RETHINKING DAYLIGHTING AND COMPLIANCE

John Mardaljevic

School of Civil and Building Engineering,
Loughborough University,
Loughborough, Leicestershire, LE11 3TU, UK
j.mardaljevic@lboro.ac.uk

Abstract: Daylight in buildings is the natural illumination experienced by the occupants of any man-made construction with openings to the outside. Our attempts to formulate some measure of daylight provision in buildings can be traced back over a century, and the daylight factor as we know it today is over fifty years old. Still the most common measure found in guidelines and recommendations world-wide, the daylight factor is used routinely and, it is fair to say, often rather uncritically. The consideration of daylight in buildings has received a new impetus from the accumulation of evidence on the wider benefits of daylight exposure. But it is continuing to prove difficult to advance beyond daylight factors towards a more realistic quantification of daylighting performance that would allow us to accommodate these new considerations in an evaluative schema. This paper examines the basis of current practice with respect to daylight evaluation, and suggests a few ways in which it can be improved with relatively modest additional effort. The paper also critiques some of the recents attempts to advance daylight evaluation by incremental means using so-called 'clear sky options'.

1. Background

Towards the end of 1990s the daylighting of buildings began to achieve greater attention than had previously been the case. There were a number of reasons for this, but the two most important 'drivers' were:

1. the widespread belief that the potential to save energy through effective daylighting was greatly under-exploited; and,

2. the emergence of data suggesting that daylight exposure has many positive productivity, health and well-being outcomes for building occupants.

The first originated with the widely-accepted need to reduce carbon emissions from buildings in order to minimise the anticipated magnitude of anthropogenic climate change. This in turn led to the formulation of guides and recommendations to encourage the design and construction of 'low energy' buildings and also for the retrofit of existing buildings. All these guides contain recommendations on daylighting, invariably founded on the daylight factor or an equally simplistic schema such as glazing factors [1].

The second driver was the gradual accumulation of data from disparate sources on the non-visual effects of daylight exposure. These effects are believed to be wide-ranging and include productivity and health/well-being, e.g. academic achievement, retail sales, recovery in hospitals, entrainment of circadian rhythm, etc. The mechanisms for these effects are not yet fully understood, and it is not yet known what the preferred exposure levels should be, nor if existing guidelines would be effective for these quantities [2]. Nonetheless, given the still relatively low cost of electric lighting – and the potential for it to be further reduced with solid-state lighting – there is evidence to suggest that the cost benefit from increased productivity due to good daylighting could be far greater than the cost saving from reduced energy expenditure [3]. Thus the second of these drivers has been promoted on both economic *and* environmental quality grounds.

Almost concurrent with the emergence of the two key drivers noted above were a major advance in the way daylight in buildings could be modelled; and, the development of numerous new glazing systems and materials to better exploit daylighting in buildings. These developments are expected to lead to significant changes in the way that daylight in buildings is both evaluated, i.e. using climate-based daylight modelling (CBDM), and exploited, e.g. by new glazing systems and materials. It should be noted that the performance of these new glazing systems often depends on their ability to shade and/or redirect sunlight. Thus they can only be reliably assessed using CBDM – the standard approaches (e.g. daylight factor) are unsuited to the task.

Notwithstanding that it is over a decade since CBDM was first demonstrated, and its effectiveness subsequently proven on a range of 'real world' projects, daylight criteria in most guidelines and recommendations are still founded on the daylight factor. More recently there have been attempts to advance the DF method incrementally using so-called 'clear sky' evaluations, though these appear unsatisfactory for reasons given below. There are a number of reasons why it has proven difficult to advance towards metrics founded on climate-based daylight modelling. Perhaps part of the difficulty in effecting this 'journey' is that we are not entirely certain regarding the point of departure: what exactly is the basis and rationale for the ubiquitous daylight factor method? This paper dissects both the basis of the method (i.e. relative values predicted using the CIE standard overcast sky) and how it is often applied to characterise a space, e.g. by giving an average daylight factor value.

The three sections that follow the note below are 'vignettes' of how I imagine the much needed deeper discussions on these matters might progress. The logic presented is mainly by way of argument illustrated with a handful of examples. The tests required to definitively confirm or disprove the various hypotheses framed here are beyond the scope of a relatively brief paper. However the author hopes that the reader will find the propositions sufficiently intriguing to engender further debate on these matters.

