

SHADERADE: COMBINING RHINOCEROS AND ENERGYPLUS FOR THE DESIGN OF STATIC EXTERIOR SHADING DEVICES

Jon A. Sargent, Jeffrey Niemasz, and Christoph F. Reinhart
Harvard University Graduate School of Design, Cambridge, MA

ABSTRACT

This paper examines several existing methods of static shading device design, and presents a new approach called SHADERADE. The approach is implemented as an eponymous tool based on Rhinoceros® and EnergyPlus, and offers flexible, novel techniques for assessing the thermal desirability of solar transmittance through any potential shading volume or surface. Using simulated sidelit offices located in Anchorage, Boston and Phoenix, it is shown that SHADERADE is able to consistently generate shading systems with improved thermal performance vis-à-vis existing methods. It is also shown that the approach can handle curved geometries with ease, and can efficiently manage the sizing of shading devices by identifying regions that matter most. The authors hope that these capacities will facilitate a more effective, less prescriptive approach to shading design.

INTRODUCTION

Static exterior shading systems such as louvers and overhangs have long been used as architectural means to control solar radiation in order to improve indoor environmental conditions. During the 1950s Le Corbusier popularized *brise soleils* while Olgyay and Olgyay provided a theoretical framework for sizing them (Olgyay et al. 1957). In recent decades, more advanced methods for generating optimal shading devices have been proposed by Arumi-Noé (1996), Kabre (1999), Kaftan (2001) and Marsh (2003). This paper presents a new form-generating method for shading systems called SHADERADE and compares shading devices generated using SHADERADE against devices generated via existing methods for a South-facing office located in Boston, Phoenix and Anchorage.

EXISTING METHODS

In general, existing methods of shading optimization divide the task of static shading design into two discrete steps: *shading period selection* and *form finding*. The first step involves identifying the periods of the year during which it is desirable or undesirable to have solar radiation directly incident on a window. The second step is concerned with

finding the form of a shading device that fulfils the requirements identified in step one.

Shading Period Selection

With notable exceptions such as Marsh's point-cloud ray-trace (2003) and Kaftan's cellular approach (2001), the vast majority of period-selection methods involve the picking of 'cut-off' dates. Most often, this involves selecting a set of contiguous dates that approximate the annual period of either mechanical cooling or potential overheating. It is also common to select, within these dates, a standard set of 'critical' hours during which direct solar radiation should be kept off a window. A shading period selected in this fashion therefore has the form: Start Date to End Date, from Start Hour to End Hour (e.g. May 2 – August 10, from 10 am – 2 pm). Note that, because the path of the sun is roughly identical on days equidistant from the summer solstice, static devices must have shading periods that are symmetric around June 21.

Equinox selection: Three methods of cut-off date selection are presented below. The first—*equinox selection*—defines the desired shading interval as that between the vernal and autumnal equinoxes (from March 20 to September 22 in the Northern Hemisphere). This 'rule-of-thumb' type method is based on the assumption that the equinoxes roughly approximate the boundaries between annual heating and cooling periods.

Degree-day selection: The second date-selection method, *degree-day selection*, uses heating and cooling degree days to identify 'overheated' and times of year, when shading is assumed to be desirable. One's initial instinct is to define this period as the annual interval over which cooling degree days exceed heating degree days. However, because weather patterns are often asymmetric around the summer solstice, this interval tends to conflict with the requirement that static devices have solstitially symmetric shading periods. In Figure 1, this conflict is illustrated for Phoenix. Using base temperatures of 18 and 20 °C, and TMY3 weather data from the Department of Energy (http://apps1.eere.energy.gov/buildings/energyplus/weatherdata_about.cfm), monthly heating and cooling degree days (*HDD* and *CDD*) are plotted over a typical year. The 'switch-over' dates between heating and cooling periods, as

defined by the intersections of the graphs, occur on March 23 and November 6 – 90 days before and 138 days after the summer solstice, respectively. For asymmetric start and end dates like these, the shading period must be defined by one date, the other, or a compromise between the two.

