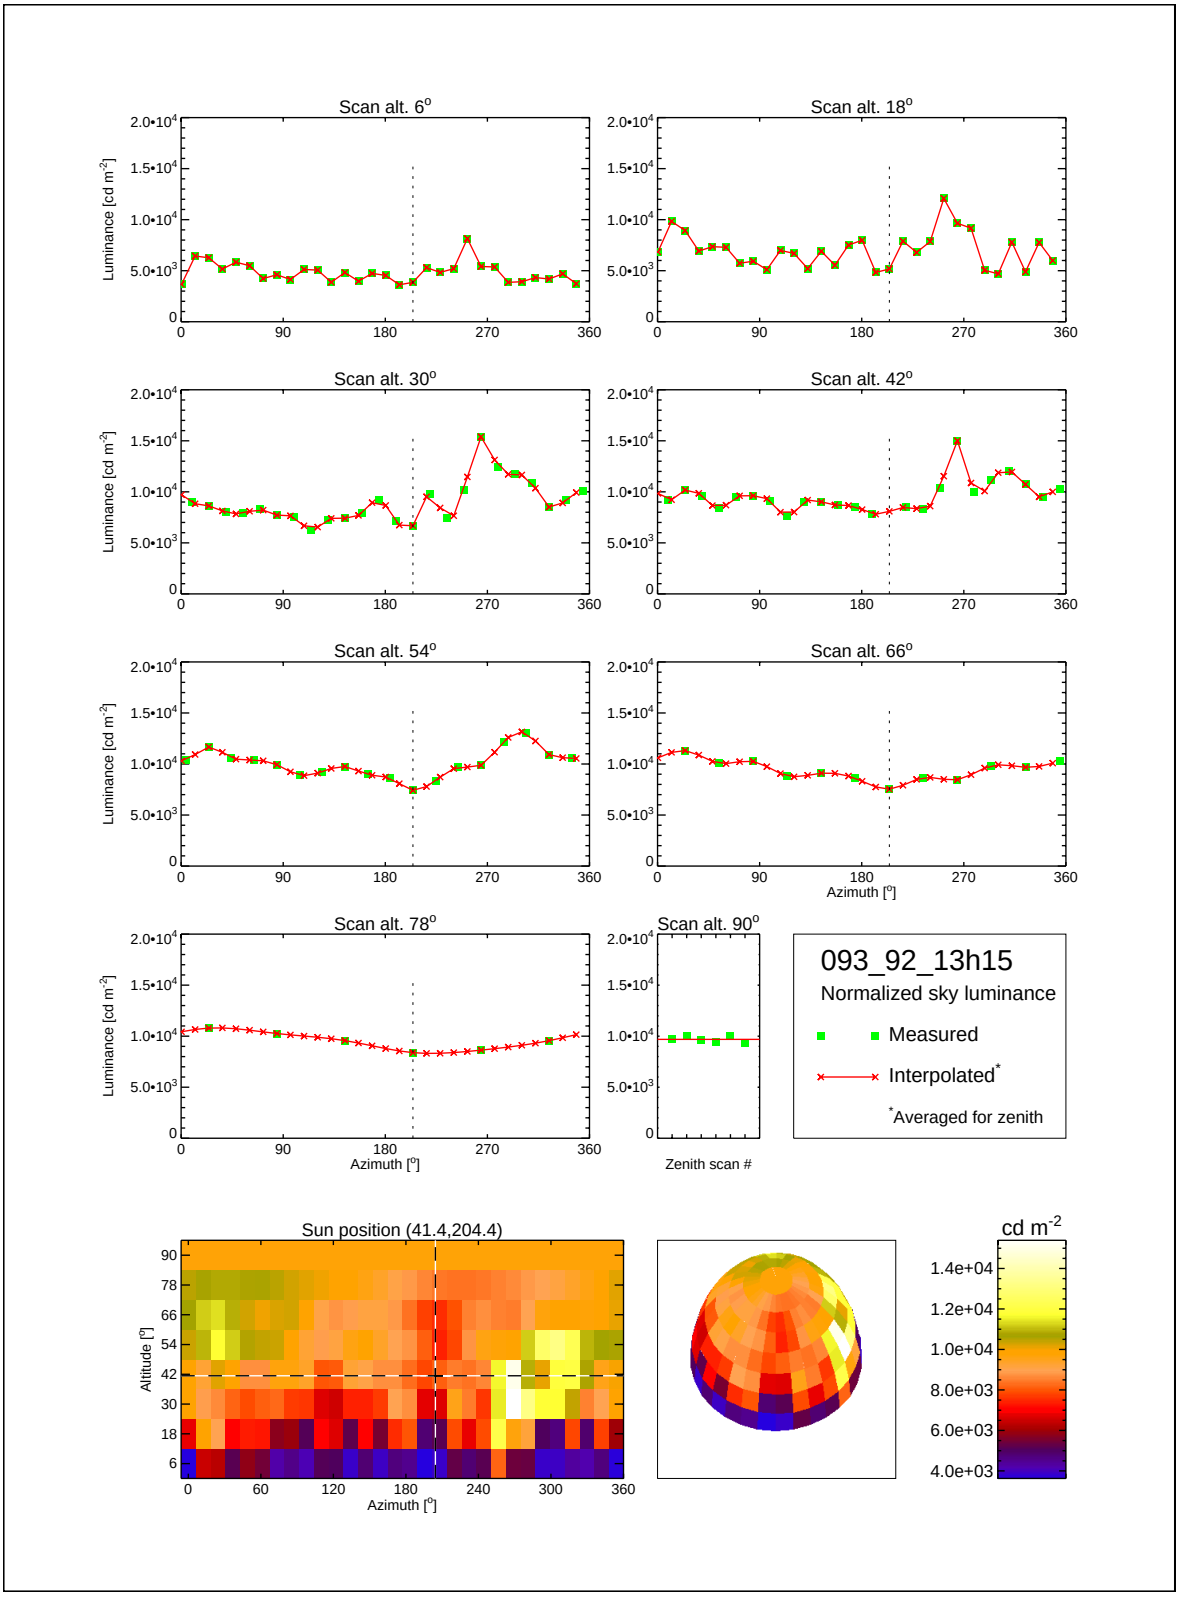

Figure 3-19. Sky 093_92_13h15

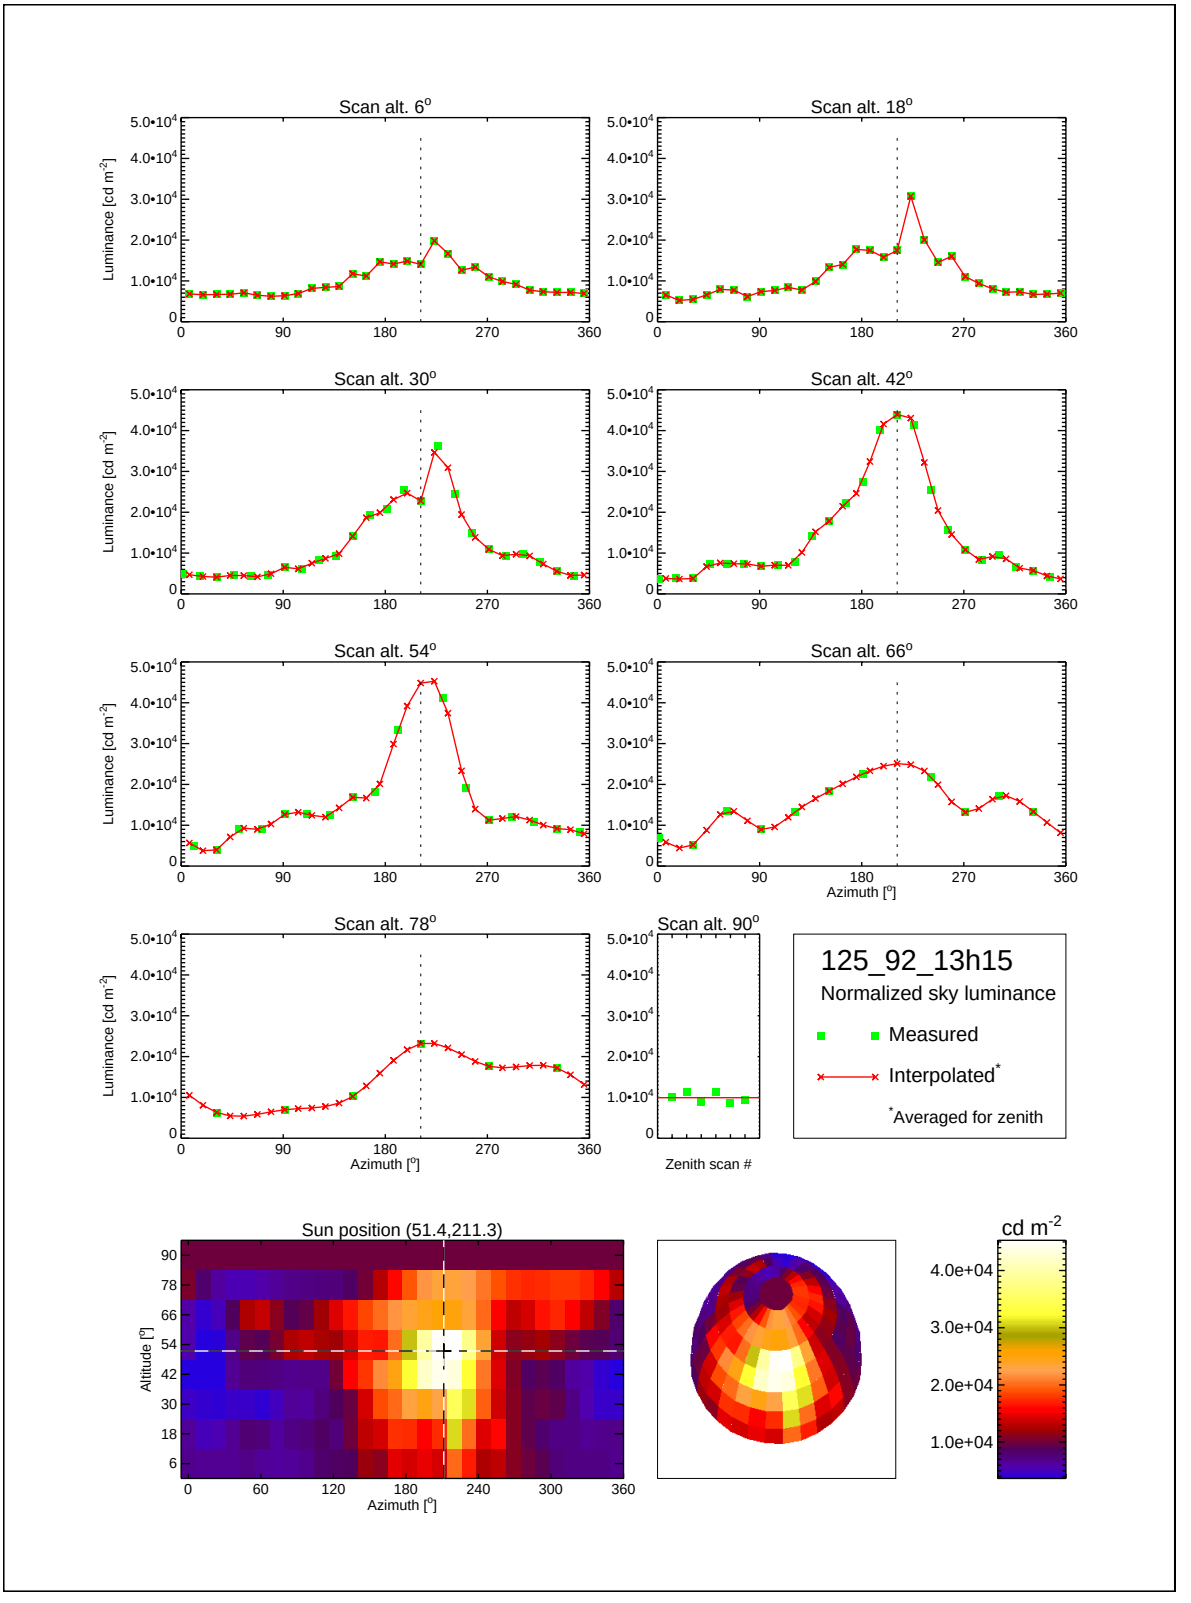

Figure 3-20. Sky 125_92_13h15

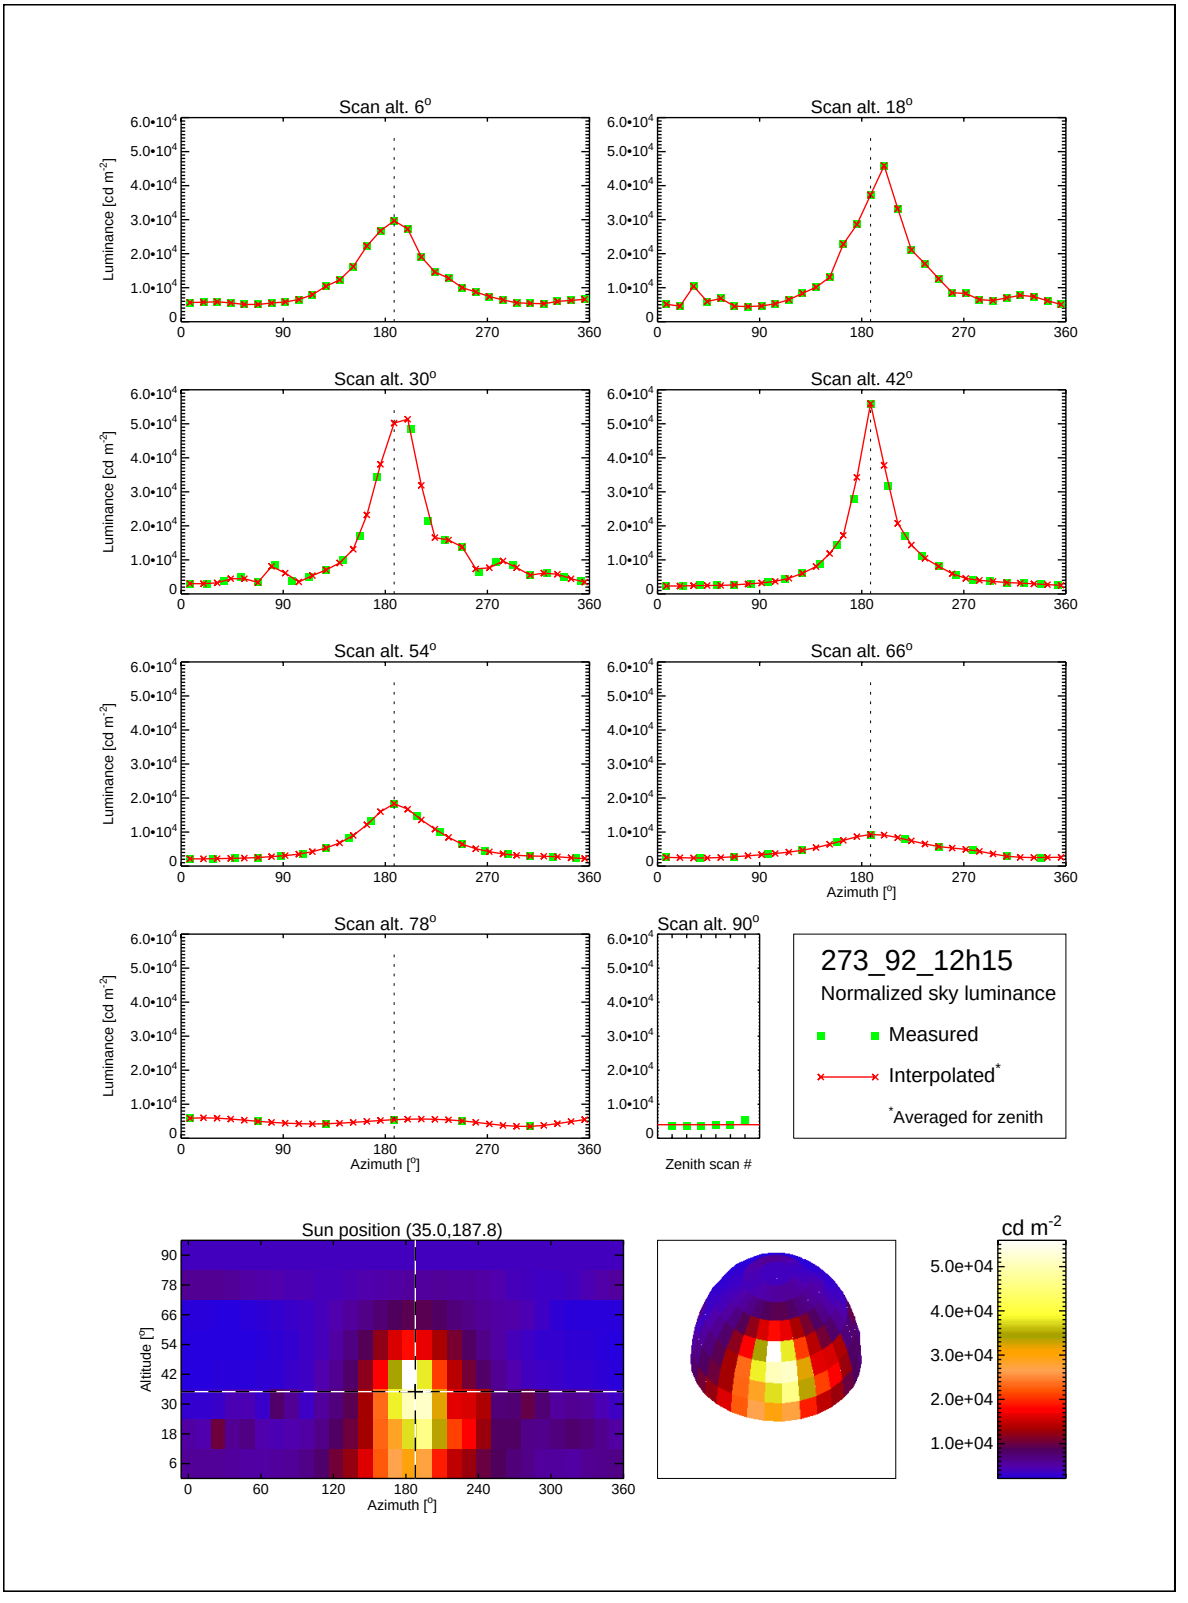

Figure 3-21. Sky 273_92_12h15

It is possible therefore that, for clear sky conditions, the sky luminance at the sun position is under-estimated by the interpolation scheme. This, when it occurs, will affect all the other sky luminance measurements, because the normalization factor (Eq 3-4) is then set to a high value to compensate. Under-estimation of the circumsolar sky luminance may lead to under-prediction of some vertical illuminances - principally the south and west orientations that most often 'saw' the solar disc. Note also that a patch of circumsolar sky, when visible, is likely to contribute proportionately more to the total illuminance on a vertical plane than a horizontal plane, Figure 3-22.