1.1 A note on the origin and formulation of the daylight factor

It appears that the daylight factor, or at least its precursor, was first proposed in 1895 by Alexander Pelham Trotter (1857-1947) [4]. The origins of the daylight factor are actually somewhat hazy since there does not appear to have been a seminal paper introducing the approach. The reference to its introduction in 1895 appears to be anecdotal and recalled a number of years later. The daylight factor was conceived as a means of rating daylighting performance independently of the actually occurring, instantaneous sky conditions. Hence it was defined as the ratio of the internal horizontal illuminance E_{in} to the unobstructed (external) horizontal illuminance E_{out}, usually expressed as a percentage, Figure 1. However, the external conditions still need to be defined since the luminance distribution of the sky will influence the value of the ratio. At the time that the daylight factor was first proposed it was assumed that heavily overcast skies exhibited only moderate variation in brightness across the sky dome, and so they could be considered to be of constant (i.e. uniform) luminance. Measurements revealed however that a densely overcast sky exhibits a relative gradation from darker horizon to brighter zenith; this was recorded in 1901. With improved, more sensitive measuring apparatus, it was shown that the zenith luminance is often three times greater than the horizon luminance for some of the most heavily overcast skies [5]. A new formulation for the luminance pattern of overcast skies was presented by Moon and Spencer in 1942, and it was adopted as a standard by the CIE in 1955. Normalised to the zenith luminance L_z, the luminance distribution of the CIE standard overcast sky has the form:

$$L_\theta = \frac{L_z \left(1 + 2\sin\theta\right)}{3} \tag{1}$$

where L_θ is the luminance at an angle θ from the horizon and L_z is the zenith luminance (Figure 1).
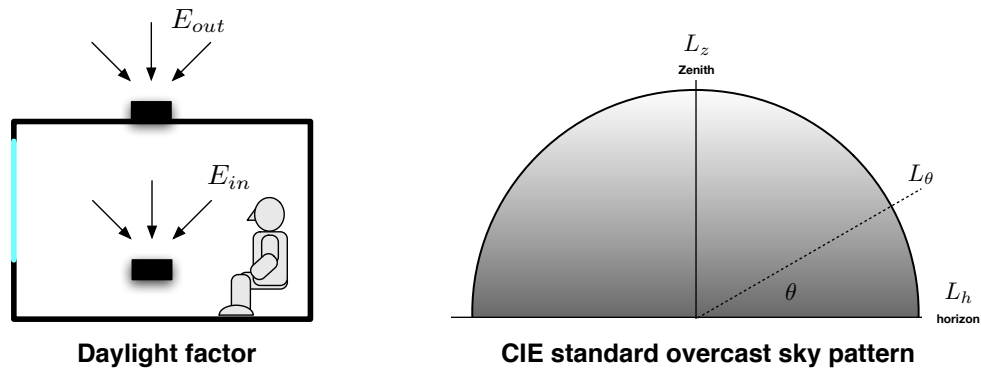
Figure 1: Definition of the daylight factor and the CIE standard overcast sky

2. Being mean to the average

It is proposed here that the average should no longer be used as a means to summarise measures such as the daylight factor distribution. The average is typically used to give a 'bottom line' number which is intended to be the sole daylight performance indicator for the space. Instead, the median (or a quartile) should be employed whenever a single quantity is required to characterise a space. The average tells us nothing about the distribution of DF in the space, whereas the median does. The average can be a quite misleading quantity when applied to daylight distributions, especially for spaces illuminated from vertical glazing on one wall where the very high DFs close to the windows can significantly influence the average DF value. Because of this, the average is very sensitive to the proximity of the sensor plane to the glazing. The closer the sensor points are to the glazing, the higher the average for the daylight factor distribution in the space. As far as this author is aware, it was not until the appearance of the 2011 revision of Lighting Guide 5 (LG5) that a recommendation for a perimeter zone between sensor points and glazing/walls has been given in a UK guideline for simulation. LG5 recommends a 0.5m gap (perimeter zone) – which seems reasonable, though it should be noted that the rationale given in LG5 (i.e. to avoid the low values at the back of the space) is incorrect.

The upper section of the plot shown in Figure 2 gives the DF distribution across (half) of a 6m wide by 9m deep side-lit space (2.7m floor to ceiling height). Here the sensor plane covers the entire 6×9m internal plan, though the glazing is located on the outer side of a 0.2m reveal (so it is not quite a worst case regarding close proximity of the sensor plane to the glazing). The average DF for this scenario was 2.8%, however the median value was only 1.1%. With the latter we know that only half the area of the space has achieved a DF of 1.1%, whereas with the average we have no such certainty. And, more worryingly, the average can in some people's minds be conflated with the median – giving a completely false impression of the DF distribution for the space. Having a reasonable perimeter (e.g. 0.5m) reduces the size of the false impression given by the average, but it does not eliminate it. In contrast, the median value is largely insensitive to the size of the perimeter, and so it is not only a more informative measure, it is also more robust since it is largely unaffected by any 'game playing' with respect to the placement and size of the perimeter.

