Three-dimensional Illuminance: a tool for lighting design?

James Casey
M.Sc. Light & Lighting
Bartlett School of Architecture
Final Report
15 September 1998

Contents

Introduction

Currently when designing a lighting scheme, the CIBSE *Code for Interior Lighting* specifies criteria for the suitability of a scheme in terms of the horizontal working plane illuminance. This implies the use of techniques such as the lumen to gauge the suitability of the design. The code also advocates consideration of the gross brightness and colour pattern in our design[1], but it does not provide the tools and methods that would allow this.

Designing in terms of the gross brightness pattern is not new advice; as far back as 1954, Waldram gave us the Designed Appearance Method[2], which allows us to work from the apparent brightness of components in the visual scene to the required direct illuminances onto those components. Unfortunately this approach can be time consuming for complex schemes, and only considers the static view; no consideration is made for the sequential experience through the space. Overall, no tools or analysis techniques are widely available which allow us to tackle the problem of examining a lighting scheme at the design stage in terms of the three-dimensional gross brightness pattern as experienced by a user of the space as he progresses through it.

One recent approach to dealing with three-dimensional illumination is Chris Cuttle's *Cubic Illumination* concept. It extends our concept of illuminance in an architectural space from that of planar illuminance distribution on a plane of interest to that of a three-dimensional vectorial representation (i.e. he considers both the magnitude and direction of illuminances). His 1997 *Lighting Research & Technology* paper[3]

- Outlines the key technical concepts involved
- Gives an introduction to measurement techniques that can be used to measure cubic illuminance in the field
- Covers how to use the cubic illuminance measures to calculate a whole range of other measures, such as semi-cylindrical, scalar and spherical illuminance.

At the end of his paper, Cuttle describes his hopes for cubic illumination,

> *"It is a component of a system that stands on a foundation of illumination vector theory, which provides the bases of the system for predictive calculation and measurement of cubic illumination….*
>
> *It can be seen that cubic illumination is the key concept in a comprehensive system of applied photometry. A lighting designer who has tools to calculate, measure and analyse the cubic illuminances would be able to select lighting criteria from a range of metrics….*
>
> *The concept is equally applicable to daylight and to electric lighting, and has the potential not merely to quantify observed aspects of lighting, but to add insight to observation leading to clearer visualisation of lighting."*

He also provides a diagram, shown in Figure 1, which outlines the areas he sees cubic illumination being applied, along with their inter-relationship. There are three distinct strands to the areas covered,

- Visual performance through task illumination
- Visual adaptation through eye illumination
- Gross illumination pattern through three-dimensional forms and the flow of light

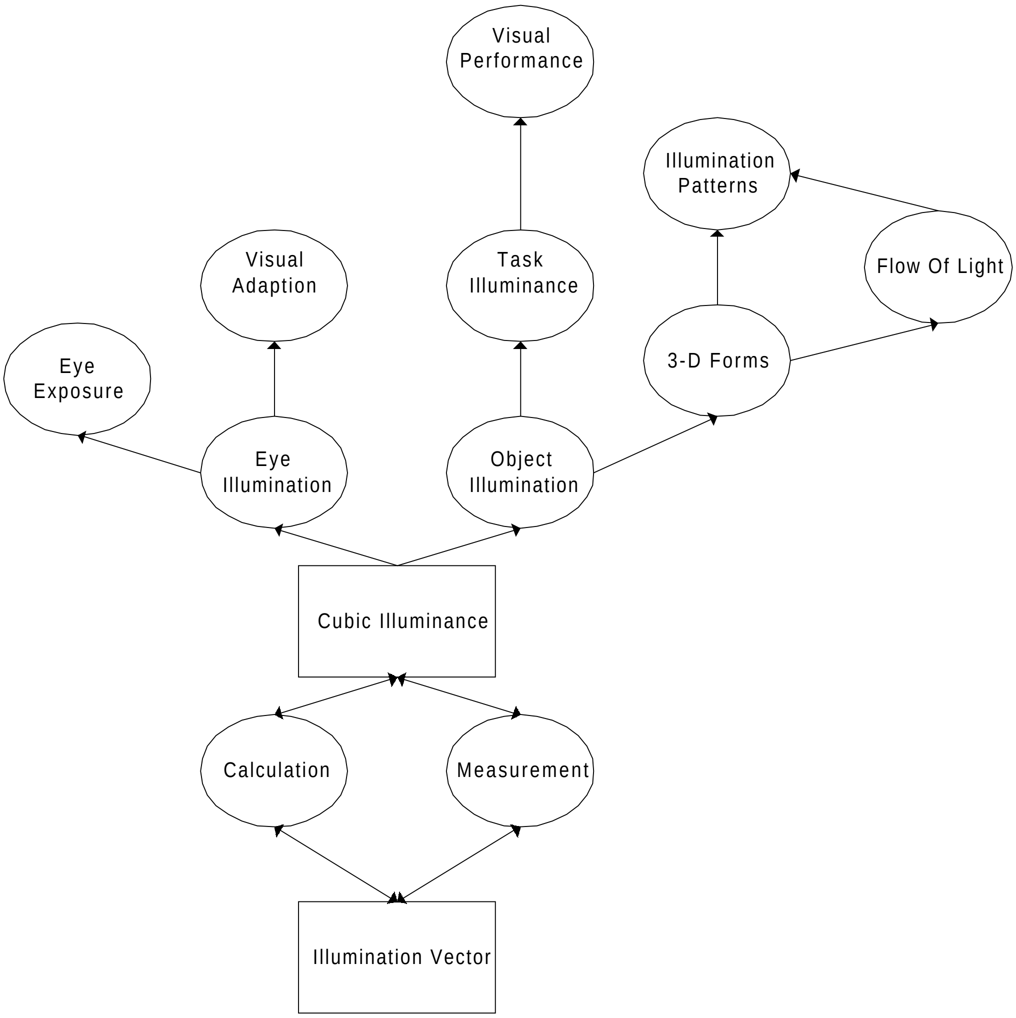

Figure 1 – Schematic representation of cubic illumination [from Cuttle (1997)]

Visual performance, in terms of task illumination, is a well-researched area and is closely linked to the illumination of the plane of interest, be it horizontal, vertical or inclined. Cubic illumination gives a methodology of obtaining those illuminances within a larger framework, and a way of easily comparing difference illumination metrics. Also, if we know something about the surface reflectances and textures, we can make an estimate of the luminance pattern viewed from any point in the room in any direction.

Since cubic illumination is inherently a three-dimensional system, it would seem logical that the areas in which it can provide added value over other systems is in its analysis of the three-dimensional view as experienced by a user in the space. This draws on the other two strands of cubic illumination, namely the state of adaptation of the user and the gross illumination pattern formed by the interaction of the lighting scheme with the three-dimensional forms in the space.

This leads us to our research proposal, an investigation of the suitability of three-dimensional illumination methods for examining the gross brightness pattern of a space.

Technical discussion

In this section, we look at the two papers that provide the starting point for our current research. They are

- Waldram's paper on *Designed Appearance Method*, a methodology for designing schemes by starting with the required relative apparent brightness and working back to the direct illumination required upon each surface
- Cuttle's paper on *Cubic Illumination* which provides a system for calculating the three-dimensional illumination distribution in a space.

Designed appearance method

J M Waldram in his seminal 1954 paper *Studies in Interior Lighting*[2], covers a wide range of topics, as shown in the start of his original summary for the paper,

> *The mechanism of interior lighting has been recently discussed in a Report of a Committee of the National Illumination Committee. The present paper records studies on three aspects of interior lighting: Adaptation, Revealing of Form and Texture, and Emphasis. It is proposed that effects of Adaptation should be taken into account in lighting calculations and values of apparent brightness accord much better with experience than figures of absolute luminance, and explain anomalies which are found with absolute figures.*
>
> *A new method of calculation is outlined, whereby the effects of inter-reflection can be found very simply without complex theory and by means suitable for day-to-day work, so that the installation can be designed from a specification of apparent brightness distribution.*

Waldram introduces the concept of *Apparent Brightness* as "a number which is possible a correlate of luminosity, obtained from curves due to Hopkinson[4] which relate it to the luminance field and to the adaptation level of an observer." This is motivated by the fact that we have a very crude sense of absolute level of luminance; having absolute luminance values for the surfaces in the field of view do not tell us how a scene will look without considering the *state of adaptation* that the viewer is in.

Figure 2 shows the inter-relation of the various factors that Waldram identifies as being a factor in interior lighting design. The most relevant part of this work is that which relates to the design of the Gross Apparent Brightness for a space. This is achieved by using his *Designed Appearance Method.* This requires the designer to specify the required apparent brightness of the components in the field of view of an observer. The adaptation state is calculated, and from Hopkinson's apparent brightness curves, the absolute luminance levels are obtained. Once the reflection properties of the components are known, we can then calculate the total illumination required. This is broken down into the required direct and indirect components, from which the lamp flux required to provide the direct illumination is calculated.