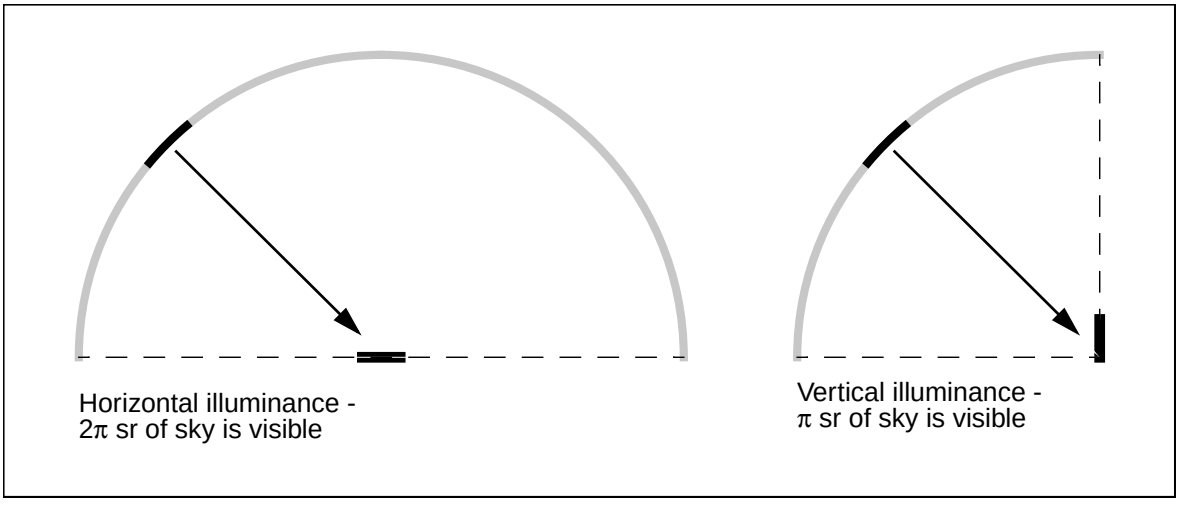

Figure 3-22. Horizontal and vertical illuminances

3.2.6 Deficiencies in the model sky representation

The BRE sky scanner measurements, although as accurate as any comparable dataset, may contain deficiencies that limit the potential accuracy of the illuminance predictions. The principal shortcoming in the measured data was the uncertainty of the sky luminance about the solar position, for both the average across the region and the luminance gradient across it. As discussed in the previous section, these quantities cannot be reliably estimated using interpolation, particularly for clear sky conditions.

The effect that this uncertainty may have on the model representation of the sky is illustrated using a (schematic) plot of sky luminance versus azimuth taken at the solar altitude and centred on the sun position, Figure 3-23. For

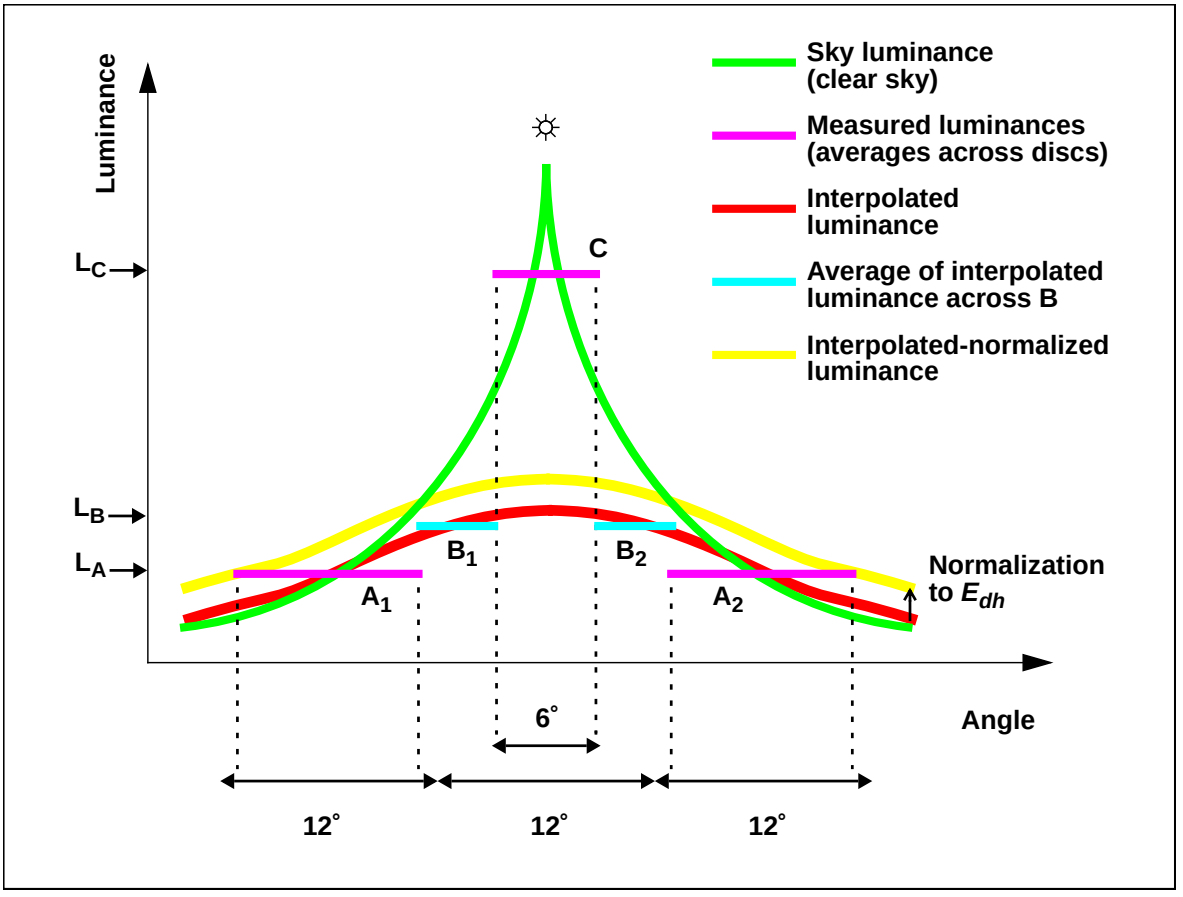

Figure 3-23. Schematic for sky luminance versus angle

this schematic plot, an idealised clear sky luminance is shown. This luminance is symmetric about the solar position and, to simplify the exposition, the sun altitude was taken to be equal to the scanner altitude. The relationship between the circular regions A and C, the annulus region B (in Figure 3-23) and the operational characteristics of the measurement instruments is shown in Figure 3-24. Due to the symmetry, A_1 and A_2 are identical, as are B_1 and B_2.

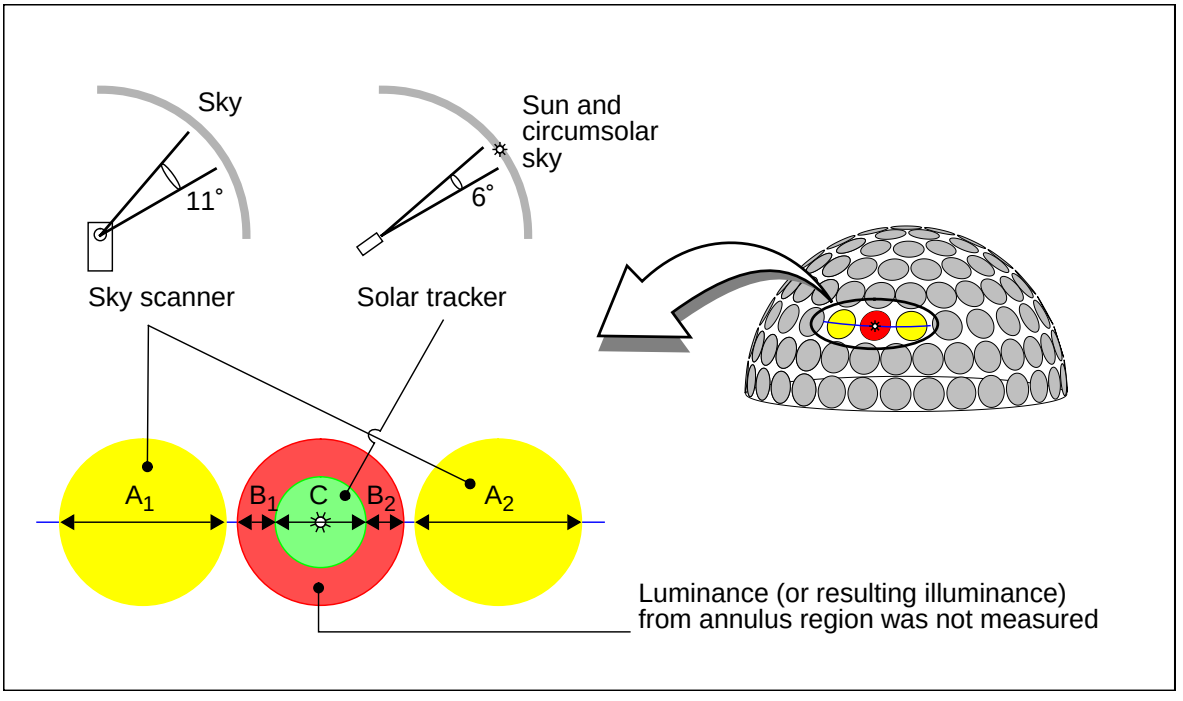

Figure 3-24. Sky scanner (a) and solar tracker (b) acceptance angles (not to scale)

The regions A_1 and A_2 show the sky patches closest to the sun that were measured by the sky scanner. For each (in-range) measurement, the recorded value was the average luminance within the sky scanner's 11˚ acceptance angle. Similarly, for region C, the (derived) measurement of luminance was the average luminance within the solar tracking instrument's 6˚ acceptance angle. The sky luminance across the annulus region (B_1 and B_2) was not measured because the scanner did not record at the sun position. For clear sky conditions, the relation between the (likely) actually occurring sky luminance distribution and the measured-interpolated quantities is summarised in Table 3-8.