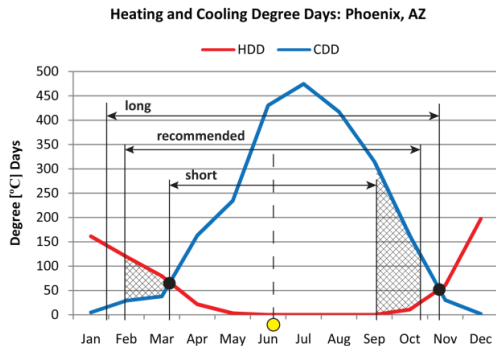


Figure 1 Degree-day date selection for Phoenix, AZ

In this study, compromise periods were established by picking the solstitially symmetric interval with the maximum time-integrated sum of $(CDD - HDD)$. For Phoenix, this approach yielded a ‘recommended’ shading period from February 12 to October 28. For Anchorage, monthly heating degree days exceeded monthly cooling degree days throughout the year, so a shading period was not recommended.

The degree-day selection method is conceptually superior to the equinox method because it begins to account for local climate. However, heating and cooling degree days do not always provide good estimates of when a space is likely to be ‘over-’ or ‘under-heated.’ These conditions are affected by building-specific parameters – such as internal loads, occupancy schedules, building form, envelope properties, and thermal mass effects – that heating and cooling degrees do not take into account.

Thermal Selection: A third date-selection method, *thermal-selection*, therefore involves conducting a “quick” annual thermal simulation of the space under consideration without any static shading device. This approach accounts not only for climate, but also for the thermally-consequential building parameters mentioned above. While properly setting up a thermal model of a space may require considerable time and knowledge on the building modeller’s part, there are a number of easy-to-use tools available to model spaces with “standard” programs, geometries, and materials within minutes. Example applications are the *EnergyPlus Example File Generator* (http://apps1.eere.energy.gov/buildings/energyplus/file_generator_about.cfm) and *DesignBuilder* (DesignBuilder Software Ltd.). In this study, hourly sensible heating and cooling loads for an example sidelit office were simulated with EnergyPlus and binned by month in order to assess heating and cooling periods. A description of the simulated space is provided in the Comparison section. Compromise

intervals for asymmetric periods were resolved using the same method described for degree-day selection. For Phoenix, cooling loads were always greater than heating loads, so the shading period was defined as the entire year.

As a side note, heating and cooling loads can also be weighted with conversion factors to assess other measures of interest, such as energy use or carbon emissions. Using different weighting functions will obviously affect the resulting time period.

Hour Selection: Each of the three date-selection methods leads to a single cut-off date that divides the year into a period of shading and a period of no shading. During the shading period, it is traditionally assumed that direct sunlight should be blocked from entering any portion a window under consideration. While the initial reaction might be to require this condition during all daylight hours in the period, such an approach will often lead to extremely large shading systems because of low solar altitudes during the beginning and end of the day. It is hence customary to define an hourly time interval during which complete shading is ‘guaranteed’ throughout the shading period. The three time intervals used in this study were 9am-3pm, 10am-2pm, and *at noon*.

Form Finding

Once a shading period has been identified using one of the methods described above, various formal strategies can be employed to keep direct sunlight off the glazed portion of the façade during the designated time intervals. Two of these methods are described below.

2-d method: The simplest strategy for guaranteeing shading is the *2-d* method. In this approach, shading design is based on a two-dimensional, vertical, perpendicular section through a glazed wall under consideration, as shown in Figure 2.




Figure 2 2d method for a simple overhang

The depth d of the shading device is taken as a function of the maximum value of θ over the shading period. θ is the angle, off vertical, of the solar vector projected onto the section plane.

For each climate, the simplest geometry resulting from the 2d method is a horizontal overhang of depth d placed at the top of the window opening. The difficulty of dealing with the near and far ends of the

window can be circumvented by assuming that the shading device extends several meters past both edges. However, one should note that this assumption effectively limits the method to a broad flat façade with continuous horizontal overhangs.

Table 1 shows the overhang depths resulting from the 2d method along with the start and end dates using the three period selection methods for Anchorage, Boston and Phoenix. The thermal-selection analysis was based on the south-facing office described in the Comparison of Methods section.