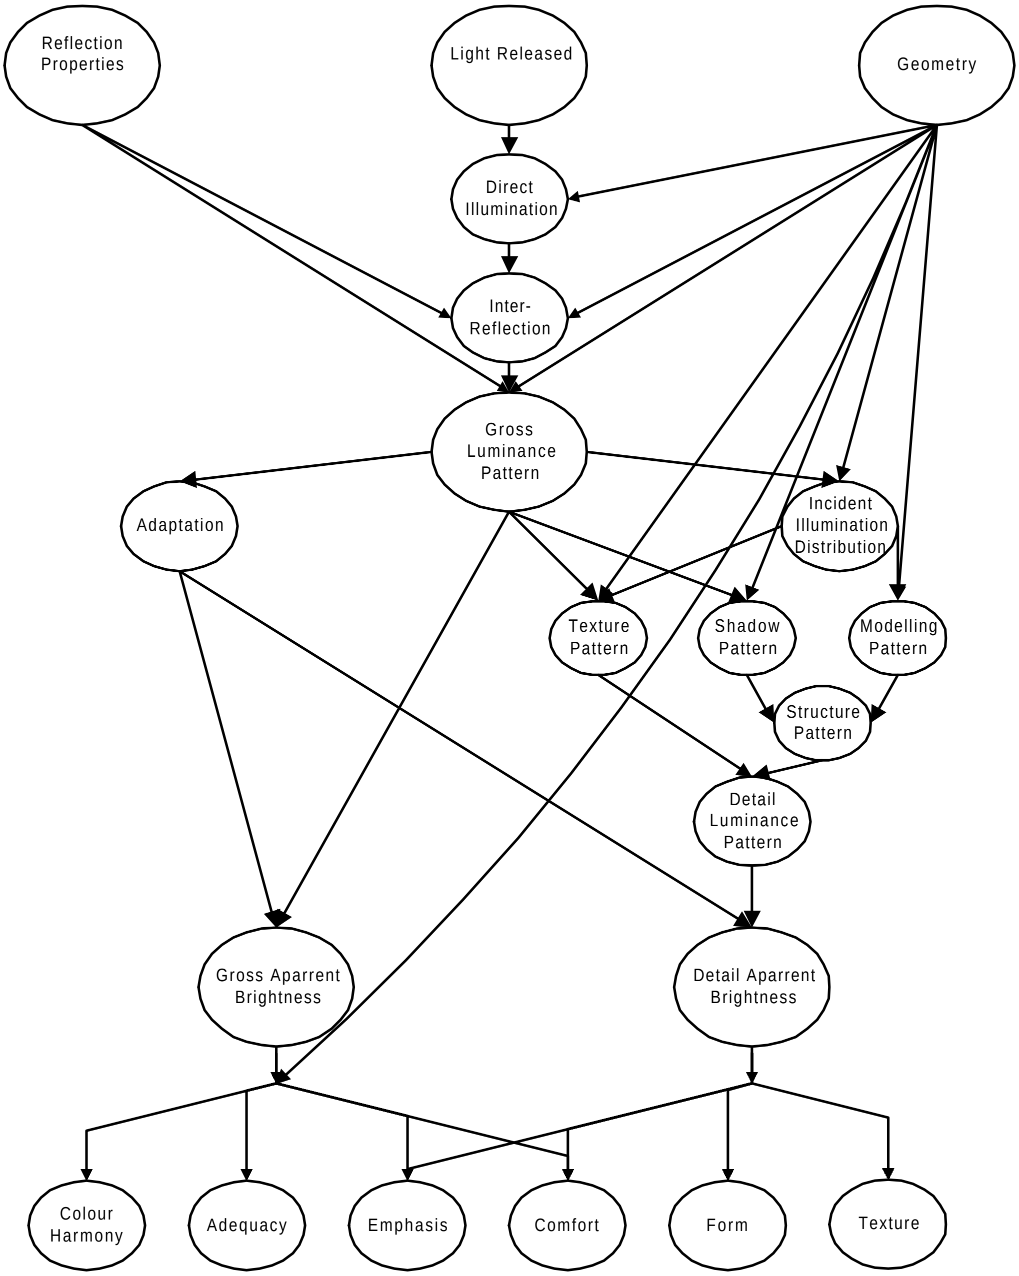

Figure 2 – Mechanism of Interior Lighting [from Waldram (1954)]

Adaptation level

Waldram defines *Adaptation level* as "the luminance of a uniform luminance field, in which the sensitivity of the eye is identical with its sensitivity in a stated luminance field.[2]" He notes that at the time of his paper the was not the means of ascertaining with accuracy this level. He takes it to be approximately the mean luminance of the field of view within 10^0 of the point of regard. Other measures commonly used are the mean luminance over the entire visual field or the illuminance at the eye[3].

Calculation of direct & indirect component

One of the main complexities in working from a preferred apparent brightness distribution back to an equipment schedule that could provide it is working out the components on each surface due to reflected light in the space. Waldram solves this without needed to do complex inter-reflection and distribution calculations as follows; each surface faces a "sky" consisting of all the other surfaces of the room. We know the luminances of each of these surfaces (since they are a function of the designed apparent brightness), and since the angular subtense of each is defined by the geometry of the room, we can estimate the illumination falling on the surface due to this "sky". He goes on to state that if a more accurate answer is required, we could treat it as a daylight analysis problem, and use a Waldram diagram or some other method.

Cubic illumination

Cubic Illuminance[3] is defined as the specification of the directional distribution of incident luminous flux at a point in space in terms of pairs of opposing planar illuminances normal to three mutually perpendicular axes intersecting at the point. A typical specification of cubic illuminance comprises six illuminances related to the surfaces of a small cube centred at the measurement point, with the surfaces of the cube aligned in accordance with the principle axes of the surrounding space, as in Figure 3.

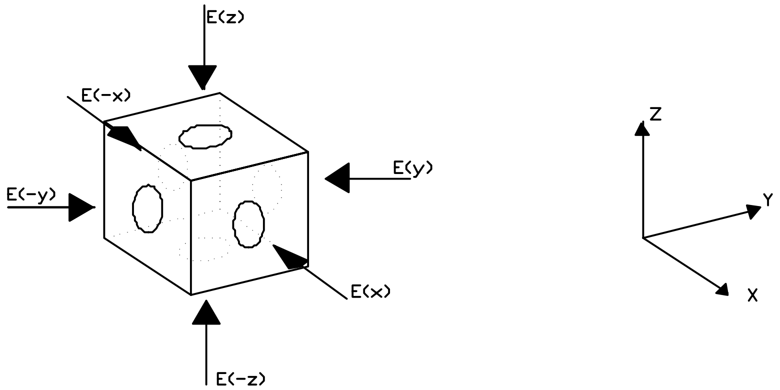

Figure 3 – Measurement of cubic illuminance at a point

There are three main vectorial concepts that are used in the paper,

- Position vector – this specifies points in vector space with respect to another point
- Illumination vector – this specifies the three-dimensional illumination distribution around a point
- View vector – this specifies the direction of view of an observer in the space

The position vector

We view the room as a three dimensional space with the origin chosen so that any point within it is defined in terms of positive values on the x, y and z axes. Then the position of point P in the space is defined by its *position vector* ***P***, such that $\boldsymbol{P} = (P_{(x)}, P_{(y)}, P_{(z)})$. The magnitude of ***P*** is indicated $|\boldsymbol{P}|$, and is the distance between the origin and P.

The distance between two points, ***S*** & ***P*** is given by vector subtraction, $|\boldsymbol{S} - \boldsymbol{P}|$. The direction of the vector ***SP*** is given by its *unit vector* ***q***, where

$$\mathbf{q} = \frac{\mathbf{S} - \mathbf{P}}{|\mathbf{S} - \mathbf{P}|}$$

The illumination vector

Considering illumination as a vector has been around since the turn of the century[5]. It is useful, since it leads to the concept of the *illumination solid* at a point, a three-dimensional surface where the distance from the point in any direction to the surface of the solid is proportional to the illuminance at the point normal to that direction. It acts as an aid to the visualisation of the three-dimensional illuminance distribution around a point. The point may be on the surface of the solid, or contained within it.

For any illuminance distribution, we can calculate the six cubic illuminances, $E_{(x)}$, $E_{(-x)}$, $E_{(y)}$, $E_{(-y)}$, $E_{(z)}$ and $E_{(-z)}$. They are the sums of the contributions on the *x*, *y* and *z*-axes of the tangent sphere illumination solids generated by every elementary source of light visible from the point. They specify six points on the surface of the illumination solid and together they define totally the vector component of the illumination solid.

The view vector

A direction of view is specified by a unit vector $v_{(x,y,z)}$. This is used to define the direction of view of an observer looking towards an illuminated object, or the direction of view of an observer who is viewing from the point.

Calculation of direct component

The calculation of the direct component of illuminance is where the power of vector calculus becomes apparent; we can use the *dot product* of two vectors to find the cosine of the angle of incidence, i.e. for unit vectors ***q*** and ***n***

$$\begin{aligned} \boldsymbol{q} \cdot \boldsymbol{n} &= \boldsymbol{q}_{(x)} . \boldsymbol{n}_{(x)} + \boldsymbol{q}_{(y)} . \boldsymbol{n}_{(y)} + \boldsymbol{q}_{(z)} . \boldsymbol{n}_{(z)} \\ &= \cos\theta \end{aligned}$$

where the direction at a point of interest *P* of a source *Q* is given by the unit vector ***q***, and the plane of the surface of interest at point *P* is given by ***n***.

From this we can restate the equation for the inverse square cosine law of illumination[6] as:

$$E = \frac{I\boldsymbol{q} \cdot \boldsymbol{n}}{|Q - P|^2}$$

To calculate the cubic illuminances, we simply calculate for three orthogonal values of ***n***, say *(1,0,0)*, *(0,1,0)* and *(0,0,1)*.