3. A gloomy view of the CIE overcast sky

At first glance, the CIE overcast sky seems a reasonable basis for the evaluation of daylight in buildings. This 'feeling' is perhaps formed, or at least bolstered, from seeing phrases such as these commonly found in documents pertaining to daylight evaluation: "the overcast sky represents worst case conditions"; "the daylight factor is defined as the worst case"; and, "the daylight factor was invented in northern Europe where the fully overcast sky is common". The implication being that, if we provide a certain measure of daylight for the "worst case", then surely it can only be better than that for the rest of the time. However, whilst such notions are suggestive, the rationale for the daylight factor has rarely, if ever, been rigorously expounded. For example, what exactly is meant by "worst case"? Is it that the absolute values provided by the sky (i.e. the diffuse horizontal illuminance) is (are?) "worst case", or is it perhaps that the luminance distribution on the sky vault is a "worst case"? Or maybe a combination of the two? And, if the daylight factor is suitable for "northern Europe", what is the extent of its zone of

applicability? The daylight evaluation in the first edition of the Estidama Pearls Design System for Abu Dhabi was founded on daylight factors, i.e. the CIE standard overcast sky [6]. A quick examination of the standard climate for Abu Dhabi reveals that it is almost *never* overcast in that region of the United Arab Emirates. This, not unexpected observation, suggests that at least in some instances the daylight factor has indeed been applied well outside of its 'zone of applicability', notwithstanding the uncertainty regarding its precise boundaries.
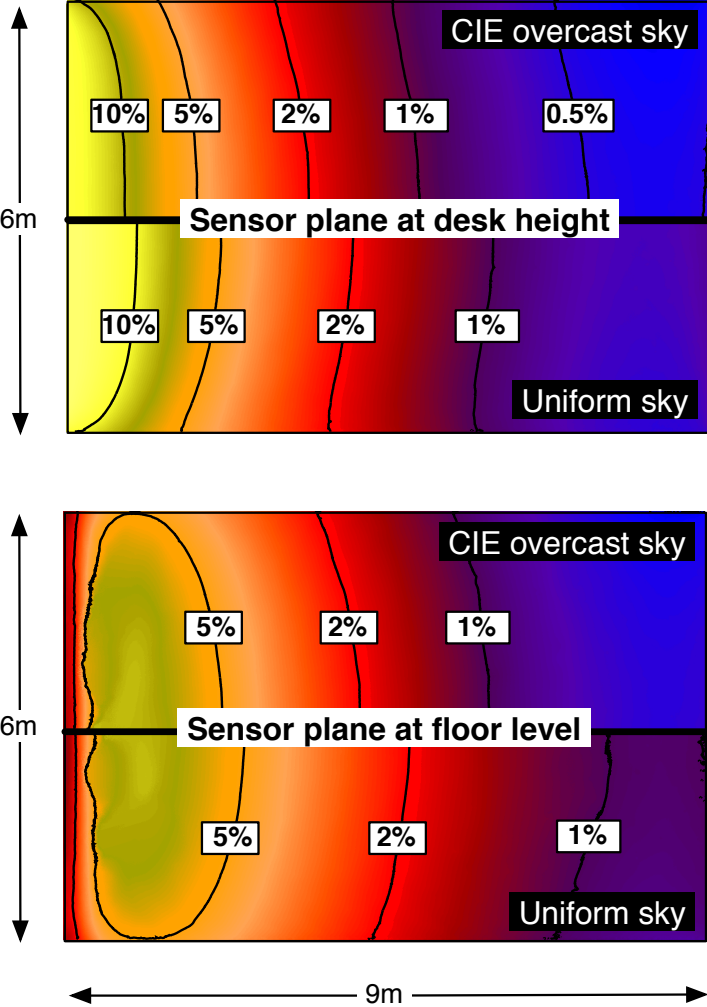

Figure 2: Daylight factor distribution for CIE standard overcast and uniform skies

The link between (relative) daylight factors (i.e. percentage values) and absolute levels of illuminance (i.e. lux) has always been tenuous. DFs are of course derived from absolute values, but the latter are often ignored thereafter. Design guides often give recommendations in terms of daylight factors, but then also suggest that daylight should provide illuminances of, say, 300 lux or more for much of the year. Building Bulletin 90 (Lighting Design for Schools) does describe how to relate DFs to estimates of achieved absolute illuminance [7], but these 'conversions' are rarely carried out.