Table 1

Start and end dates returned by the different period selection methods along with the overhang depth according to the 2d method for Anchorage, Boston and Phoenix

Anchorage, AK	Start	End	<i>d</i> (m)
Equinox 9-3	3/20	9/22	3.38
Equinox 10-2	3/20	9/22	3.38
Equinox noon	3/20	9/22	3.38
Degree Day 9-3	none	none	none
Degree Day 10-2	none	none	none
Degree Day noon	none	none	none
Thermal 9-3	3/28	9/14	2.97
Thermal 10-2	3/28	9/14	2.97
Thermal noon	3/28	9/14	2.97
Boston, MA	Start	End	<i>d</i> (m)
Equinox 9-3	3/20	9/22	1.71
Equinox 10-2	3/20	9/22	1.70
Equinox noon	3/20	9/22	1.70
Degree Day 9-3	5/15	7/28	0.81
Degree Day 10-2	5/15	7/28	0.81
Degree Day noon	5/15	7/28	0.81
Thermal 9-3	2/13	10/27	3.34
Thermal 10-2	2/13	10/27	2.94
Thermal noon	2/13	10/27	2.72
Phoenix, AZ	Start	End	<i>d</i> (m)
Equinox 9-3	3/20	9/22	1.23
Equinox 10-2	3/20	9/22	1.23
Equinox noon	3/20	9/22	1.23
Degree Day 9-3	2/12	10/28	2.82
Degree Day 10-2	2/12	10/28	2.31
Degree Day noon	2/12	10/28	2.01
Thermal 9-3	12/22	12/21	4.69
Thermal 10-2	12/22	12/21	3.48
Thermal noon	12/22	12/21	2.87

Ecotect: A second form-finding method, which has the advantage of considering solar vectors in three dimensions, is available through the Autodesk® *Ecotect™ Shading Wizard*. To use the Wizard, a window-geometry must be specified in Autodesk *Ecotect* 2010 (Autodesk® 2010). After feeding the window geometry into the tool, the user must specify a climate location and the cut-off dates and hours of the desired shading period. In this study, each period-selection result (listed in Table 1) was entered and run separately. The Wizard offered solutions for several different types of shade, including a

rectangular overhang, an ‘optimized’ horizontal overhang (by ‘critical date’ or ‘critical period’), and a ‘surround shade,’ which combined horizontal and vertical projections from the top and sides of the window frame. In this study, both the period-optimized overhang and the surround shade were selected for each climate and shading period. According to Marsh, the *Ecotect* shades are derived to guarantee direct shading over the entire window during the specified time periods (Marsh 2003).

While both *Ecotect* and the 2d method can provide full shading over a specified set of time intervals, the *Ecotect* approach is preferable because it can do so while theoretically blocking fewer solar vectors *outside* these times. Moreover, as a 3-dimensional tool, the *Ecotect* method is applicable to a wider set of window and shading geometries.

Limitations of Existing Methods

The methods described so far involve two steps: first, the selection of periods of shading; and second, the attempt to find forms that guarantee complete shading of a window throughout these periods.

As others have pointed out, one significant limitation of these approaches is the problem of overshading (Kaftan and Marsh 2005). By guaranteeing complete shading at the boundary of the desired period, one almost inevitably guarantees partial shading outside that boundary – i.e., during time segments when shading may be *undesirable*. One could compensate for this ‘spill-over’ effect by changing the boundaries of the shading period, but the methods offer no guidance as to the necessary size of such an adjustment. In this sense, the problem is not overshading *per se*, but rather the methods’ reliance on a single boundary condition in the first place. In reality, any shading device located near a window will cast shadows on different parts of the window at different times of year – *all* of which must be considered in weighing the costs and benefits of the device.

Another shortcoming of the methods is the notion that a window should be fully shaded during all time segments in which the reference space is ‘too warm’ (or is being actively cooled). In some situations – as when direct solar gains are greater than the cooling load – shading the entire window to avoid cooling may create a need for heating. If the induced heating load is greater than the replaced cooling load (adjusted for factors of interest such as energy use or carbon emissions), the net effect of full shading may be undesirable.