Calculation of indirect component

Cuttle's method of dealing with the indirect component of illumination is to estimate it as a constant for the space. He uses the indirect component of the *mean room surface illuminance* as a measure of ambient illumination. This is given by

$$M_{rs} = \frac{FRF}{A(1 - R)}$$

where *FRF* is the first reflected flux for the room, A is the area of the room and R is the average area weighted reflectance. Cuttle covers this in more detail in his paper on Sumpner's principle[7].

He justifies this approach by stating that when reflectances are low, the indirect component is small as compared to the direct component. Conversely when reflectances are high, there will be highly diffuse indirect illumination, so it will not vary much with respect to position or direction. Therefore he sacrifices some accuracy for simplicity in the calculation, much as Waldram did in his designed appearance method.

Vector and symmetric components of illumination

When the cubic illuminance at a point has been calculated by summing the components from all visible light sources and the indirect illumination, we can then derive a number of illumination measures from it. Again, this is another strength of the system – given a single set of illuminance values for points in a room, a wide range of metrics can be readily obtained by simple calculations. To do this we break the illumination solid at a point into two components

- The vector component *'E* – this represents the illuminance difference in the three directions of measurement, e.g. in the *x*-axis, $'E_{(x)} = |E_{(x)} - E_{(-x)}|$. In this way it represents all the *asymmetric* components of illumination.
- The *symmetric* component *~E* – This is obtained by subtracting the vector component from the illumination solid.

Cuttle points out that these two components have distinctly different properties,

> *"– The vector component is always equivalent to the illumination solid produced by a single light source that is located in the direction of the vector, and for which the normal illuminance* E_n *is equal to the magnitude of the illumination vector* $|\mathbf{E}|$.

> *– If a plane passing through a reference point is rotated, then at any angle of rotation the illuminances due to the* symmetric component *on opposite sides of the plane are equal…*
>
> *The fact that the illuminance distribution about any point in space can be analysed in terms of these two components, each having clearly defined properties, is, in the author's opinion, the single most remarkable finding to emerge from the study of illuminating engineering. Its direct relevance to this paper is that it provides the basis for estimating a variety of indices relating to characteristics of spatial illumination, and also it provides insight into the visualisation of the dynamic effects of lighting, sometimes referred to as the 'flow of light' concept.*[8]*"*

Designed appearance & cubic illuminance – the link?

The important point to realise in discussing Waldram and Cuttle's papers is that Waldram's paper gives us a design methodology for specifying how we want an interior to look, while Cuttle's paper purely gives us measurement and calculation system in which we are able to examine aspects of a design. Cubic illumination per. se., will not allow tell us how to design a scheme, but will provide us with tools that can give us more information about how a scheme performs. Waldram does present us with a methodology which goes "beyond the working plane", but it not simple enough to supplant the traditional two-dimensional view of working plane illuminance as the measure of a scheme. What cubic illumination may give us is a system that lets us express Waldram's concepts in a method that is more amenable to calculation, and that is the technical problem that this paper tries to examine.

Now look back to Waldram's map of concepts involved in designing lighting for interiors on page 5. We see that since surface reflectances, room geometry and pattern of light released all can be handled under the cubic illumination system, and therefore we should be able to examine all aspects of the mechanisms of interior lighting. Cuttle presents methods for calculating the incident illuminance distribution and adaptation state. We want to extend down another level in the diagram, and try to develop methods for examining the gross apparent brightness pattern.

Design

Luminance or illuminance?

As said previously, Waldram's designed appearance method works from a statement of the apparent brightness ratios of surfaces in the room, as seen from one viewer position, to the required direct illuminances needed. Conversely cubic illumination starts with a lighting scheme and calculates the illuminance distribution in three dimensions. Since the human eye directly perceives luminance patterns not illuminance, we must justify how we can use a system based on illuminance in examining the gross brightness pattern as seen by a human observer. We use arguments by Jay[9,10], Lynes[11] and Cuttle[12] to do this.

Jay[9] reviews experimental data that looks at the relationship between apparent brightness, luminance and adaptation. This includes the work by Hopkinson and Waldram that underpins the apparent brightness to luminance mappings used in designed appearance as well as work by Marsden and Saunders. The general conclusion he draws is that perceived brightness is not simply a function of luminance and adaptation level but is also affected by the luminance of adjacent elements in the visual field. Furthermore, he states that the perception of brightness is affected by contour. Lowry and da Palma[13] set up an experiment where they created a visual field consisting of an area of high luminance to one side and low luminance on the other, with a steady gradient between. The subjective brightness was enhanced at the transitions from the luminance gradients to the areas of constant luminance. This is done by *lateral inhibition*, which has the effect of sharpening edges, introducing Mach bands.

Jay draws certain conclusions from all the experimental work that I think strongly influence this work, and provides the basic justification for the methodology. He states

- As illumination varies, adaptation will ensure that variations in apparent brightness will be much less than variations in luminance
- All contours in the visual field, and variations in brightness on either side of a contour, will be emphasised
- Within a contour both adaptation and persistence of vision will tend to reduce the visual significance of any relatively small changes in luminance

The ordinary scene therefore will consist of sharp luminance gradients associated with object boundaries, within which there will be gradual changes in luminance, due both to the fine structure of objects and changes in illumination across the object. Jay concludes that where gradients of illuminance are much more gradual than gradients of reflectance, the perceptual process will tend to ignore the slow gradients. Therefore, when objects in the visual field vary widely in reflectance, perception will therefore relate more closely to relative reflectance than to relative brightness.

Lynes[11] looked at circumstances required for maintenance of perceptual constancy in the visual field. He concludes that they are equivalent to those found in a normal building that would meet the current recommendation for good lighting design – adequate uniformity, relatively high reflectances of major surfaces, high diffuse component, avoidance of glare and so forth. When the visual field is full of strong contrasts and the sources of illumination are not obvious, such as stage lighting or display lighting, constancy may break down completely.

But where does this leave Waldram's Designed Appearance Method, which is based, as Cuttle states[12], on the premise that the appearance of a surface is determined primarily by its luminance? Waldram demonstrated his application of the designed appearance method using his lighting schemes for a series of British Cathedrals such as Ely. Cuttle suggests[12] that in these situations the breakdown of constancy has been engineered, and in such a situation a luminance-based approach is applicable. For situations where modelling and emphasis are not of concern, and we want uniformity and avoidance of glare and shadows, an approach based on the manipulation of the reflectance and illumination of surfaces is more appropriate.

Lynes' relative surface illuminance method

Since we want to look at the general case of examining the gross brightness pattern of normal interiors such as an office, we obviously now cannot use a purely luminance based method. Lynes[14] describes a method that specifies a design in terms of the relative surface illuminances. It does not limit itself to a single point of view, but calculates the flux distribution for all the surfaces in the room. This takes the same design approach as Waldram, but instead of specifying the apparent brightness of each surface, the designer specifies the required surface illuminance and reflectance. We use a version of this from Lynes[15], summarised in Appendix A.

CalcuLuX analysis

Once we have the required flux distribution on the surfaces in the room, we can then design a luminaire layout that will achieve this distribution. This scheme then can form the basis for our cubic illumination analysis. In order to create the scheme we do the calculations by hand, or use a computer tool such as CalcuLuX. This is a software tool, made by Phillips, that aids lighting designers select and evaluate lighting. It can produce such data as luminance or illuminance distributions for any surface in the room. The value in us using such a tool is that it provides a validation mechanism for the cubic illuminance analysis.

What we now seem to have is a design method that starts with the required illuminance and surface reflectances, and calculates the required direct flux on each surface. From this we work out the required luminaire layout. Then using the room and surface geometry and reflectance information in combination with the luminaire layout, we need to compute the three-dimensional illuminance distribution in the room. From Figure 4 below, we see that we are only one step away from coming full circle in the design process and getting back to the illuminance ratios of the surface, but this time in a three-dimensional form.