Another feature of the interpolation procedure described in Section 3.2.5 is that the sky luminance peak, for clear skies, may not coincide with the solar position. This can be seen in Figure 3-21 where the measured-interpolated peak occurred at (scanner) altitude 42˚ and not at altitude 30˚ which was closest to the sun altitude. This displacement of the interpolated sky

Region	Average luminance	Estimate of luminance gradient across region (clear sky conditions)
A_1 and A_2 (sky patches on either side of sun position)	Measured by sky scanner (11˚ acceptance angle)	Likely to be small or moderate
B_1 and B_2	Not measured - estimated from interpolation of neighbouring values A_1 and A_2	Potentially significant
C	Evaluated from measurement of direct normal illuminance (6˚ acceptance angle)	Likely to be quite large

Table 3-8. Likely luminance gradients across regions

luminance peak from the sun position can be appreciated better from the maps and plots shown in Figure 3-25.

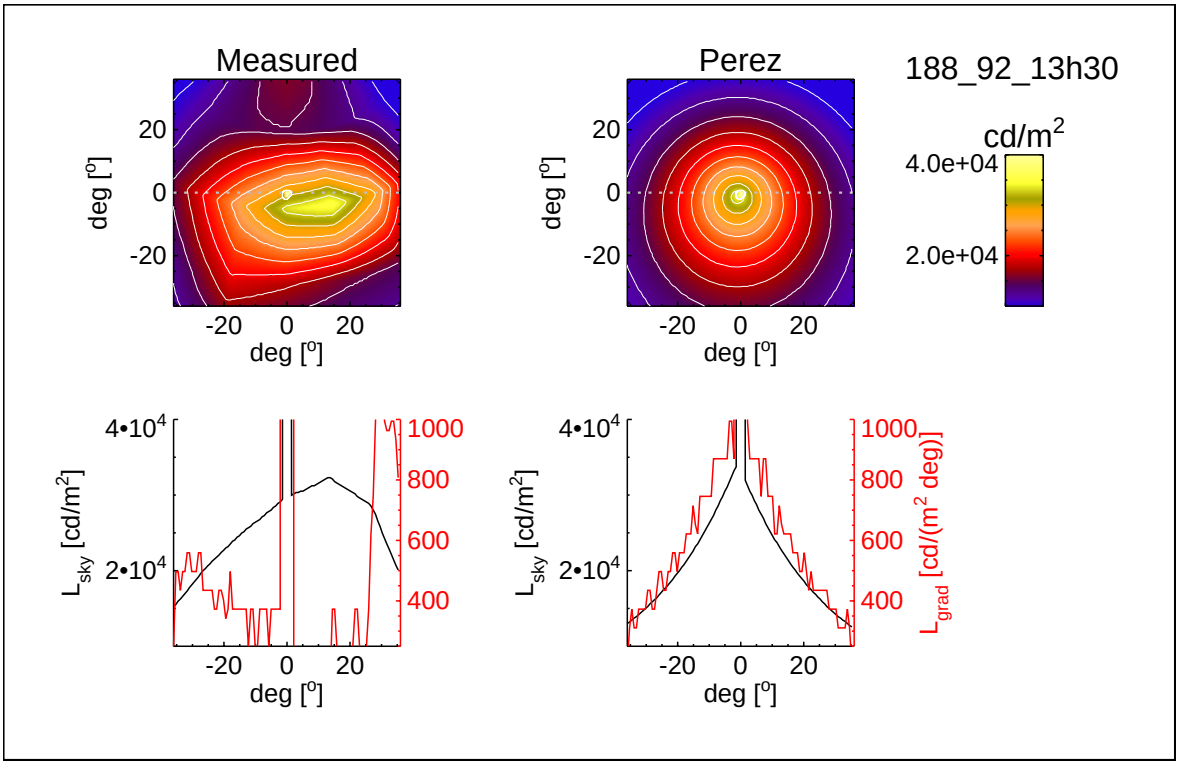

Figure 3-25. False colour maps and profiles of the circumsolar luminance for a measured-interpolated sky and a Perez model sky for case 188_92_13h30

For this illustration, renderings for a 60˚ by 60˚ region centred on the sun position were generated for sky 188_92_13h30 using a luminance distribution based on the scanner data ('Measured') and the Perez-All-Weather model ('Perez'). Each sky was normalized to the same diffuse horizontal illuminance. An angular fish-eye projection was used to generate the renderings from which these maps were derived. In this projection, the distance from the centre of the image is proportional to the angle from the central view direction.

Sky 188_92_13h30 had one of the highest sky clearness indices in the validation sample, and their is no evidence of cloud structure from the measured luminance distribution. It can be reasonably expected therefore that the sky luminance peak should be coincident with the sun position. This was not the case however with the measured-interpolated sky. Here, the interpolation algorithm could not reproduce the luminance peak at the sun position. In contrast, the Perez model predicted an approximately symmetrical luminance distribution centred on the sun position. The luminance gradient in each map can be gauged from the overlaid contours. Below each map is a plot of the luminance profile and luminance gradient across the sun position for the dashed line shown in the maps.[10] Here it can be seen that, across the sun position, there is marked difference between the measured-interpolated sky and the Perez model sky, in both the magnitude and gradient of the luminance profile. It is not intended that any inference regarding the accuracy of the Perez model should be drawn from this illustration.

3.2.7 A hypothesis concerning potentially unreliable photocell-sky combinations

In this section, a class of potential sources of imprecision in the model representation and program operation are identified. These sources of

10. Pixel sampling effects are responsible for the small peaks etc. in the luminance gradient plots.

imprecision are distinct in origin, and yet strongly inter-related in effect. Any one of these could be the cause of occasional gross over or under-prediction in internal illuminance. The class are referred to here as 'source visibility related errors' (SVRE). The reason for this name will become apparent in the discussion that follows. The class comprises four separate types of error: three are related to imprecision in the model representation and one to the operational mode of the lighting simulation program. A description of the four types, their cause, their principal effects and an assessment of the scope for improving or fixing the errors are given in Table 3-9.

Type	Description	Cause	Principal effect	Scope for improvement/fix
A	Imprecision in the geometrical specification of the office model, i.e. inaccurate placement of window bars	Finite resolution of measurement accuracy for linear dimensions - hand measurement by ruler	Photocell actually in shade may be predicted to be in sun, or vice-versa	Repeat measurements using better accuracy techniques[a]
B	Uncertainty in the sky luminance distribution about the solar position	Operational characteristics and finite resolution of the sky scanner and solar tracker	Direct component of illuminance resulting from circumsolar region maybe in error	None[b]
C	Single-ray light source sampling of sun	Default operational mode of *Radiance* sampling	Solar penumbra not computed in simulation	Multiple-ray light source sampling is possible
D	Point source representation of photocell in model	Default calculation mode	Partial shading - and therefore partial response - of the photocell is not modelled	An array of calculation points could instead be used

Table 3-9. Source visibility related errors - type, circumstance and effect

a. In practice this could prove to be a significant task, requiring digital photogrammetry etc.

b. For existing validation set.

Improving on or fixing the type C and D errors would be relatively straightforward, The type A error would be more difficult to improve on; in practice this could be a significant task, requiring digital photogrammetry etc. The type B error however, would still be a major source of inaccuracy,

and so remedial action to correct the type A, C and D errors was not considered justified. Note that the potential for inaccuracy in the illuminance predictions resulting from all four error types - acting independently or in combination - is greatest for sunny conditions when the circumsolar region (and sun) 'come into view' from the photocell position.

It is proposed that:

1. The four error types have the potential to affect only certain photocell-sky combinations.
2. The photocell-sky combinations at-risk are those where the photocell can 'see' all or some of the circumsolar region.
3. Illuminance predictions from the at-risk combinations may contain gross errors which are due to imprecision in the model representation rather than the underlying accuracy of the program.
4. If these at-risk cases are identified and treated separately, then a true assessment of the absolute accuracy of the program can be made.

These propositions form the hypothesis concerning potentially unreliable photocell-sky combinations. In Chapter 4, the error characteristics of the illuminance predictions are analysed, and evidence to support the hypothesis is presented.