SHADERADE METHOD

In order to address these limitations of existing methods, a new approach called *SHADERADE* was developed. *SHADERADE* traces back individual rays from different parts of a window along different solar angles in order to quantify their competing

effects over a 3-dimensional ‘shading volume.’ The approach resembles the ‘cellular method’ described by Kaftan and Marsh (Kaftan and Marsh 2005), but uses a vector-based volumetric analysis, and a fractional rather than a binary evaluation of beam desirability for each time segment. I.e., rather than asking whether shading is “beneficial” or “undesirable,” the method asks, “What *fraction* of transmitted solar beam energy is desirable at any hour of the year?” In order to answer this question, SHADERADE bases its analysis on load calculations of an investigated space without an external shading system. In the SHADERADE implementation described below, EnergyPlus is used as the thermal simulation engine.

Valuation of Transmitted Beam Energy

According to the method, the quantity of transmitted solar beam energy that is *undesirable* equals the cooling load at each time segment. (If the transmitted beam energy is less than the cooling load, then the undesired quantity is simply 100%.)

The *desired* fraction of transmitted beam energy at any time step, $T_{desired}$, can therefore be expressed as follows:

$$(1) T_{desired} = \begin{cases} 1, & E_{cooling} \leq 0 \\ \frac{E_{beam} - E_{cooling}}{E_{beam}}, & 0 < E_{cooling} < E_{beam} \\ 0, & E_{cooling} \geq E_{beam} \end{cases}$$

In this study, $E_{cooling}$ in Equation 1 was replaced by the quantity $(E_{cooling} - E_{heating})$. This formulation was necessary because the considered time segment (an hour) comprised multiple steps in the thermal simulation, making it possible to have both heating and cooling loads in one time segment.

Resolution of Shading Conflicts

SHADERADE circumvents the ‘spill-over’ effect at shading-period boundaries by finding *all* the time segments during which a location in space has a potential shading effect on a window. The method works by projecting the full set of hourly solar vectors from hundreds of sample points distributed evenly across the reference window. These vectors are projected into a user-defined exterior volume, which can be of any shape and defines the spatial boundary of all possible shading locations to be considered. The volume is divided into a closely-spaced uniform grid, and each cell in the grid is checked for intersections against all projected vectors.

Determining the desired transmittance of a particular cell is then an act of conflict resolution between the cell’s intersecting vectors. SHADERADE resolves the conflict by picking a cell transmittance, T_{cell} , that minimizes the *sum* of the absolute value, $|E_{cooling} - E_{heating} + E_{beam}(T_{cell} - 1)|$, over the cell’s full set of intersecting vectors (V). To do this, T_{cell} is simply

varied from 0 to 1 in increments of 0.01 to find the value producing the smallest sum.

It should be noted that this approach ignores any interaction between cells. That is, the desired beam transmittance of a cell is not adjusted according to the recommended transmittances of other cells. The arguments in favor of this approach are twofold. First, treating each cell in isolation is a much *faster* solution than recursively negotiating communications between hundreds or thousands of cells. And second, a reductive analysis of each cell, unattached to a particular test surface, is the only way to approximate qualities *inherent* to a particular point in space. This makes it possible to evaluate and visualize a *volume* of space comprising an infinite number of possible surfaces, rather than having to conduct analyses one test surface at a time.

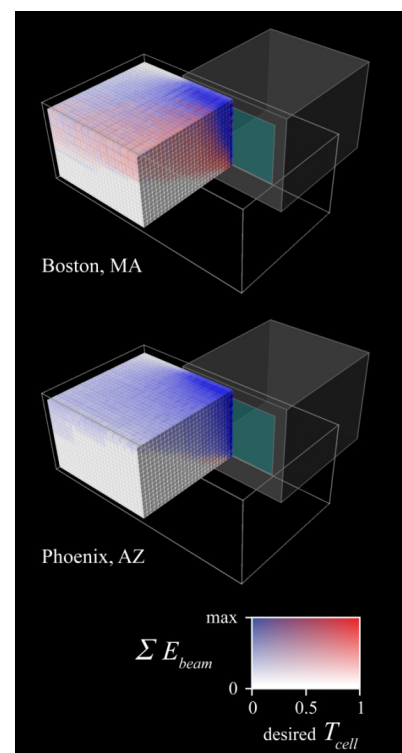


Figure 3 Perspective view of sections through 3-d shading volumes

Figure 3 shows sections through the ‘shading volumes’ for Boston and Phoenix. The load-optimized transmittance of each cell is mapped as a color from blue ($T_{cell} = 0$) to red ($T_{cell} = 1$). In order to visualize the amount of energy ‘at stake,’ the sum of vector beam energies ($\sum_V E_{beam,i}$) for each cell is mapped as the color’s saturation.