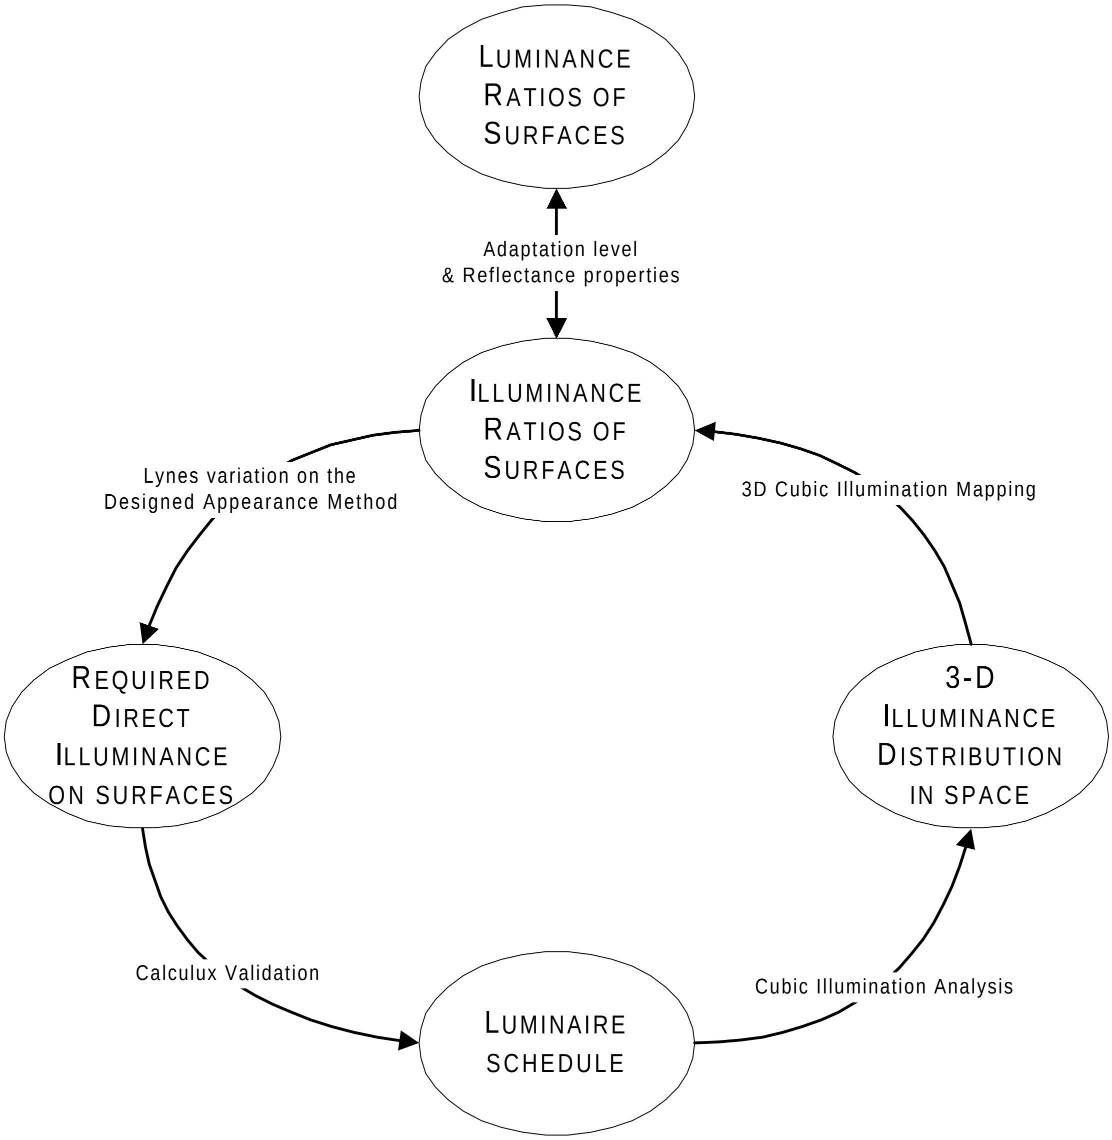

Figure 4 – The steps of design process, showing a loop

Three-dimensional Illuminance

We have now defined the general problem we want to tackle, i.e. going from a design in terms of relative surface illuminances to the gross brightness pattern as seen by an observer in the space, and the main steps in the procedure. Now we look in more detail at the final step in the procedure, the mapping from the three-dimensional illuminance distribution to the gross luminance pattern from a point in the room.

Currently researchers wishing to visualise a space in terms of the three-dimensional luminance pattern use two main computational methods – Ray Tracing and Radiosity.

Ray tracing

Ray tracing works on the principle of tracing rays of light from a source (i.e. photons), and seeing which of them would hit our eyes if we were standing in the space, as in Figure 5. This is highly inefficient, since only a small proportion of the light bouncing around in a room actually ever reaches our eyes. This leads us to the concept of *backward ray tracing* where we start at the eye, sample all possible rays leaving it and see which objects they hit. This means we only follow rays that we can actually see. A good overview of ray tracing is given in Glassner[16]. Ray tracing essentially measures the direct component of illumination, which allows it to accurately model specular reflection, producing photo-realistic images. RADIANCE[17] is a ray tracing software package that uses a physically correct model of light that can produce photometrically accurate architectural renderings. It has been used extensively to examine both electric light and daylighting schemes. It has components that accurately model the response of the eye in terms of adaptation, glare and contrast. Such a ray tracer is able to show three-dimensional representations of the space in terms of both the luminance and illuminance distributions achieved.

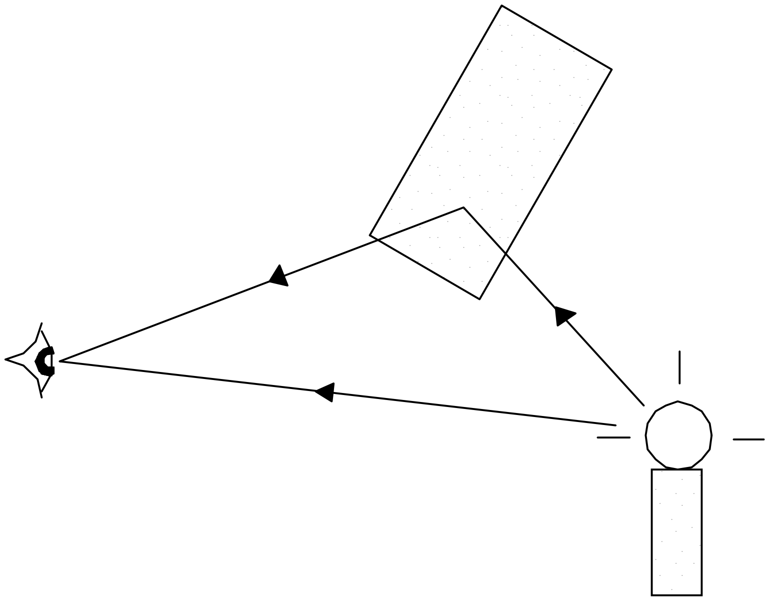

Figure 5 – An example of forward ray tracing

Radiosity

Radiosity is a different approach to rendering explained fully in Ashdown[18]. It makes three major simplifying assumptions,

- All surfaces are ideal diffuse and opaque reflectors
- Light sources are ideal diffuse emitters of light
- Each surface is subdivided into a mesh of elements called patches

The key assumption here is that all surfaces are diffuse (Lambertian) reflectors and emitters. This means that the surface has a constant radiance or luminance independent of the viewing direction. We can use the *Cosine Law* to predict the luminous intensity at any angle to the surface, i.e.

$$I_\theta = I_n \cos\theta$$

where I_θ is the luminous intensity measured normal to the surface. Since we know the geometry of the surfaces in the room, and each light source is Lambertian, we can quickly calculate using the cosine law, the amount of flux that is received by each surface in the room due to the light sources.

Also, these surfaces, if they have non-zero reflectance, will re-emit some of this light back into the room. Since they are Lambertian, we can easy compute their re-emitted flux as the product of the received flux and their reflectance. If we iterate this process until there is no more reflected flux, we finish knowing the luminance of each surface. For Lambertian surfaces it is simple to calculate the illuminance, given the luminance,

$$E = \frac{\pi \cdot L}{\rho}$$

where E is the illuminance in lux, L is the luminance in cdm^{-2} and ρ is the reflectance of the surface.

Radiosity versus ray tracing

In ray tracing the view position is very important. All rays that are traced start at the view position and make their way into the space. If the view position is changed then it usually requires for the scene to be re-traced from scratch; ray tracing is an example of a *view-dependant* process. In contrast, radiosity is *view-independent*. It deals solely with the geometric distribution of sources and surfaces within the space. Once we calculate the distribution of flux on the surfaces, then we can position the viewer anywhere in the room, and move it around at will without having to recalculate the luminance distribution on the surfaces in view.

Of course, radiosity is not perfect; it can only model Lambertian surfaces efficiently. It can only model semi-specular surfaces with about the same amount of computational effort as ray tracing, and cannot handle specular surfaces at all[18].

Cubic illumination analysis

Now let us look at where this leaves the cubic illumination method proposed by Cuttle. In calculating the illuminance at a point he separates out the direct calculation from the indirect calculation.

For the indirect calculation we only have to one flux transfer calculation for the entire space, since we use a mean room surface illuminance as an estimate for all points in the room. This is easy to calculate, since we know the size and reflectance of all the surfaces in the room.

For the direct calculation, three separate point source illuminance calculations are required for each grid point since we use three different orthogonal surface normals to resolve the illuminance into its three component parts.

Upon looking closer at this calculation, we can see certain flaws. Cuttle does not deal with objects in the space occluding the light sources when calculating the distribution through the room. He gives the basic formulation for the cubic illumination at a point – the sum of the direct illuminances due to all the light sources in the room. This is the same method used by programs such as CalcuLuX. It would seem that we could quite easily update the algorithm to compute the cubic illumination at a point to include a shadow test – we calculate the vector direction from the point of interest to the light source, and then follow this vector back to the light source. If we pass through an object in the room before getting to the light source, then the point of interest would be in shadow with respect to the source, and the illuminance contribution from the source should be discounted.

Comparison of cubic illumination versus ray tracing and radiosity

Cubic illuminance, on closer examination, seems to be an simplified integration of ray tracing and radiosity which glosses over all the difficult implementation details that are required to accurately model real spaces. The indirect calculation is equivalent to carrying out the first step of the radiosity calculation, and then computing the mean value over all the surfaces. While this should be acceptable for office spaces with diffuse surfaces and average reflectances, it is not going to produce accurate results when we have a wide range of reflectances in a space.

Also, if we look at the direct component calculation. In order to place objects into the room, we need to do shadow testing for each point in the space. Glassner states that this shadow test calculation is the bulk of the computation work required in a ray tracing solution[16]. Therefore, why not just use ray tracing and gain all its added advantages, such as the ability to handle specular and semi-specular objects.