3.3 The lighting simulation - preparation

3.3.1 Simulation parameter settings and accuracy

The potential accuracy of the illuminance calculation may not be realised if the simulation parameters are not set correctly. The key simulation parameters for daylight illuminance calculations are those which control the depth (i.e. number of reflections) and resolution of the inter-reflection calculation. In the *Radiance* system these are referred to as the ambient parameters.[11] The inter-reflection calculation, in *Radiance*, progresses recursively. Rays are spawned at the evaluation point(s) to sample the

luminous environment, when a ray intersects with a material surface, additional rays may be spawned and so on. In this scheme, each level of sampling is equivalent to one (diffuse) reflection of light.

The computational cost of an illuminance calculation (and rendering) is very sensitive to the setting of the key ambient parameters. For the work described here and in later sections, it was necessary to carry out many thousands of lighting simulations. For this to be achieved on what is now considered to be a relatively low powered workstation[12] each simulation, ideally, needed to take no longer than a few minutes. A preliminary to the validation simulations was a parameter optimization study where the sensitivity of the accuracy and the simulation time to variation in six ambient parameters was investigated. These parameters were:

- **ad** the number of ambient divisions
- **as** the number of ambient super samples
- **ar** the ambient resolution
- **aa** the ambient accuracy
- **ab** the number of ambient bounces
- **av** the constant ambient approximation

The large number of parameters requiring investigation presented a problem: even if the range and the number of values for each parameter was restricted to say five, the total number of possible combinations would be large (i.e. $5^6 \cong 15.6 \times 10^3$). An additional complication is that, at low resolution, Monte-carlo calculations can give seemingly accurate predictions through 'chance hits' rather than from reliable convergence. 'Chance hits', when the occur, are by their nature unreliable - a small change in parameter value in either direction can give very different results.

11. In computer graphics, light not received directly from a source of illumination is usually referred to as the 'ambient component'.

12. Sun SPARC station 2.

And, a chance accurate prediction for one luminous environment may not necessarily be repeated for another. This is illustrated in the following example where the direct sky component (expressed as daylight factor) at a point in the room was predicted using a wide range of ambient divisions (i.e. initial sampling rays). Genuine convergence in the predicted value is apparent for ad > 128. Note however that for ad = 2 and 4, these samples produce 'chance hits' which result in predictions that are close to the converged value, Figure 3-26.

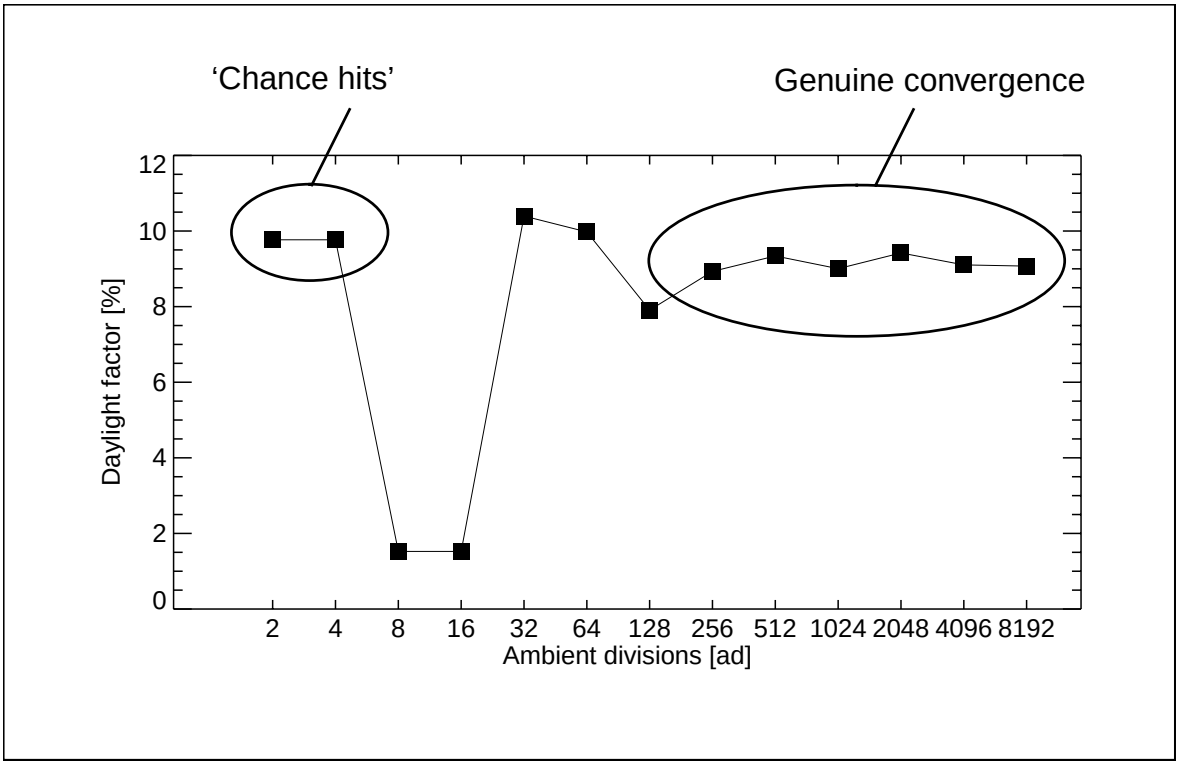

Figure 3-26. Predicted sky component (daylight factor) versus number of ambient divisions

As a result, it was not sufficient to select one parameter combination, which happened to give an accurate result for one sky, and hope that the success would be repeated for the entire validation sample. Instead, what was needed was a robust parameter selection method which ensured that an accurate result, when achieved, was relatively insensitive to moderate

changes in any of the parameter settings. And also that the prediction was not highly sensitive to any particular sky and sun conditions. This goal was achieved by examining the trend in the accuracy of the predictions as a single parameter was varied, with the other parameters held constant. This was done for each parameter in turn.

A positive ambient value can be used to approximate the contribution of higher order reflections in a rendering or illuminance calculation. If ab is set to zero, the ambient value is used directly to approximate the (essentially) infinite number of light reflections that can occur. For ab > 0, the ambient value is the remainder contribution at the final (calculated) reflection. It is clear however that for all normal spaces under varying natural illumination conditions, an appropriate ambient value is both time (that is, illumination) and position dependant: it will be greatest near to the windows and least at the back of the room. A constant ambient approximation can, at best, be appropriate for only a limited range of sky conditions and for only a small fraction of the workplane surface in a typically glazed space - such as that used in this validation study, Figure 3-27. It is not reliable therefore to use

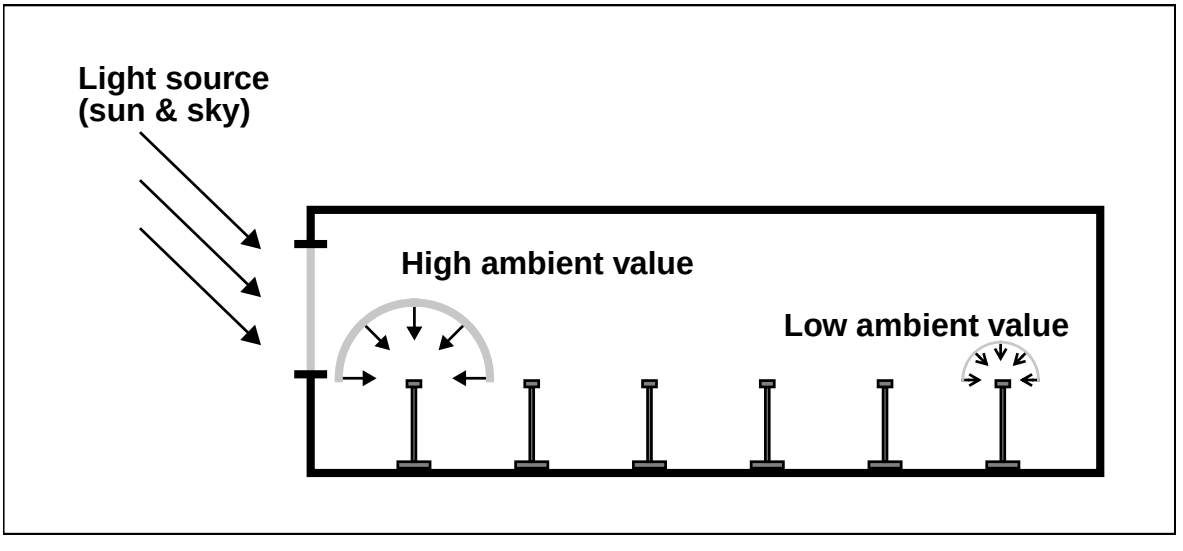

Figure 3-27. Constant ambient value approximation

the (constant) ambient approximation when high accuracy predictions are needed.[13] For the parameter optimization exercise that is described below, the ambient value was set to zero and the total light contribution due to inter-reflection was calculated explicitly. This parameter was therefore eliminated from the optimization exercise. As a consequence, it was necessary to carefully examine the convergence characteristics when increasing the ab parameter. If this parameter is not set sufficiently high, the calculation is likely to consistently under-predict illuminance regardless of the resolution of the other ambient parameters.