In Australia and New Zealand a uniform sky is used for what are in effect 'daylight factor' calculations, though the sensor plane is set at floor level rather than at desk height introducing another dissimilarity when comparing methods. The differences in predicted distributions between 'classic' daylight factor (i.e. overcast sky and sensor plane at desk height) and the option recommended in Australia/NZ (i.e. uniform sky with sensor plane on the floor) is revealed by comparing the plots shown in Figure 2. The metrics derived from each of the four distributions are given in Table 1.

As we might expect for a side-lit space, a uniform sky produces higher ratios (i.e. DFs) than a CIE overcast because, with the latter, the sky vault luminance is 'concentrated' around the zenith – the average, median and minimum DFs are markedly higher for the uniform sky, Table 1. Placing the sensor plane at floor level results in a vastly different DF pattern compared to when at desk height – irrespective

Sky type (sensor height)	Average [%]	Median [%]	Max [%]	Min [%]
CIE overcast (desk)	2.8	1.1	15.3	0.38
Uniform (desk)	3.4	1.6	15.9	0.60
CIE overcast (floor)	2.5	1.6	8.1	0.47
Uniform (floor)	2.9	2.0	8.3	0.71

Table 1: Metrics derived from the distributions shown in Figure 2

of the sky type, Figure 2. Because of the sill, the window from the perspective of a sensor plane on the floor appears more like a clerestory window, i.e. the peak in both DF distributions is displaced *away* from window rather than closest to it. Note that, if the glazing in the space was floor-to-ceiling, then the DF distributions at floor level would appear similar to those in the upper plot, but with higher absolute values since the sensors now 'see' a greater expanse of sky. Furthermore, the values would be highly misleading because the sensors now include the contribution of light that enters the space *below desk height*, i.e. heading directly for the (typically) low-reflectance floor where most of the light will be immediately absorbed. I have not been able to locate any documents that describe the rationale for placing the sensor plane at floor level.

Whilst it is hard to see any benefit in having the sensor plane at floor level, might there nevertheless be a case for basing estimates of internal illuminance availability on the uniform rather than the standard overcast sky? Consider the cumulative diffuse availability curve shown in Figure 3. The following estimates can be derived from the curve: a 2% DF gives ~100 lux for 85% of the year, whereas the same DF gives an illuminance of ~300 lux for about 55% of the year. Occupants will, of course, invariably prefer daylight illuminances around the 300 lux mark compared to those around 100 lux. Some of the skies around the 5,000 lux (diffuse horizontal illuminance) mark are likely to conform to varying degrees to the standard overcast pattern. However, the 15,000 lux skies needed to produce 300 lux internally (for a DF of 2%) are much more likely to have luminance distributions that diverge significantly from the standard overcast pattern. And that will be even more the case for the remaining 45% of the skies in the distribution that have higher diffuse illuminances. In other words, when the DFs are predicted using the standard overcast sky, the basis for the estimate of the occurrence of internal illuminances is strictly self-consistent only for those skies in the annual climate file that conform to the CIE standard overcast pattern. But what is the proportion of annual occurring skies that are a good match for the CIE standard overcast sky? That is not an easy question to answer. It is possible to determine the annual occurrence of *essentially* overcast skies in standard climate datasets using, say, the Perez clearness index. However, those will include the whole gamut of overcast skies, many of which it seems do not exactly match the standard pattern: Enarun and Littlefair suggest that "...if a general cloudy sky is all that is required, the CIE may not be the best option" [8]. In the same paper they suggest that the "quasi-overcast sky" may serve better as a "general cloudy sky". The quasi-overcast sky has a more gradual gradation between horizon and zenith compared to the CIE standard overcast. But, it also includes a small component which varies with angle from the sun. Thus, it could <u>not</u> replace the use of the CIE standard overcast in a daylight factor evaluation because the sun position is now a factor in the evaluation.