In order to help us balance the pros and cons of all these three methods and choose what would be best, let us look again at the aims of this final stage of the process. The goal is to quickly produce rendering of the three-dimensional gross luminance pattern as seen by an observer in a typical office space in a reasonable amount of time. Since we are aiming to look at simply the gross luminance pattern, and not the fine detail luminance pattern, ray tracing techniques are not needed. This means that the cubic illumination grid calculations are too expensive as well, if we want to accurately model the effects of objects in the space. Since these are really needed in order the model the luminance pattern in complex spaces, this leaves us with the radiosity method. It satisfies the requirements in that it can provide fast renderings of the gross luminance pattern in a space containing objects. Also since it is view independent, we can produce the luminance pattern from different observer positions and orientations without having to do expensive recomputations.

The HELIOS software system

Ashburn[18] provides guidance on how to efficiently implement the radiosity method. He provides an example of this in a software package called HELIOS. It can model arbitrarily complex rooms that adhere to the three guidelines mentioned on page 16. This can produce renderings of the room from any viewpoint and orientation. It is designed to produce visually pleasing renderings rather than photometrically correct ones.

Given a geometric distribution of surfaces with initial reflectance and radiant exitance, HELIOS performs the following steps,

- Perform a radiosity calculation to determine the final flux distribution on all the surfaces in terms of spectral exitance
- Determine extent of visible scene at a given view position and orientation
- Render the scene. This uses a process called *polygon scan-conversion*, described in Foley[19], to map from the polygons that make up the surfaces in view to the scan lines that make up a two-dimensional picture.
- Perform tone reproduction on the image. This is designed to compress the wide range of luminances that appear in the normal environment to the dynamic range of the eye. This normally needs to map luminances in a range of 10,000:1 to a range or 100:1. HELIOS uses a process known as *exitance normalisation*. All surface exitances are scaled according to the vertex with greatest *reflected* exitance. This is assigned a value such that the greatest value of its colour bands is slightly less than 1. If the exitance of a light source exceeds this value, it is scaled to equal one. Thus light sources will appear as the "brightest" objects in the scene. Exitance normalisation has no physical justification – it simply produces results that are visually appealing[18]. Obviously this may not be good enough for us if we want to extract information out of the scene such as surface luminance and illuminance.
- Perform gamma correction to compensate for the non-linear output response of the output device that is being used.

We end up with a 24-bit RGB colour image that is now suitable for display on a computer screen. An example of such an image is shown in Colour Plate 1 located in Appendix B.

Overcoming the tone reproduction problem

We would like to be able to extract numerical data from our output images, such as the actual luminance at a point in the scene. Obviously in the exitance normalisation approach taken by HELIOS, we lose some of this information. Nonetheless, we can extract a reasonable amount of information from the images, at least about the reflected surfaces. Since these have all been scaled down by a constant factor, the maximum reflected exitance, it possible, by multiplying the image values by this scale factor, to obtain the actual spectral exitance, at any point in the picture. Unfortunately we cannot find the spectral exitance for the luminous surfaces of the light fittings by this method.

Once we have spectral exitance in terms of the red, green and blue bands of an RGB pixel, we can obtain the luminance by the following formula, from Ward[20],

$$L = 179 \cdot (0.3R + 0.59G + 0.11B)$$

that models the scotopic response of the human eye with respect to the RGB primaries. 179 lumens W^{-1} is taken as the luminous efficacy of equal energy white, and R, G, B are the spectral exitance of the red, green and blue bands respectively.

Simplification of light sources by using diffuse sources

Since the radiosity technique only allows for diffuse light sources, we approximate all light sources in our radiosity model by a diffuse source with the same luminous area and lamp flux. This should be able to model the effects of most regular arrays of linear luminaires, as are common in office spaces. More problematic are those luminaires, such as spotlights, which have a highly directional output. Obviously, this is a major weakness of the radiosity technique as it stands.

Adaptation level and human response

The RADIANCE software package provides tools that let us model the effects of the human eye upon viewing a scene. They take as input a set of spectral radiances. RADIANCE also provides tools that convert from the output of HELIOS into such a file. In particular RADIANCE can model the effect of the adaptation state of the observer. It also can use the human contrast sensitivity function in order to determine the exposure for the picture, and apply more accurate tone reproduction algorithms[21]. Using these, it is possible to produce more physically realistic renderings than those produced by HELIOS.

Radiosity step of the design procedure

After that brief, but necessary, excursion into computer graphics techniques, let us summarise the steps in the experiment concerned with the generation of the three-dimensional gross brightness pattern and its analysis. This uses parts of the HELIOS and RADIANCE software packages.

- We start off with the room geometry, and the luminaire pattern supplied to us by CalcuLuX. We create an input file for HELIOS containing this information, approximating the photometric output of the luminaires by diffuse emitters of the same luminous area and initial lamp flux
- Using HELIOS, we perform a radiosity calculation, which gives the spectral radiant exitance distribution for all the surfaces in the room

- HELIOS then renders a scene as seen from a particular viewpoint and orientation. It performs exitance normalisation as a form of tone reproduction in order to produce a 24-bit RGB image file. We can change the point of view and orientation easily, and the new scene will be rendered quickly, giving us a form of interactivity with the space
- We convert this RGB image file into a form that RADIANCE can understand, and then undo the exitance normalisation performed by HELIOS
- We modify the image, modelling the response of the human eye in terms of adaptation state, exposure and dynamic range
- Finally we can use RADIANCE's image viewing facilities to look at the image. At this stage, we can also produce false-colour images that show us the luminance distribution on the surfaces. Since we know the reflectances of the surfaces, and they are Lambertian, we can find the surface illuminances too.

Results

A sample room

Now we describe our findings when we carried out this design procedure on a simple scheme. Firstly, let us define a sample room, shown in Figure 6. It is a typical small office space and its geometric properties are as follows,

- Dimensions, Width 5m, Length 8m, Height 3.2m
- Surface reflectances, ceiling 0.7, side and front walls 0.5, back wall 0.6, floor 0.3
- The working plane height is 0.8m

There are several objects in the room that we want to highlight,

- Two office desks, one with low reflectance 0.2, and one with high reflectance 0.6
- A picture on the back short wall

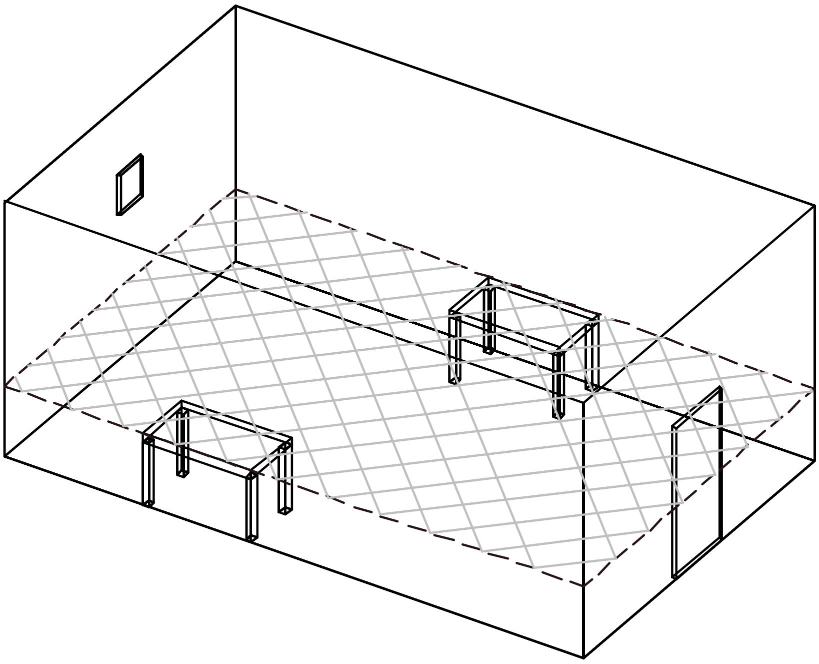

Figure 6 – Three-dimensional representation of a sample room

Lynes variation on the designed appearance method

We now calculate the required direct flux for the surfaces in the office. This is shown in Table 1. The assumptions we have made are as follows

- We treat the back wall as the key surface around which we design the relative illuminances. The picture on it, along with a change in surface reflectance can form an area of visual rest.

- We choose the illuminance level on the ceiling such that it is lit purely by ambient illumination

- We set the desks (i.e. the working plane) at 500 lux, per CIBSE requirements. The floor is set at 280 as an approximation of how much of this 500 lux would travel to the floor. This is calculated assuming 500 lux is provided on the working plane as direct illuminance from a downlighter. Then the amount that will hit the floor according to the inverse square law of illumination is

$$E_{floor} = E_{wp} \cdot \frac{d_{wp}^2}{d_{floor}^2} = 500 \cdot \frac{2.4^2}{3.2^2} = 281\,\text{lux}$$

- The floor area used is assumed to be the total area of the floor minus the area of the desks, and the area of the back wall used is the area of the actual wall minus the area of the picture hung on it.