Opaque Shading Surfaces

For the special case of opaque shading objects, a decision about inclusion or non-inclusion of a cell can be made by comparing the initial loads to those ‘predicted’ for the cell having a transmittance of zero. This difference is referred to as the *opaque cell inclusion value (OCIV)*. If OCIV is positive,

inclusion of the cell is considered desirable. If negative, inclusion is considered undesirable.

For any base-surface, an opaque device can be designed by accepting all—or some fraction—of cells with an *OCIV* greater than zero (Figure 4 middle). There are two reasons a designer might wish to accept only a fraction of the *OCIV*-positive cells. First, issues such as constructability, visual comfort, aesthetics, and view may call for a device that is smaller than the *OCIV*-positive region. And second, testing a series of smartly-selected sub-regions can help negotiate the effects of artificial lighting and diffuse solar gain – elements ignored by the simple arithmetic model described above.

Because the degree of importance is already built into value of the cells, a sensible selection strategy involves rejecting the lowest-value cells first. In order to produce the largest reduction in area while maintaining, say, 90% of the value of the device, one can dispose of the lowest-value cell until the integrated surface value is reduced by 10%. This trimming process is illustrated in Figure 4 using an arbitrary curved surface generated in Rhinoceros. For this surface, a value reduction of 10% resulted in a 51% reduction in surface area.




Figure 4 Diagram showing method of trimming a surface using opaque cell inclusion values (*OCIV*)

In this study, the SHADERADE valuation and trimming strategies were applied to overhang- and perimeter-plane geometries. A 3x13-meter (*d* x *w*) overhang above of the window, and a 3-meter-deep

“perimeter” shade extruded from the top and side edges, were used as base surfaces. From each, surface cells with negative opaque inclusion values were discarded to minimize either direct loads or equivalent carbon emissions. The *OCIV*-positive region was then trimmed, as previously described, using an integrated surface value retention of 85%.




Figure 5 Theoretical translucent shading surface composed of cells of energy-optimized transmittance.

Translucent Shading Surfaces

Translucent shading devices can be also designed by using the desired values of T_{cell} as actual surface transmittances. In this study, a translucent shading option was tested in the form of a theoretical box composed of hundreds of squares of different opacities. The load-minimizing translucent box for Boston is depicted in Figure 5.

Table 2 Reference-Space Thermal Model Properties

Zones	Single, 3.5 x 4.0 x 3.0 meters (<i>w</i> x <i>d</i> x <i>h</i>)
Glazing	South façade, 50%
Adiabatic Surfaces	E, W, N, floor, ceiling
Shadow Calcs	Full interior/exterior with reflections
Weather	TMY3 epw files, Department of Energy
Construction	ASHRAE 90.1-2007 (for location-specific, commercial building -- EnergyPlus example file generator)
Natural Ventilation	Off
Cooling Set	24 °C, Mon-Fri, 9 – 5
Heating Set	22 °C, Mon-Fri, 9 – 5
Heating Setback	12 °C, all other hours
People Load	0.0391 m ⁻²
Electric Plug Load	8.07 W/m ²
Gas Plug Load	0.30 W/m ²
Lighting Load	10.80 W/m ²
Daylighting Control	398 lux
Target Illuminance	

COMPARISON BETWEEN SHADERADE AND EXISTING METHODS

In order to compare the effectiveness of static shading devices generated with SHADERADE as opposed to the previously discussed existing methods, an example analysis was conducted for a south-facing private office in the U.S. cities of Anchorage (61.18° N, 150.00° W, 35m), Boston

(42.37° N, 71.02° W, 6m), and Phoenix (33.45° N, 111.98° W, 337m) (Table 2). For each location, up to fifteen different 'conventional' shading geometries were generated from a combination of three date selection methods (equinox, degree-day, and thermal selection), three hour selections (noon, 10-2 and 9-3) and two form finding techniques (Ecotect and 2d method). Using SHADERADE, three additional shading devices were generated for each location: a horizontal overhang, a perimeter shade, and a translucent box. A "standard" integrated-value retention of 85% was used for the overhang and perimeter shades.