To recap. The CIE standard overcast sky is in fact – to quote Enarun and Littlefair – an "extreme" case of overcast sky. Thus, skies that conform to the CIE standard overcast sky pattern are likely to be rarer than is generally imagined, and in any case produce internal illuminances at or below the lower end of what is generally preferred by occupants. A sky luminance distribution with a smaller ratio between horizon and zenith is believed to be a better fit to the more typical gamut of overcast skies, i.e. the brighter overcast skies that deliver more useful levels of natural light for occupants. Given that the "quasi-overcast" cannot replace the CIE standard overcast in a DF-based evaluation, perhaps the uniform sky is in fact the 'best' simple sky condition on which to base estimates of daylight provision using diffuse illuminance curves. In fact, the uniform sky is probably a closer fit to an average of the "quasi-overcast" (for varying sun positions) than the CIE standard overcast pattern. Furthermore, it is perhaps not unreasonable to describe the CIE standard overcast sky pattern an one that exhibits *bias* when used to estimate the occurrence of internal illuminance from DFs. This is because the luminance pattern – maximum at the zenith – deviates from the gamut of typically occurring overcast patterns in a consistent manner. The effect on the prediction of ratios at the sensor plane (i.e. DFs) is evident in Figure 2 and Table 1. Compared to the uniform sky, the maximum DFs for the standard overcast sky are more tightly packed closer to the window where the sensors have the best 'view' of high altitude sky close to the zenith.
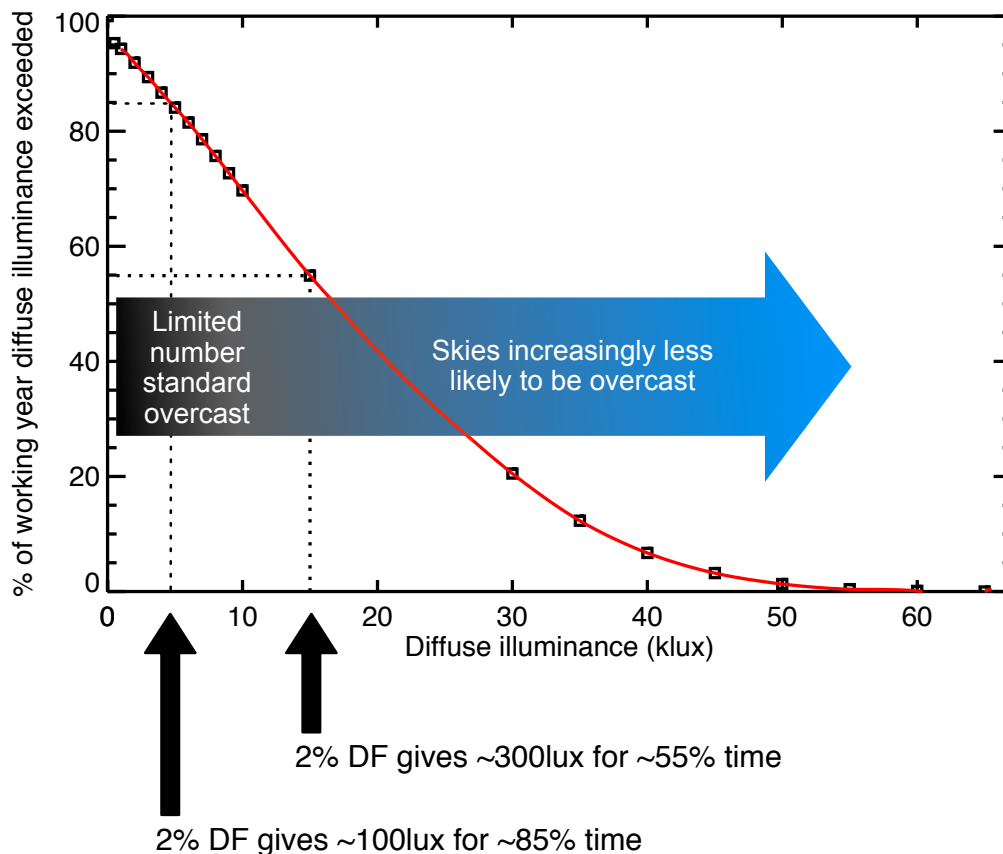
Figure 3: Cumulative curve showing diffuse illuminance availability

The case for suggesting that a uniform sky might actually be a sounder basis for daylight design than the CIE standard overcast can be *reasoned*, as demonstrated above. But it is not at all clear at this stage how it might be *tested*. In large part this is because we do not yet have a robust notion regarding an agreed upon datum against which we can discriminate outcomes. A somewhat idealised datum is of course some measure that, if achieved, ensures "good daylighting". One proposed measure made by the IES Daylight Metrics Committee is the annual occurrence of 300 lux across the workplane.