(1) Surface Name	(2) Area A_s	(3) Reflectance R_s	(4) Illuminance E_s	(5) Reflected Flux $A_sR_sE_s$	(6) Direct Illuminance $E_s(d)$	(7) Direct Flux F_s
Roof	40	**0.7**	**100**	2800	7	272
Back Wall	15.75	**0.6**	**200**	1890	107	1682
Front Wall	16	**0.5**	**200**	1600	107	1709
Left Wall	25.6	**0.5**	**200**	2560	107	2734
Right Wall	25.6	**0.5**	**200**	2560	107	2734
Floor	37	**0.3**	**280**	3108	187	6911
Picture	0.25	**0.5**	**750**	93.75	657	164
Dark Desk	1.5	**0.2**	**500**	150	407	610
Light Desk	1.5	**0.6**	**500**	450	407	610
	Total Area		:	163.2		
	Total Reflected Flux		:	15212		
	Indirect Illuminance		:	**93**		

Table 1 – Required direct illuminances on room surfaces

CalcuLuX analysis of scheme

Into this room, we are going to put a lighting scheme typical for such an office; it will consist of several subsystems,

- A regular array of linear fluorescent luminaires to provide general lighting to the space

- Tungsten-halogen spotlights to highlight objects such as the picture, giving areas of visual interest

Since the radiosity method works best with diffuse sources, we pick a linear fluorescent that has a reasonable approximation to a Lambertian emitter. This is a Zumtobel Miral RAL[22]

1/58W, which is a standard surface mounted louvre luminaire. It takes a 58W T26 fluorescent lamp with an initial luminous flux of 5,200 lumens.

For the spotlight onto the picture, we pick a low voltage directional recessed spotlight, the Zumtobel CNS70 1/50W[22]. This uses a 50W QR-CBC51 lamp that has an initial luminous flux of 990 lumens. The required luminaire layout in the room is shown in Figure 7. We summarise the results in Table 2.

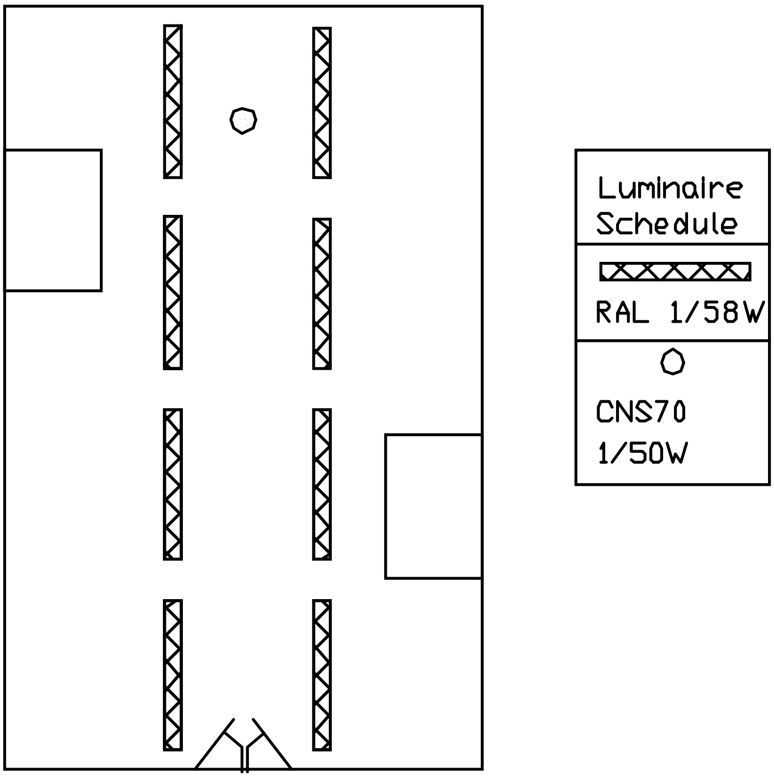

Figure 7 – Luminaire layout in sample room

	Average Illuminance (lux)			Average Luminance (cd/m^2)
	Total	Direct	Indirect	
Floor	360	264	95	22.9
Working Plane	433	342	91	N/A
Ceiling	107	0	107	23.8
Front wall	208	121	87	33.1
Left wall	208	114	94	33.1
Back wall	229	143	86	43.8
Right wall	208	114	94	33.1

Table 2 – Summary of photometric results from CalcuLuX analysis of sample room

Radiosity analysis

Firstly we create input files or HELIOS that represent the room and luminaire geometry from the previous stage. We approximate the linear fluorescent by a diffuse source of 5,200 lumens, and the QR-CBC spotlight by a diffuse source of 990 lumens. We carry out the radiosity calculation for the scene as seen from the position in Figure 8, and obtain a scale factor of 105 for the rendered image. This means the luminance of the brightest reflected surface is 105 cd/m^2. We use this value to undo the exitance normalisation, and input the rendered image into RADIANCE. As output, we produce two images – one that shows how the scene will look to the human eye, Colour plate 2, and one that shows the luminance distribution in the scene, Colour plate 3.

We repeat the above process for a different viewpoint, as shown in Figure 9. This produces two more images, Colour plate 4 and 5. Finally we compute a view as seen by an observer standing just inside the room, to the left of the door, Figure 10. The human eye corrected image and false colour image are shown in Colour plate 6 and 7 respectively. The view point and orientation information used to generate these views is shown in Appendix C.