3.3.2 Optimization methodology

The methodology for the optimization was as follows. One clear sky case and one overcast sky case were selected at random from the validation data. For the clear sky case the internal illuminance was predicted using a balanced set of high resolution parameters (with av=0). If the simulation did not complete within 1hr. of CPU time, it was terminated and the simulation was restarted with one or more of the parameters relaxed.[14] Once an accurate prediction was achieved using 1 hr. or less CPU time, the simulation was repeated for the overcast sky to ensure that equivalently accurate predictions were obtained. There was an element of luck here, the few cases that were chosen at random for the initial tests all yielded accurate illuminance predictions at each of the photocell locations.[15] The results presented in Section 4.5 however show that this would not have been the case for all of the skies in the validation data.

The high resolution ambient parameter settings which gave an accurate result were called the ‘slow-basecase’ combination. The next stage was to examine the trend in results as, one at a time, a parameter was varied from

13. This is not usually the case for renderings. See “Setting -av and -aw” on page 38.

14. Initial tests using ‘guesstimate’ parameter combinations showed that accurate results could be achieved using 1hr or less CPU time.

15. Fortunately, the skies randomly selected for this preliminary exercise did not contain occurrences of photocell - sun position combinations that proved to be unreliable (see discussion on SVRE Section 3.2.7).

a very low resolution value to its (high resolution) 'slow basecase' setting. Where possible, the increments were chosen to cause an approximate doubling in the complexity (and therefore computational cost) of the calculation. For example, the number of ambient divisions was increased from 2, 4, 8, 16, 32, and so on. The trend in CPU usage was compared alongside the trend in the accuracy of the predictions. The hypothesis governing this approach was based on the assumption that, for each parameter, a value could be found which gave accurate results quickly, and which were stable to moderate changes in parameter value. Albeit, in combination with, in each case, the 'slow basecase' settings. The individual parameter values determined in this way were collected together to form a new combination called the 'fast basecase'. The illuminance predictions were repeated using this fast combination of parameters. The final stage in the optimization was to 'fine-tune' the 'fast basecase' parameters by incrementing them - one at a time - to higher resolution values, trading off gains in accuracy against increases in CPU time. This resulted in the final 'basecase' set of parameters that was used for most of the lighting simulation work described in this and later chapters.

Flexible optimization criteria were employed at various stages, and the process was steered to some degree by the intuitive sense for predicting outcome that a user often develops from working with a complex simulation model. An example set of plots from the optimization exercise are shown in Figure 3-28. In this test, the number of ambient divisions (ad) was the parameter that was varied (from 16, 32, 64, etc. to 4096), the others were held constant. For each value of ad tested, the illuminance (measured and predicted) is shown versus distance from the window. Below each illuminance plot there is a bar graph showing the relative error in the illuminance prediction at each photocell location. The graph titled 'Convergence', plots the average of the relative errors (absolute values) for the six illuminance predictions versus the processor (CPU) time used for each of the ad values tested. Here, the fastest simulation (ad=16) took only

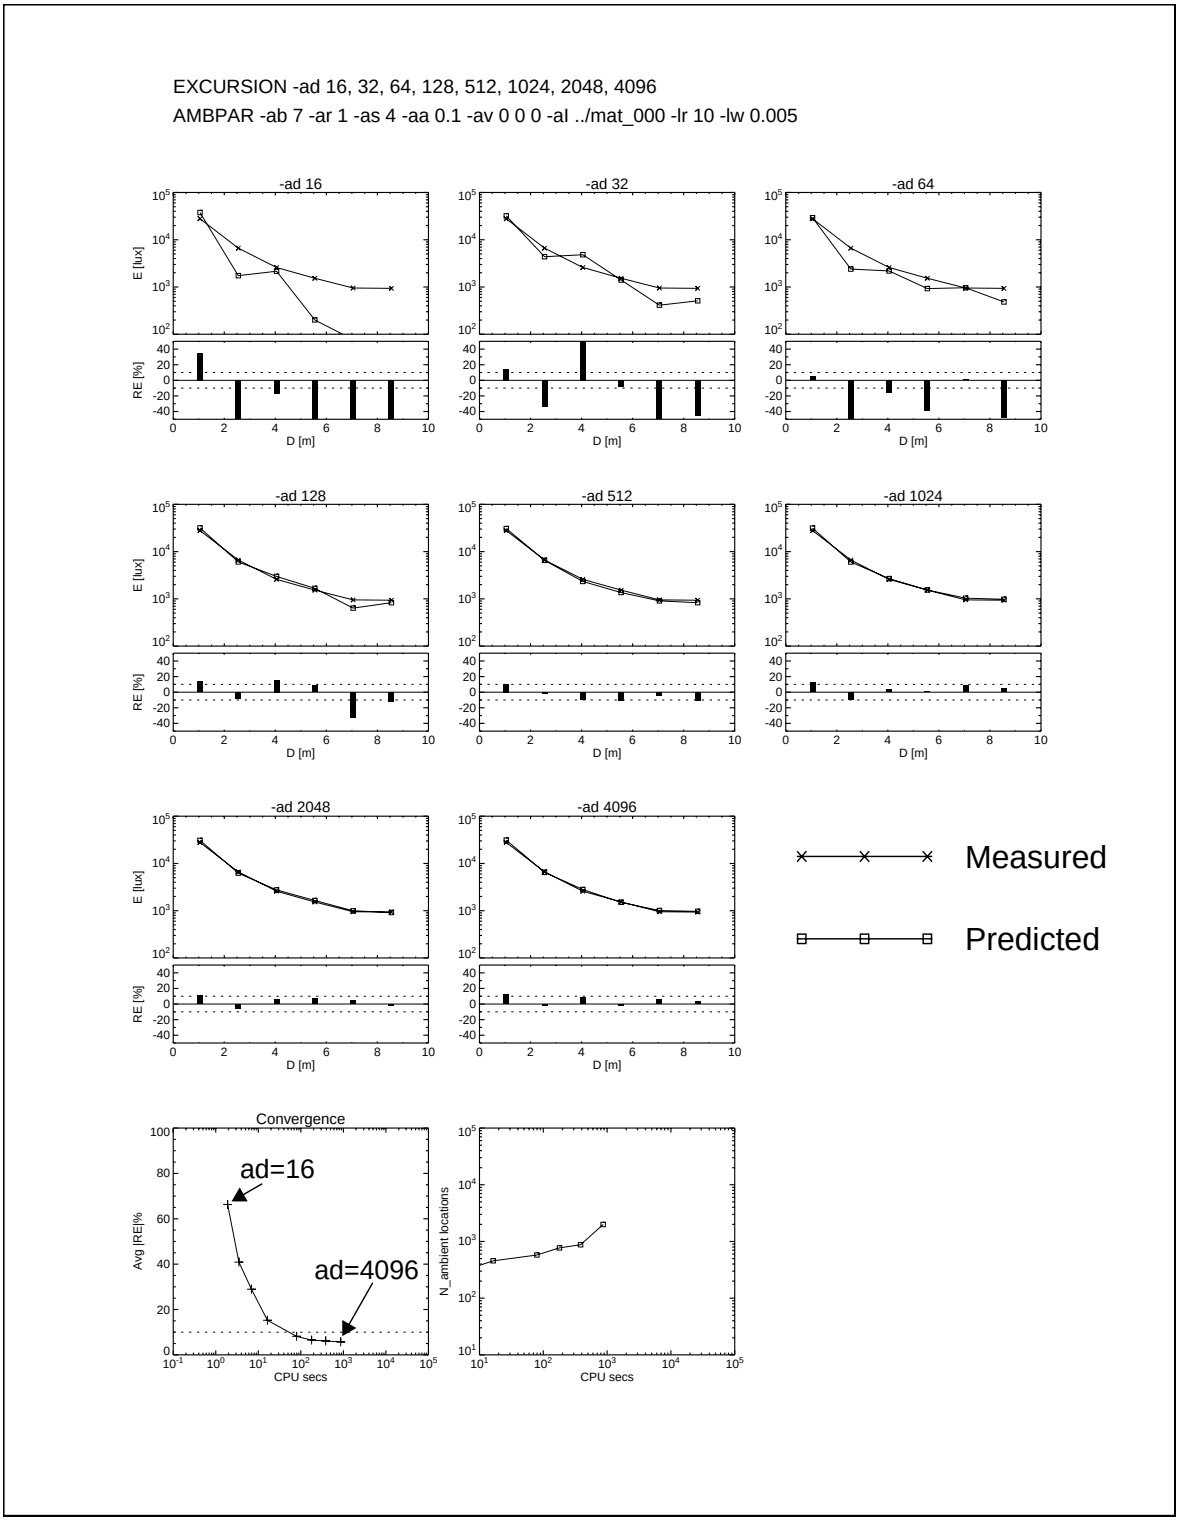

Figure 3-28. Results for ambient divisions excursion

a few seconds but produced large errors, whereas the slowest (ad=4096) took ~1000 secs and gave very accurate predictions. The last plot shows how the number of ambient locations (that is, points were an indirect irradiance gradient was calculated) was related to the CPU usage. Each series of simulations for an excursion was initiated from custom C-shell scripts, which in turn were initiated from an 'executive' script that controlled the entire simulation sequence. The optimization exercise was therefore highly automated, and much of the available processor time was used to thoroughly investigate the convergence characteristics of the illuminance calculation.

For reasons of brevity, the majority of the simulation data resulting from the optimization study (dozens of sets of plots) have not been included in this thesis. For a practical guide to how to set the ambient parameter values, see the author's chapter in Rendering with *Radiance* (Chapter 2 in this thesis).