4. New approaches: Do 'halfway' methods work?

The "clear sky option" appears to have been introduced in LEED version 2.2 as a means to overcome the limitations of the climate/orientation insensitive glazing factor and daylight factor methods. To achieve credit 8.1, the requirement can be:

> *Demonstrate, through computer simulation, that a minimum daylight illumination level of 25 footcandles has been achieved in a minimum of 75% of all regularly occupied areas. Modeling must demonstrate 25 horizontal footcandles under clear sky conditions, at noon, on the equinox, at 30 inches above the floor.*

Whilst this may appear, at first, reasonable, the LEED v2.2 documentation gives no supplementary data for the evaluation. This omission all but renders the evaluation meaningless since there is no statement regarding the diffuse horizontal illuminance that the sky should be normalised against. The user, it seems, is to trust the default value that is provided by the sky generator program. The default value is an extremely coarse approximation with some latitude dependance (and of course time of day/year), but no basis whatsoever in local, prevailing climatic conditions. Many users are unaware that the key input parameter for their simulation is of dubious provenance and has been automatically selected on their behalf. It gets worse. Nor indeed is there any mention of what the sun luminance (usually derived from direct normal illuminance) should be. This too is surprising, since the sun contribution will greatly

add to the illuminances resulting from the diffuse sky (which will depend on the unspecified diffuse horizontal illuminance anyway). Given the relatively modest target illuminance (around 250 lux) it seems likely that the evaluation is meant to be carried out using a clear sky distribution <u>without</u> a sun. Which, of course, is a physical impossibility in reality. Anecdotal evidence has confirmed that users of LEED have indeed 'demonstrated compliance' with the recommendations and obtained Daylight Credit 8.1 by using a physically impossible luminous environment (i.e. clear sky without sun) that is normalised to an unknown diffuse horizontal illuminance.

ASHRAE standard 189.1 (2009) has a similar clear sky option to LEED. As with LEED, there is no mention of normalising the sky to a specified diffuse horizontal illuminance, so the same shortcomings (outlined above) apply. As with the LEED Clear Sky option, the ASHRAE draft guidelines suggest (but do not clearly state) that the clear sky modelling is to be done without a sun – which is, as noted above, a physically impossible illumination condition in nature. The ASHRAE draft states that *either* the CIE Overcast *or* the CIE Clear sky model may be used. This offers intriguing possibilities to the artful compliance chaser, since the outcome it turns out depends to a large degree on what default values 'drop out' of the sky generator program. Since many practitioners use the *Radiance* lighting system, either in its raw (UNIX) form or in one of the many bundled packages, its instructive to see how different the outcomes can be depending on the choice of sky.

The *Radiance* (UNIX) command 'gensky 3 20 12 -c' generates a description of the brightness distribution of the CIE standard overcast sky for noon, 20 March (i.e. month 3). A similar command generates the description for the CIE clear sky pattern. The guidance gives no recommendation regarding normalisation of the skies to a known diffuse horizontal illuminance (E_{dh}). So, the diffuse horizontal illuminance of the resulting sky depends entirely on how the sky model generator program gensky was devised to produce skies of either type, i.e. its default behaviour. The diffuse horizontal illuminance for the two sky types turns out to be very different – almost by a factor of two, Table 2. It may even seem counterintuitive that, without any user intervention by way of supplying normalisation values, the diffuse horizontal illuminance for the overcast sky should be nearly *twice* that of the clear sky. However, the reason is quite straightforward. The sky model generator program does not have any knowledge of local meteorological conditions. What it does know are: latitude/longitude i.e. location (the default of Berkeley, USA is used in the example), time of day/year and therefore sun position, and also the incident extraterrestrial solar radiation. This last part is apportioned between sky and sun (if present) depending on the selected sky type. For an overcast sky the incident extraterrestrial solar radiation is reprocessed into diffuse sky radiation (using default values for turbidity etc.). But, for a clear sky distribution, the extraterrestrial solar radiation is apportioned between the sky radiation *and* the (now expected but not included) sun. Thus, the diffuse horizontal illuminance for the clear sky is lower (typically just over half using the gensky defaults) than the diffuse horizontal illuminance for the overcast sky. Notwithstanding the differences in the sky luminance patterns, the designer 'chasing' the attainment of an absolute level of interior illuminance would be advised to opt for the overcast sky because of its much higher diffuse horizontal illuminance.

CIE sky type	*Radiance* command	E_{dh} [lux]
Standard overcast	gensky 3 20 12 -c	14,679
Standard clear	gensky 3 20 12 -s	8,454

Table 2: Diffuse horizontal illuminance (E_{dh}) for standard overcast and standard clear skies generated without normalisation

In recognition of what must be viewed as a less than ideal state of affairs regarding the lack of normalisation in the 'clear sky option' of Version 2.2, the 2nd Public Comment Draft on LEED (July 2011) contains the following:

> *Demonstrate through computer simulations that the applicable spaces achieve illuminance levels between 300 lux and 3000 lux for both of the following sky conditions:*
> *− 9:00 am equinox on a clear-sky day (solar time)*
> *− 3:00 pm equinox on a clear-sky day*
> *Illuminance intensity for sun (direct component) and sky (diffuse component) for clear sky and overcast conditions for those time periods shall be derived from the local weather data, or TMY weather tapes for the nearest city, first by selecting the two days within 15 days of September 21st and March 21st that represent the clearest sky and most overcast sky condition, and then averaging the hourly value for the appropriate spring and fall hour.*