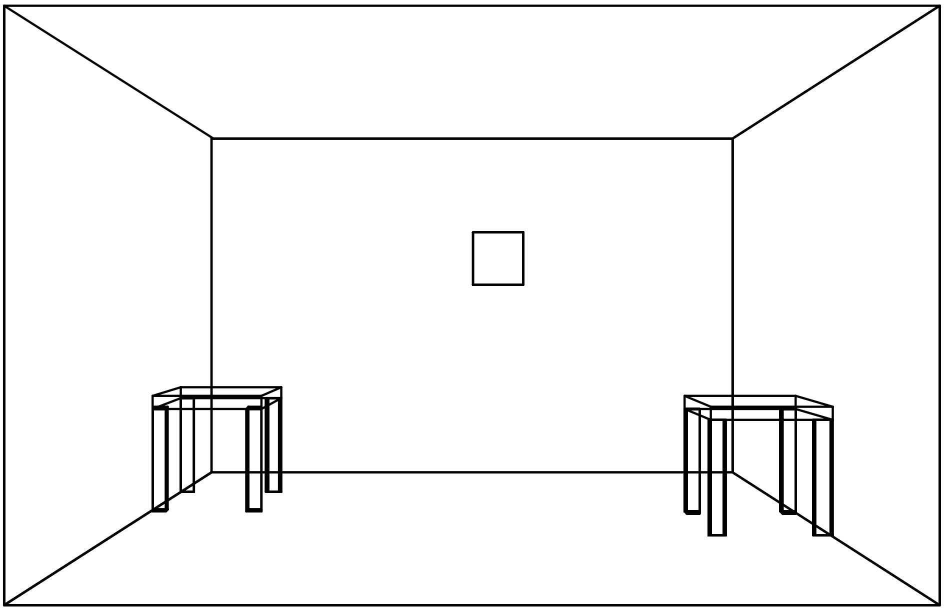

Figure 8 – Front view of room

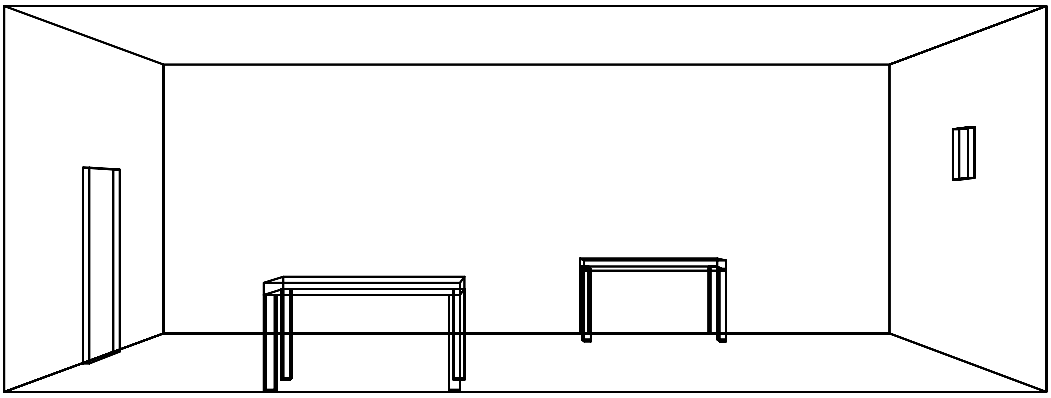

Figure 9 – Side view of room

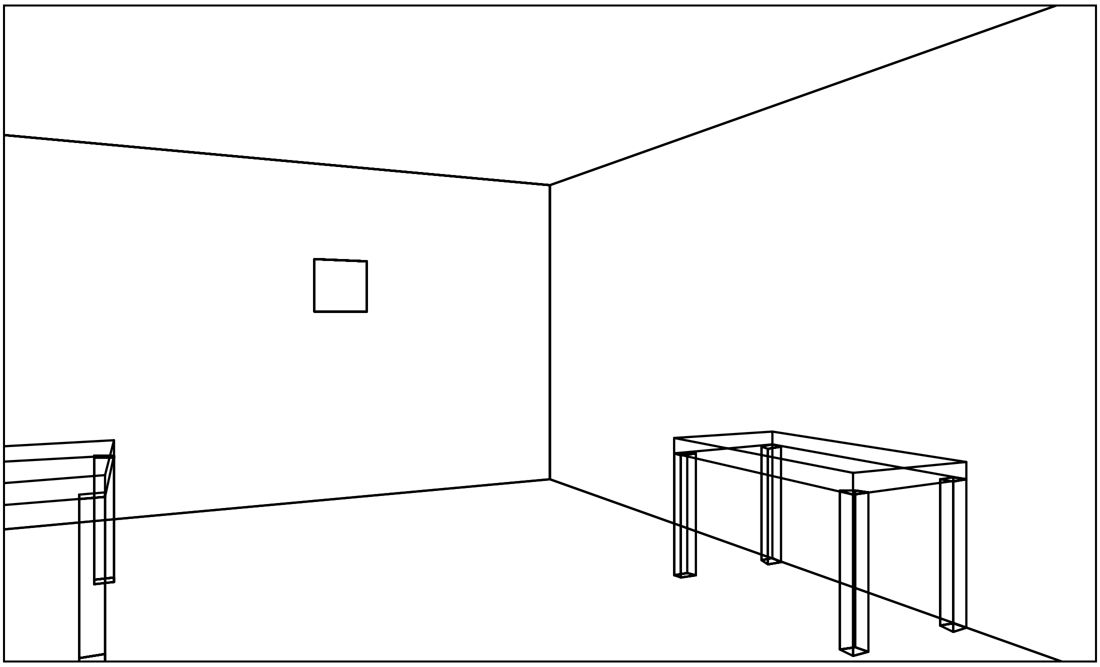

Figure 10 – View of sample room from inside

Analysis

We would now like to look at what photometric data we can extract from the images and examine its accuracy. Using the RADIANCE viewing program XIMAGE, we can select areas of an image and it computes average luminance. We carried out this procedure for all the surfaces visible in each of the three images. The results are shown in Table 3.

Luminance(cd m^{-2})	*Front View*	*Side View*	*Interior View*	*Average*
Ceiling	23.7	22.6	15.6	20.6
Back wall	46.4	48.1	49.4	48.0
Front wall	N/A	36.9	N/A	36.9
Left wall	38.0	32.7	35.6	35.4
Right wall	39.8	N/A	41.5	40.7
Floor	29.1	26.7	26.7	27.5
Picture	63.6	63	69.5	65.4
Dark desk	21.4	22.0	15.5	19.6
Light desk	64.2	66.6	58.7	63.1

Table 3 – Luminance of elements in field of view

There are two things we now can compare this against; the predicted luminances obtained from CalcuLux, and the required average surface illuminances used to initiate the design process. This is shown in Table 4 and Table 5.

The first thing we notice is that the average error between the CalcuLuX predicted luminances, and the luminances calculated is only 10% for the major room surfaces, i.e. ceiling, walls and floor. This corresponds to an average error of about 3 cdm^{-2}, which would not be noticeable. Comparing the illuminances calculated, in Table 5, these concur with both the CalcuLux produced results and the original design criteria for all the major surfaces, with about 50 lux extra on the walls. When we then go to look at the illuminances on the picture and desks, we notice more of a problem. The picture is about 50% (350 lux) too low, while the desks are 40% (200 lux) too low.

The cause of this, I suspect, is the modelling of luminaires in HELIOS. Since we had to use a diffuse emitting source for the general array, this means less of the flux is getting pointed down onto the working plane, and more on the walls. In fact if the extra 50 lux per wall was redirected downwards properly, the illuminance figures for the desks would then be correct. A similar story applies for the spotlight on the picture – a diffuse emitter of the same luminous intensity as the spotlight would not deliver anyway near as much flux onto the picture.

Luminance(cd m^{-2})	*HELIOS*	*CalcuLux*	*Error*	*Error (%)*
Ceiling	20.6	23.8	-3.2	- 13.4
Back wall	48.0	43.8	4.2	+ 9.6
Front wall	36.9	33.1	3.8	+ 11.5
Left wall	35.4	33.1	2.3	+ 6.9
Right wall	40.7	33.1	7.6	+ 22.9
Floor	27.5	22.9	4.6	+ 20.1
		Average:	3.22 cd m^{-2}	9.6%

Table 4 – Comparison of predicted luminance between HELIOS and CalcuLuX

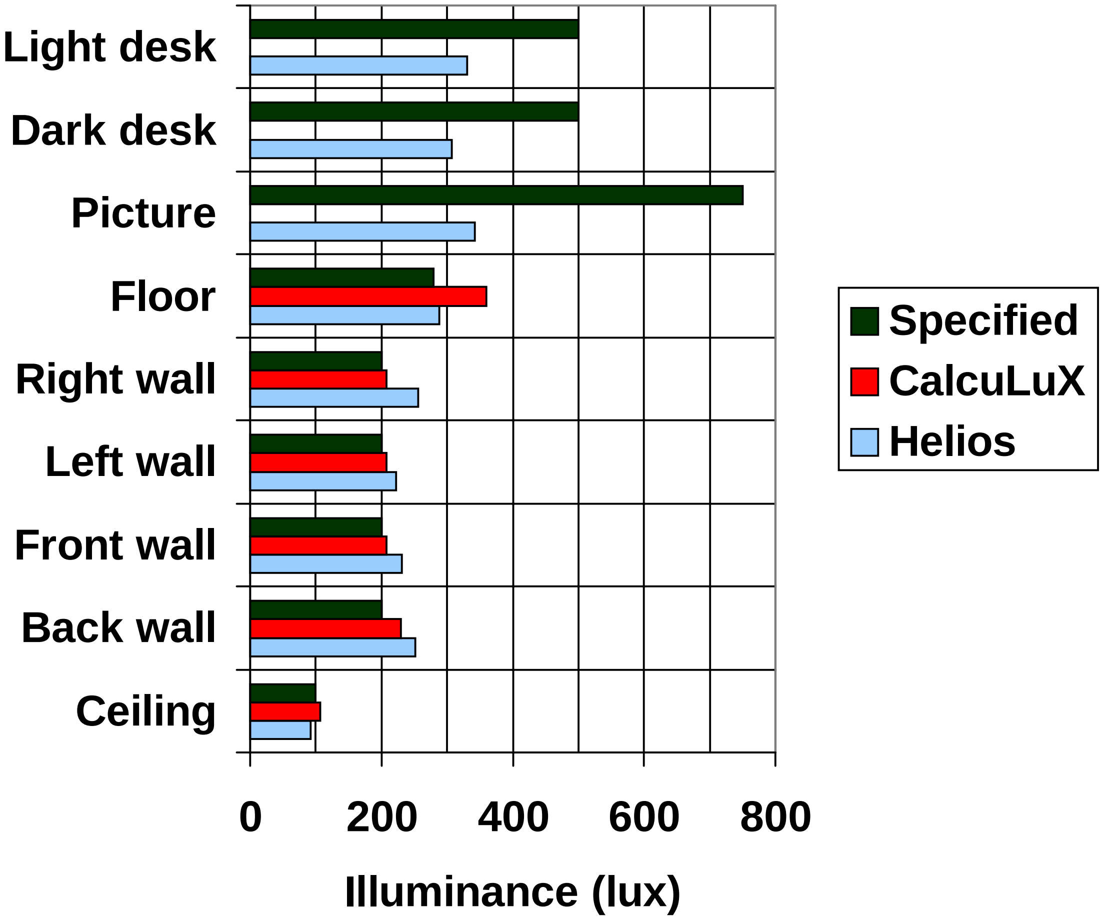

Table 5 – Comparison of Illuminances

Conclusion

We now look back at the technologies we have considered in this paper, and how they measured up to the task of examining the gross brightness pattern in an architectural space.

Designed appearance method

The Designed Appearance Method, grounded firmly in apparent brightness and luminance based specification, never quite managed to become the solution to the design of the visual field that it hoped to be[12]. Nonetheless, Waldram's paper[2] still remains one of the best statements of the principles of interior lighting. It obviously is a very successful design methodology for some purposes, as seen in Waldram's Cathedral designs. But the fact that it relies on apparent brightness, and the breakdown in perceptual constancy that this requires, has been the mean reason that it cannot be applied to many problems. Lynes[11] shows that the maintenance of constancy is a major requirement for the diffuse, uniform, limited glare lighting strategy that is required for most of our buildings.

Lynes's relative surface illuminance method

Lynes's modification to the Designed Appearance Method takes us into the realm of illuminance based design methods. Being based on design in terms of both illuminance and surface reflectance, it works well in those situations when constancy is upheld. In our example, the specifications for illuminance on the walls, floor and ceiling were quite easy to fulfil, but the areas of highlight in the example space, such as the picture, provided more of a problem. It was quite hard to find the luminaire needed to place the correct amount of extra flux required after the main array has lit the object. In the end the easiest way was trial and error with a range of luminaires and luminaire mounting positions.

Lynes himself notes that the method has drawbacks[15], in that it may lead the designer to place too much emphasis on the appearance of the bounding surfaces of an empty space at the expense of objects in the space. This was not too much of a problem in our example, since radiosity also puts more emphasis on the large surfaces in the room, since they tend to receive and emit the bulk of the flux in the room.

A better design method for this stage would have a been a two step process –

1. Use the Lynes spreadsheet to calculate the general lighting from the main array of luminaires in the space
2. Once the direct and indirect flux on the areas of interest due to the main array have been determined, use point source calculations to work out the luminaires needed to provide the local highlighting.

Overall, I think that this method is an excellent way to quickly bulk out the main elements of a design and see what effect they will have. If we only deal with the main arrays, it would be possible to even place these into a space automatically with our computer design system and see the effect they have.

Ray tracing

Ray tracing was quickly discounted in this paper, due to the requirement that we wanted a quick view-independent method for calculating the gross brightness pattern. This is not to say that it has no use in lighting design. In fact, RADIANCE is the only tool widely available that allows us to model all forms of lighting photometrically, including daylighting. It can be seen as complementing radiosity, in that it allows you to paint in the fine detail into images over the gross brightness pattern created by radiosity methods.