Heating and cooling loads were simulated using the EnergyPlus thermal simulation engine, version 4.0. The input data files (IDFs) for EnergyPlus were generated using custom Rhinoceros/Grasshopper components. The SHADERADE shading devices were generated using another custom component, also in Rhinoceros/Grasshopper.

RESULTS

Figure 6 graphs the percentage load and carbon-emission reductions compared to the unshaded office offered by all eighteen shading devices for Anchorage, Boston and Phoenix. The load graph (top) shows the percentage reduction of combined heating and cooling loads. The carbon graph (bottom) shows the percentage reduction of equivalent carbon emissions, taking into account electric lighting and plug use as well as COPs of 1.67 for cooling and 0.83 for heating. Carbon equivalent fuel factors of 0.232 and 0.758 kg/kWh were used for gas and electricity respectively, in accordance with section 7.5.3 of ASHRAE Standard 189.1 2010.

The SHADERADE group produced the largest average savings across the three climates, with a mean load reduction of 26% and a mean carbon reduction of 14%. The method was also the most consistent, with the three SHADERADE solutions always being within 3% of the best performing shading device for all climates.

The *degree-day, 10am-2pm* method also produced effective results, with mean load and carbon savings of 20% and 11% for both 2d-section and Ecotect-Shading-Wizard overhangs. In Anchorage, however, the method simply failed by not identifying any shading period – leading to zero savings. The method's indifference to building size, materials, or internal loads could mean that its relative success in this study is may be difficult to reproduce for other building types.

DISCUSSION

Suitability

Shading design inevitably involves considerations other than thermal performance. Cost, material




Figure 6 Percentage load and carbon reductions offered by simulated shading devices

efficiency, aesthetics, access to natural light, and view, can – and should – play a role in the final determination of form. In this regard, many of the large. The *9am-3pm degree-day* Ecotect shade in Phoenix, for instance, produced a 19% reduction in carbon emissions, but had a surface area of 21 m². The SHADERADE 85%-value devices, at 4-13 m², were more reasonably sized (no conventional device was smaller in all three climates), but were hardly minimal. In order to produce suitably sized devices, and adapt to the constraints of a particular project, a designer ought to know which parts of a shade are most necessary for thermal performance, and which are most expendable. In this regard, the

SHADERADE method was uniquely useful because it permitted shading-size reductions that minimized degradations in load or carbon reductions.

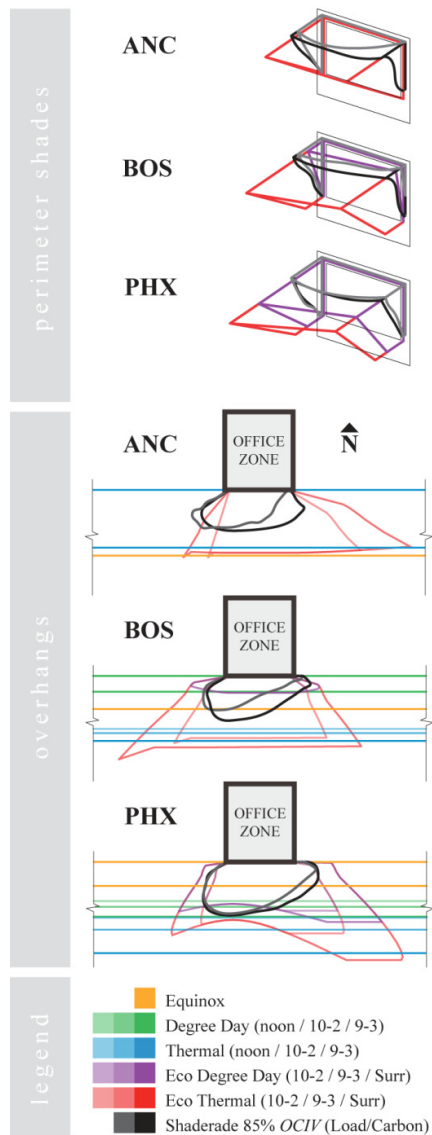


Figure 7 Axonometric view of perimeter shades (top) and plan view of overhang shades (bottom)