Whilst this revision might, at first, be seen as 'heading in the right direction', it too has potential problems and confounding issues. The patterns of hourly values in the illuminance datasets are unique and, because of the random nature of weather, they will never be repeated in precisely that way, Figure 4. Climate datasets are however representative of the prevailing conditions measured at the site, and they do exhibit much of the full range in variation that typically occurs, i.e. they provide definitive yardstick quantities for modelling purposes – provided that the <u>entire</u> year is used in the evaluation. The solid lines on the plots in Figure 4 mark the times of the equinoxes, and the dashed lines mark the date 15 days either side of each equinox. As is evident from the pattern, whilst it might be *likely* that a sunny (i.e. clear sky) period occurs within 15 days either side of the equinox, it is by no means certain because of the random nature of weather. Also, how "clear" is clear? That is not specified. Thus, it is highly problematic to attempt to extract and define supposed 'representative' illuminance data from climate files. Furthermore, "averaging" of any climatic illuminance data is risky since the user must ensure that the conditions to be averaged are indeed similar.
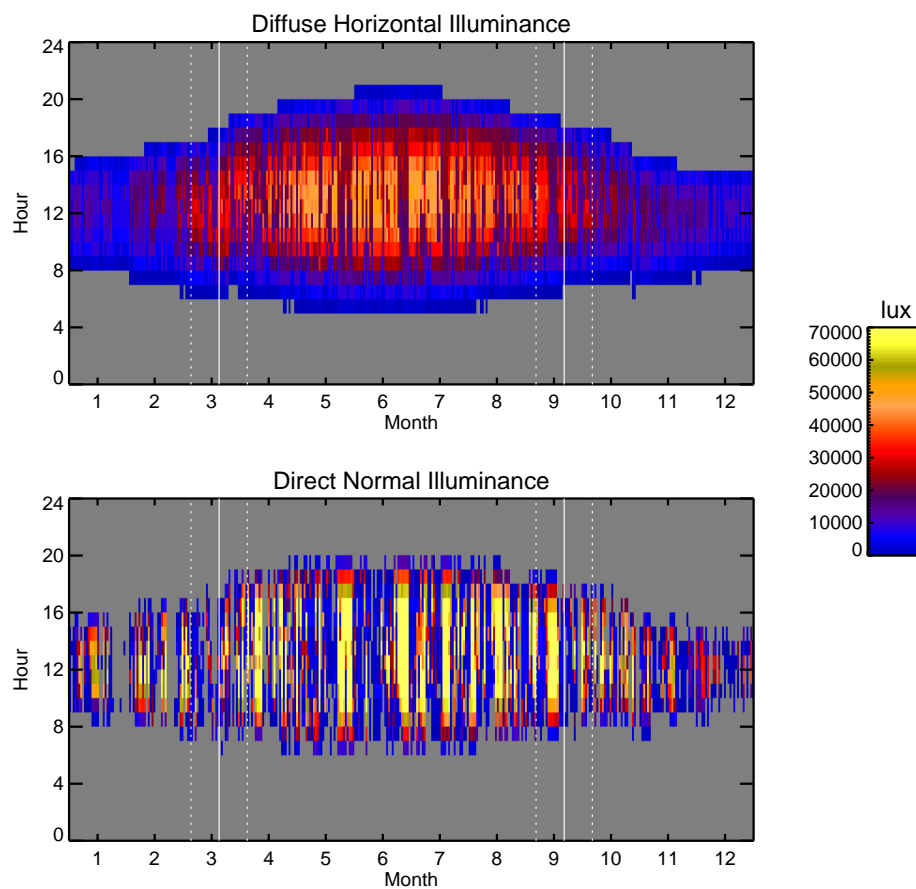

Figure 4: Illuminance data from the standardised climate file for London

Based on the attempts made thus far, it does not seem possible to advance the DF approach by incremental means, i.e. evaluations based on 'clear sky options', 'snap-shots' or 'salami-slicing' of climate data. Efforts in this direction have resulted in methods that are one or more of the following: confusing, inconsistent, prone to the vagaries of patterns in climate data, and/or without a proven rationale. There seems to be no half-way house between a DF-based evaluation (e.g. in conjunction with cumulative diffuse illuminance curves) and a full-blown annual evaluation using climate-based daylight modelling.