The 'basecase' parameters that were determined using the optimization methodology described above are listed in Table 3-10. With this parameter combination, each simulation used approximately 5 minutes CPU time. The total CPU time for one pass of the validation data was therefore approximately 2.6 days. The 'slow basecase' parameter combination in comparison, would have required about 1 month CPU time.

Parameter	Value
ad	2048
ab	7
ar	2
as	32
aa	0.10

Table 3-10. 'Basecase' parameter values (av=0)

3.3.3 Ambient calculation - progression and convergence characteristics

The progression of the ambient calculation can be appreciated from the renderings shown in Figure 3-29. For these images, a red marker was added

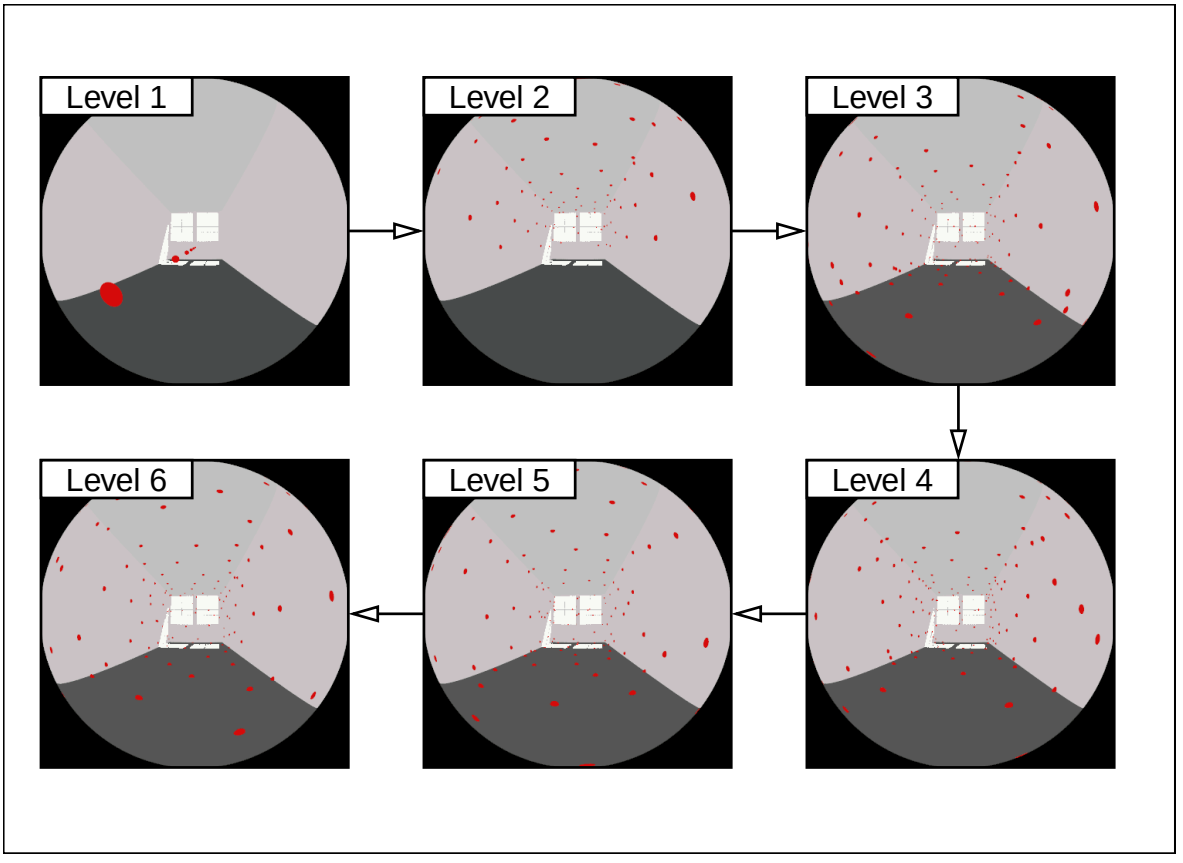

Figure 3-29. Recursive progression of ambient calculation; levels 1 to 6

to the model at those places in the scene where an indirect irradiance value was calculated. These locations were extracted from the ambient file which resulted from a seven bounce simulation for case 102_92_13h00. For ab = 1 (level 1), there were six points from which indirect irradiance sampling was initiated. These were the six photocell locations.[16] At these points, hemispherical sampling rays were spawned. Some of these rays will sample

16. It is not necessarily the case that the number of initial sampling points will always equal the number of calculation points - See “Case Study III: Introducing Complexity” on page 24.

the sky luminance through the glazing. Most however will intersect with opaque surfaces, and from some of these points the next level of sampling was initiated. These are the red markers for 'Level 2', which of course, are all above the (horizontal) plane of the photocells. For this illustration, the number of initial sampling rays was 2048. But the number of sampling points at higher levels is much lower than this because most of the spawned rays use nearby cached values, that is, already determined indirect irradiance values. The number of sampling points at each level is given in Table 3-11. This caching and reusing of indirect irradiance values is one of the keystone features of the *Radiance* program. Without this and other optimizations, the total number of rays spawned would grow geometrically and soon become unmanageable.

Level	Number of sampling points
1	6
2	87
3	199
4	202
5	188
6	151

Table 3-11. Number of points at each level where hemispherical sampling was initiated

The convergence characteristics of the illuminance calculation for one case (121_92_14h15) are shown in Figure 3-30. Here, the RER in the illuminance prediction at each photocell is shown for ambient bounces equal to 1, 2, 3, 4, 5, 6 and 7. It can be seen here that inter-reflection is generally more important at the back of the room where ab > 5 is required to achieve a |RER| < 10%.

3.3.4 Automation of the simulations

A scheme for the management and automation of a large number of simulations needs to be both efficient and extensible. Efficient, because the sequence of simulations should ideally be executed with minimal user

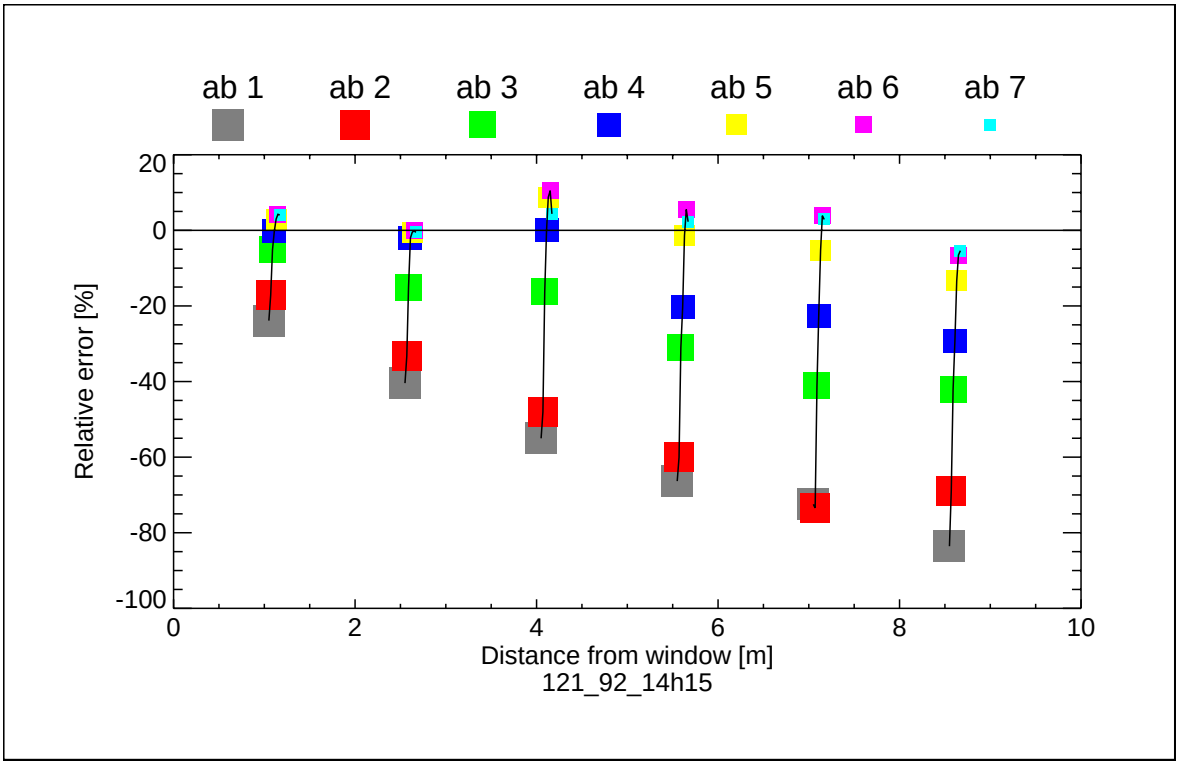

Figure 3-30. Convergence characteristics of the illuminance calculation

intervention, and extensible so that no practical limit is placed on the range or scope of the investigation(s). Moreover, for the purposes of validation, it is advantageous to maintain the measurements and the predictions in a common format.