Radiosity

Radiosity, which is the cornerstone of our design method goes a long way to solving our problem, in that it is a view-independent calculation technique that can quickly calculate the approximate distribution of all the luminous flux in a room. Its big falling down point, insofar as producing photometrically correct results currently lies in its lack of ability to model luminaires with non-Lambertian distribution. Some headway has been made to rectify this. Ashdown has introduced the concept of *Near-field Photometry*[23], which allows modelling of arbitrarily complex light sources in computational systems such as ray tracing and radiosity[24].

Ashdown[25] has also very recently looked at the use of three-dimensional illuminance measures in modelling complex environments. He presents an algorithm for implementing a *Virtual Photometer*, which can measure the illuminance at any point in a space. He quotes errors in illuminance prediction of less than 5% when compared to inverse square calculations on a point by point basis. This seems to be a step ahead for radiosity based systems, since it allows now allows accurate prediction of the direct component on top of the indirect component.

Cubic illumination

So where does this leave Cubic Illumination? Well, as an analysis system, it gives us easy access to a wide range of illumination metrics. Unfortunately, to my mind, it compromises its potential as a real three-dimensional illumination system by trying to provide ease of calculation. The simplifications necessary for this means its use is limited when the goal is rendering the three-dimensional gross brightness pattern.

The important concept to keep in mind is that the features of Cubic Illumination can be seen as a subset of those available in more complicated three-dimensional illumination systems. As we said earlier, in order to compute the cubic illumination distribution when there are objects in the room, we need to carry out the shadow testing that we would carry out in a ray tracing system. Thus, it would also be easy to gather up the cubic illumination information when computing scenes in a ray tracer.

The ideal solution?

I think the ideal computer based design system would be a combined radiosity and ray tracing system with illuminance metric analysis tools such as provided by cubic illumination. It would,

- Have a simple graphical user interface similar to CANDLE, an interface to the RADIANCE system developed at the Bartlett.
- Use Lynes relative surface illuminance method to calculate layouts of luminaire grids to provide general lighting
- Use standard radiosity techniques to compute a model that let the user interactively review how the main room surfaces would be illuminated with respect to a variety of view points
- Use Ashdown's Virtual Photometry methods to model directional light sources that can provide highlighting and local illumination.
- Use ray tracing methods to produce final photometric accurate renderings for presentation to clients.

This does not exist as a single product, but all the individual components do. In my opinion, this would provide a powerful aid in letting the lighting designer go beyond the Working Plane, and design in all three dimensions of a space.

Appendix A – Procedure for design in terms of relative surface illuminances

- List all significant visible surfaces in column (1) of Table 6 and fill in details of areas and reflectances in columns (2) and (3)
- Pick an 'anchor' surface for which you judge a certain illuminance to be appropriate, and specify that illumination in column (4)
- Complete column (4). Your skill in design is how you relate the illuminance to be provided on each surface to the anchor illuminance and to the illuminance of its surrounding surfaces. This is the statement of the lighting design
- These steps lend themselves to computer implementation
  - Sum (2) for total surface area A
  - For each surface, reflected flux (5) = $A_s R_s E_s$
  - Sum (5) for total reflected flux, then:

    $$E_{rs} = \frac{A_s R_s E_s}{A}$$

  - For each surface $E_s(d) = E_s - E_{rs}(I)$. Complete (6)
  - Direct flux for each surface $F_s = E_s(d) \times A_s$. Complete (7)

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Surface Name	Area A_s	Reflectance R_s	Illuminance E_s	Reflected Flux $A_s R_s E_s$	Direct Illuminance $E_s(d)$	Direct Flux F_s

Table 6 – Flux distribution table [after Lynes(1987)]

Appendix B – HELIOS renderings

The next few pages contain the following computer generated images.

- Colour plate 1 – Example HELIOS rendering of simple room with objects
- Colour plate 2 – Front view, rendered as seen by human eye
- Colour plate 3 – Front view, false colour luminance rendering
- Colour plate 4 – Side view, rendered as seen by human eye
- Colour plate 5 – Side view, false colour luminance rendering
- Colour plate 6– interior view, rendered as seen by human eye
- Colour plate 7– Interior view, false colour luminance rendering

Appendix C – Radiosity scene information

This contains the parameters and settings needed in order to duplicate the sample views used in this document.

Radiosity Information

- Scene Information
    - Objects: 20
    - Surfaces: 103
    - Elements: 920
    - Vertices: 1452
    - Patches: 356
- Radiosity processing information
    - Number of Iterations: 250
    - Convergence: 0.0377
    - Time taken: 33 seconds on a Pentium 233 MHz processor

Scene Information

Front View (Figure 9)

- Camera Distances
    - View:3.5
    - Front Plane:7.25
    - Back Plane:10000
    - Tilt: 0.0
- Window Parameters
    - scale: 1.1
    - X size: 820
    - Y size: 530
- Eye Position
    - X: 2.5
    - Y: -10.0
    - Z: 1.6
- Focus Position:
    - X: 2.5
    - Y: 0.1
    - Z: 1.6

Interior View (Figure 10)

- Camera Distances
    - View:1.2
    - Front Plane:1.7
    - Back Plane:10000
    - tilt: 0.0
- Window Parameters
    - Scale: 1.9
    - X size: 900
    - Y size: 600
- Eye Position
    - X: 0.6
    - Y: 0.1
    - Z: 1.6
- Focus Position
    - X: 5.0
    - Y: 8.0
    - Z: 1.6

Side View (Figure 8)

- Camera Distances
    - View: 2.0
    - Front Plane:-1.99
    - Back Plane:10000
    - tilt: 0.0
- Window Parameters
    - Scale: 1.1
    - X size: 800
    - Y size: 380
- Eye Position
    - X: 14.0
    - Y: 4.0
    - Z: 1.6
- Focus Position
    - X: 2.5
    - Y: 4
    - Z: 1.6

Appendix D – Figures and Tables

Figures

Tables

References

[1] *CIBSE Code for interior lighting*, section XXX (London: Chartered Institute of Building Services Engineers) (1994)
[2] Waldram J, Studies in interior lighting *Trans. Illum. Eng. Soc.* (London) **19**(4) 95–134 (1954)
[3] Cuttle C, Cubic Illumination *Lighting Res. Technol.* **29**(1) 11–4 (1997)
[4] Hopkinson R, Stevens W and Waldram J, *Trans. Illum. Eng. Soc.* (London), **6**, 39 (1941)
[5] Mehmke R, Über die mathematische Bestimmung der Helligkeit, *Z. Math. Phys.* **43** 41 (1898)
[6] *CIBSE Code for interior lighting*, section 5.13 (London: Chartered Institute of Building Services Engineers) (1994)
[7] Cuttle C, On Sumpner's principle *Lighting Res. Technol.* **23**(2) 99–106 (1991)
[8] Lynes J, Burt W, Jackson G and Cuttle C, The flow of light into buildings *Trans. Illum. Eng. Soc.* (London) **31**(3) 65–91 (1966)
[9] Jay P, Lighting and visual perception *Lighting Res. Technol* **3**(1) 133–146 (1971)
[10] Jay P, Lighting and the Search for Meaning, Developments in Lighting – 1 (ed. J Lynes), London: Applied Science Publishing (1978)
[11] Lynes J, Lightness, colour and constancy in lighting design *Lighting Res. Technol*. **3**(1) 24–42 (1971)
[12] Cuttle C, The New Photometry, Developments in Lighting – 1 (ed. J Lynes), London: Applied Science Publishing (1978)
[13] Lowry E M, da Palma J, *J. opt. Soc. Am.* **51**, 740 (1961)
[14] Lynes J, Patterns of Light and Shade *Lighting in Australia* **7**(4) 16–20 (1987)
[15] Lynes J, Bedocs L, Professional studies in British architectural practice – Lighting. RIBA: London
[16] Glassner A, An introduction to ray tracing, London: Academic Press (1990)
[17] Ward G J, The Radiance Lighting Simulation and Rendering System *Computer Graphics Proceedings* (ACM SIGGRAPH '94) (1994)
[18] Ashdown I, Radiosity: A Programmer's Perspective. New York: Wiley & Sons (1994)
[19] Foley J D, van Dam A, Feiner S K, Hughes J F, Computer Graphics: Principles and Practice (2nd ed.). Reading MA: Addison Wesley
[20] Ward G, Radiance Digest **1**(2) *Available from* http://radsite.lbl.gov/ (Dec 1990)
[21] Ferwerda J A, Pattanaik S, Shirley P, Greenberg D, A Model of Visual Adaptation for Realistic Image Synthesis *Computer Graphics Proceedings* (ACM SIGGRAPH '96) 249-258 (1996)
[22] Zumtobel Lighting, Zumtobel Lighting Catalogue 97/98
[23] Ashdown I, Virtual Photometry – New directions in Photometry, Ledalite Architectural Products Inc. *Available from http://www.ledalite.com/* (1993)
[24] Ashdown I, Near-field Photometry – Measuring and Modelling Complex 3-D Light Sources, Ledalite Architectural Products Inc. *Available from* http://www.ledalite.com/ (1993)
[25] The Virtual Photometer: Modelling the flow of light. Ledalite Architectural Products Inc. *Available from* http://www.ledalite.com/ (May 1998)