5. Discussion

Notwithstanding the more than occasional tone of a *jeremiad*, this article is in fact intended to accentuate the positive – we do have the means now with climate-based modelling to greatly advance the basis of daylight evaluation for buildings. However, CBDM tools are still largely the preserve of lighting simulation experts/researchers. For CBDM to become mainstream the software to do it needs to be taken up and

supported by one or more major software houses. Here lies a classic 'chicken and egg' conundrum. On one hand, those who draft guidelines are reluctant to recommend metrics founded on CBDM because tools to predict the metrics are generally not available, at least not as supported software. On the other hand, the software vendors are understandably loathe to dedicate the resources to develop and maintain CBDM tools because – inasmuch as climate-based metrics are not in the guidelines – there will be no real market for these new tools. This presents something of an impasse to all those who strive to advance daylighting standards beyond the current guidelines.

A suggested way around this impasse follows. In order to obtain 'buy-in' from all stakeholders (e.g. standards bodies, designers, end-users, tool developers, etc.) it is important that first they recognise the benefit of the changes proposed. These benefits should include: a more robust approach to evaluating daylight in buildings using existing tools with only modest enhancements; a methodology that allows for later progression to more reality-based evaluations; and, a transition roadmap with clear market horizons to ensure that software vendors invest the necessary resources to develop the next generation of modelling tools (i.e. CBDM for 'end-users'). To this end, it is proposed that current standards based on daylight factors should be upgraded as soon as possible to evaluations founded on the annual occurrence of an absolute value for illuminance (e.g. 300 lux) estimated from the cumulative availability of diffuse illuminance as determined from standardised climate files. For example, a daylighting 'target' could be that half of the sensor points in a side-lit space should achieve 300 lux for half of the time when the sun is above the horizon. This is an application of an established but largely neglected approach [9]. Such an upgrade requires only a modest extension to existing DF software, and, importantly, it provides some 'connectivity' between the daylight availability and the prevailing climate. Note also that there may be good cause to use a uniform rather than a standard overcast sky for this evaluation. Of course, this is not a full-blown climate-based solution since direct/indirect sun is not accounted for. However, unlike the 'halfway' measures described in this article, the cumulative illuminance approach has a defensible rationale. Furthermore, by shifting the analysis to measures of absolute illuminance, it prepares the ground for a relatively smooth transition to eventual, full-blown CBDM evaluations. One could envisage, say, a three year 'overlap' period in standards/guidelines during which either the cumulative illuminance approach or CBDM could be used to demonstrate compliance. And then, at the end of that period, only evaluations founded on CBDM would be permitted. Such a provision would encourage software houses to invest the time and resources to develop end-user CBDM tools in the certainty of a guaranteed market for the product by a due date – thus solving the 'chicken and egg' conundrum noted above. Note that, although similar, or even identical, targets would be used with either approach, with CBDM it would be necessary to model user deployment of blinds etc. since direct (and indirect) sunlight now figures in the evaluation.

The author hopes that the issues raised here will be progressed in wider discussions within the daylighting community and relevant stakeholders. This article is the first of several that are in preparation in support of the activities of CEN TC 169/WG11. Though it should be noted that the views expressed in this paper are those of the author alone.

Acknowledgements

References

[1] J. Mardaljevic, L. Heschong, and E. Lee. Daylight metrics and energy savings. *Lighting Research and Technology*, 41(3):261–283, 2009.

[2] M. Andersen, J. Mardaljevic, and S. W. Lockley. A framework for predicting the non-visual effects of daylight – Part I: photobiology-based model. *Lighting Research and Technology*, 44(1):37–53, 03 2012.

[3] L. Heschong. Daylighting and human performance. *ASHRAE Journal*, 44(6):65–67, 2002.

[4] J. A. Love. The evolution of performance indicators for the evaluation of daylighting systems. *Industry Applications Society Annual Meeting, 1992., Conference Record of the 1992 IEEE*, pages 1830–1836 vol.2, 1992.

[5] P. Moon and D. E. Spencer. Illuminations from a non-uniform sky. *Illum. Eng.*, 37:707–726, 1942.

[6] Estidama. Pearls Design System - New Buildings Method. *Abu Dhabi Urban Planning Council*, 2009.

[7] BB90. Lighting design for schools. *Building Bulletin 90 (2008 revision), Department of Education and Skills, School Building and Design*, London 2008.

[8] D. Enarun and P. Littlefair. Luminance models for overcast skies: Assessment using measured data. *Lighting Research and Technology*, 27(1):53–58, 1995.

[9] D. R. G. Hunt. Improved daylight data for predicting energy savings from photoelectric controls. *Lighting Research and Technology*, 11(1):9–23, 1979.