As demonstrated in Section 3.2.5, the sky luminance measurements needed to be reformed to be compatible with the *Radiance* brightdata format. This was achieved using a set of procedures and functions written in IDL. For the initial tests and parameter optimization (Section 3.3.2) just a couple of skies were prepared using the IDL programs, and all the simulations were initiated from (UNIX) shell scripts. Having settled on a basecase set of ambient parameter values, a scheme was conceived to manage both the execution of the simulations and the updating of the results dataset. Furthermore, the same program environment would be

used for the analysis and the plotting of all the results. The IDL environment was selected for this task.[17]

Prior to the simulations, it was necessary to prepare the validation data and convert it to IDL variables. The first stage was to create a single 2D floating-point array which contained all of the relevant matched entries in the validation data files.[18] In total, the BRE supplied 81 ascii data files (27 days, and 3 types). The sky luminance distribution was recorded every 15 minutes, but the other measurements were given as 5 minute averages of 1 minute data. It was necessary therefore to 'time-align' the measurements: readings taken at the same instant were identified and formed into a row vector for insertion into the array. Each row vector of the array therefore contained all the (unique) entries in the measured data that were taken at the same instant. Some of the measured quantities were of type integer, these were converted to floating-point. The 5 minute data was maintained in a separate array structure.

The array of measured quantities was of size 754 x 178, that is, 178 measured quantities (and identifiers) taken at 754 instants. The contents of a row vector are given in Table 3-12. The simulation results for each

Index	0-1	2-3	4-5	6-12	13-15	16-165	166-171	172-177
Quantity	day year	solar azimuth altitude	hour minutes	glb.horiz. glb. diff. vertical N,E,S&W dir.norm.	zen.lum. temp. humid.	150 sky luminance meas.	innov. office illuminance	single glazed office illuminance

Table 3-12. Measured quantities by vector index

17. IDL is a high-level, interpretive programming language with powerful data analysis and visual display features. IDL variables, procedures, operators and functions operate on scalar, vector and array data with no change in notation or meaning. Additionally, IDL can communicate with the UNIX operating system. It is relatively straightforward therefore to execute shell scripts etc. from within a IDL program.

18. Irradiance quantities in the ALR files were excluded.

individual sky were concatenated to the validation array thereby increasing the number of columns in the array. For example, the first quantities to be predicted for all 754 skies were the global horizontal illuminance and the four vertical illuminances. These five predicted quantities were - for each sky - concatenated to the row vector for that sky. The array size was then increased to 754 x 183. The index numbers for these predicted quantities were, Table 3-13.

Index	178	179	180	181	182
Quantity	glb.horiz.	vertical N	vertical E	vertical S	vertical W

Table 3-13. Vector index for first predicted quantities

At the time of completion of this thesis, the validation array had grown to size 754 x 405: that is, 227 lighting and lighting-related quantities - for each sky - were predicted using *Radiance*. The quantities added to the validation array at various stages included the following:

- Internal and external illuminances using measured sky luminance distributions (Chapter 4).
- Visibility tests and components of internal illuminance (Chapter 4).
- Internal illuminance predictions using sky model generated luminance distributions (Chapter 5).
- Internal illuminance predictions derived using daylight coefficients (Chapter 6).

A full description of the contents of the validation array is given in Table A-2.

In Chapter 4, the sensitivity of the relative error in the internal illuminance predictions to several measured and predicted quantities is analysed. This process was greatly facilitated by maintaining all the measurements and predictions in a simple common format. As the range and scope of the validation grew, so did the size of the validation array. Because the

validation array grew by concatenation of the row vectors, backwards compatibility with analysis programs was preserved: already existing programs could use the newly updated validation array without modification.

Each sequence of 754 (or more) *Radiance* simulations was initiated from an 'executive' IDL program, specially written for the task. Although each executive program was different they all shared a basic program structure, Figure 3-31.

3.4 Conclusion

The preparatory work for the validation of the *Radiance* lighting simulation program has been described. Each stage in the processing of measured sky luminance distributions has been presented and example file formats etc. given. A hypothesis concerning potentially unreliable photocell-sky combinations was formulated in Section 3.2.7. This hypothesis is tested in Chapter 4.

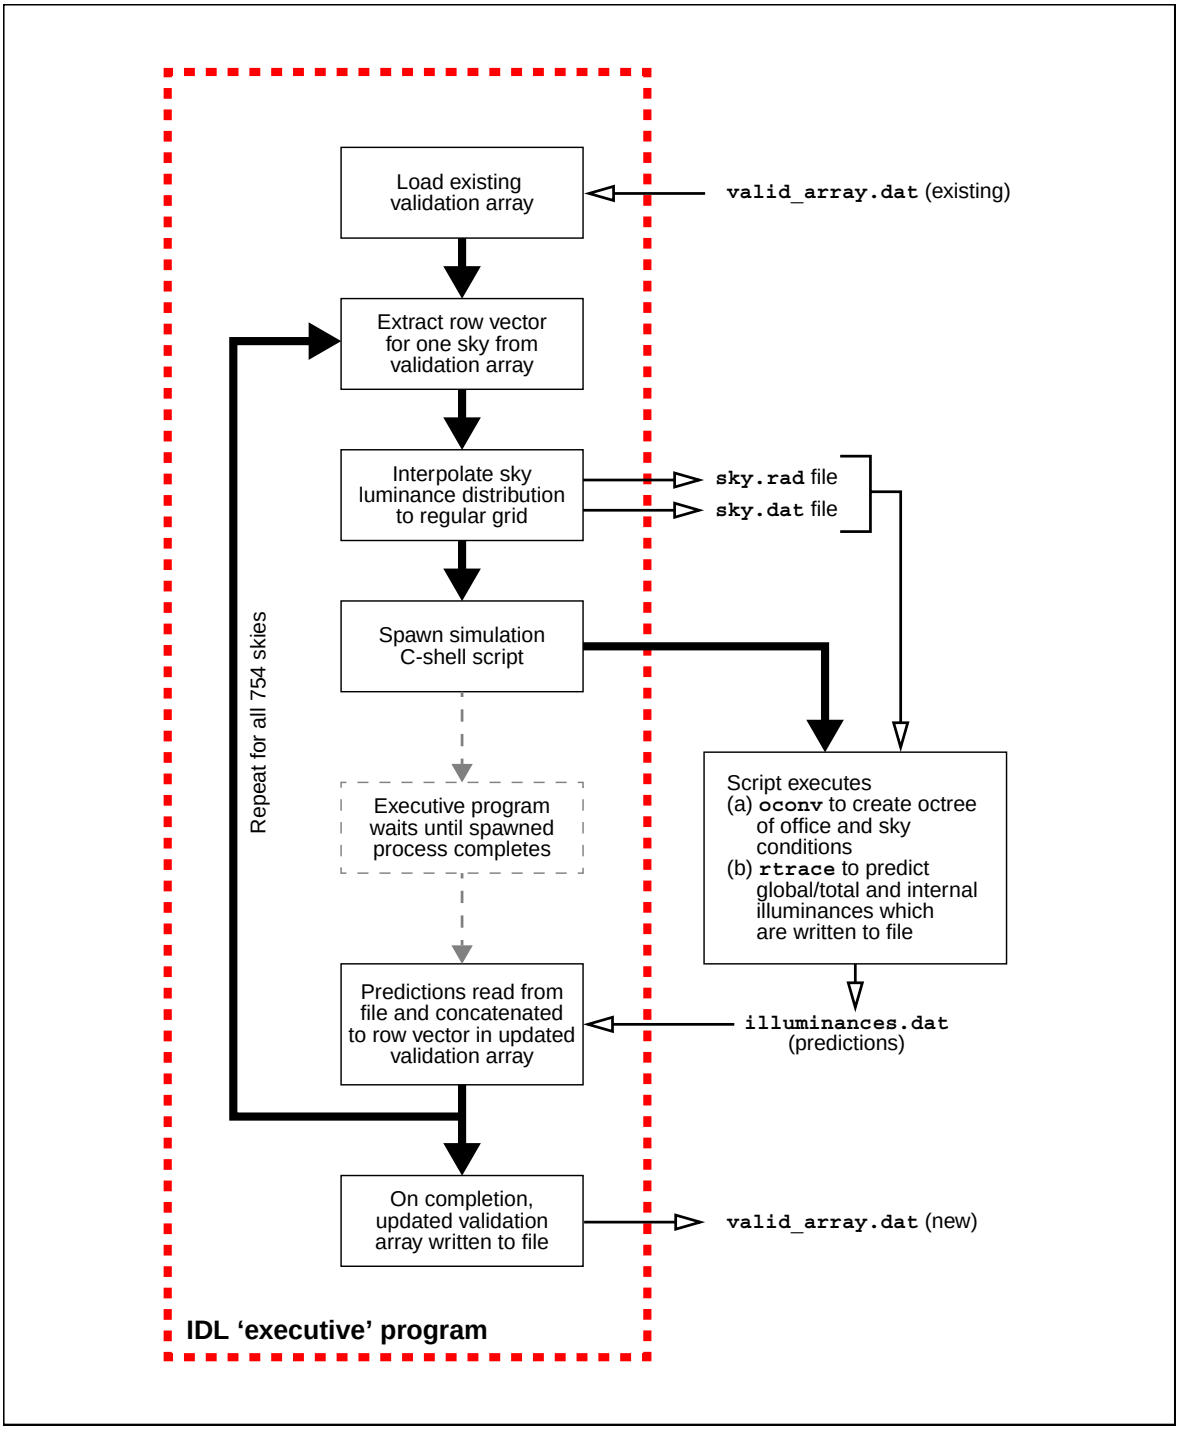

Figure 3-31. Structure of the 'executive' program