To illustrate this capacity, a 2-meter-deep overhang in Phoenix was subjected to size-reductions using two methods: (1) simple depth modulation, and (2) SHADERADE's OCIV-biased trim. The area reductions were varied from 0 to 100% in increments of 5%, and the resulting geometries run in EnergyPlus to assess carbon savings. In addition, the "predicted" effect of area-reduction was calculated for each trimmed surface, based on its retention of integrated, carbon-weighted OCIV. Recall that OCIV assigns an importance to each shading cell, and hence can be used to chart the expected drop-off in performance as portions of the shade are lost.

The results of the simulations and predicted model, shown in Figure 8, suggest two important lessons.

First, the discrepancy between predicted and simulated behavior is a reminder of the model's limitations. In particular, the failure to account for diffuse gains and artificial lighting can cause an over- or under- valuation of the shade as a whole, hence necessitating simulations of a series of shade sizes. Second, for any particular shade size, the discrepancy between depth-modulated and SHADERADE-trimmed performance shows the latter to be a superior method of area adjustment. At an area-reduction of 60%, for example, the SHADERADE trim maintained a carbon savings of 16%, while the savings provided by the depth-modulated overhang had dropped to 13%. This separation occurred because SHADERADE is better at determining the relative value of the shade at each point on the surface, and can therefore remove portions in a manner that minimizes performance losses.

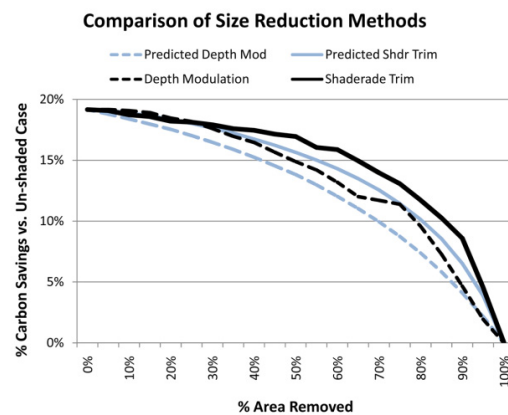



Figure 8 Shade with OCIV trim isocurves, plan view (top). Consequent carbon reductions (bottom)

Flexibility

In adapting to the intentions of a particular project, a shading methodology should allow not only intelligent size adjustment, but also a high degree of formal flexibility. On this criterion, the performance of the three form-finding methods varied greatly. The 2d-section method was the least flexible, unable to consider irregularly shaped windows and limited to the extrusion of curves drawn in two dimensions. The Ecotect Shading Wizard was more flexible, able to accept arbitrarily shaped window openings, and providing an option for the trimming of any user-specified, 3-dimensional mesh. The SHADERADE

method was most flexible, offering the same freedoms as Ecotect, but operating on the parameterizable NURBS curves of Rhinoceros/Grasshopper rather than imported dxf meshes. In addition, the SHADERADE method distinguished itself through the visualization of gradients of shading desirability and importance in three-dimensions—which allowed the designer the option of making informed *manual* decisions about the shaping and trimming of devices.

Current Developments

Despite increased freedom in modelling and visualization, SHADERADE's flexibility was hindered by its inability to consider redundancy in beam-object intersections. This limitation stemmed from the evaluation of a contiguous volume of shading cells, in which redundancy is a given. For the case of specific input surfaces (such as arbitrarily shaped lovers), the method is currently being expanded to handle multiple intersections.

Also in development is a method of “building” shading devices that picks one cell at a time and reassesses solar gains at each step. This method, which introduces cell-to-cell communication, has the potential to account for diffuse solar shading, and to coordinate the casting of non-uniform shadows over the face of a window. Both methods, however, must contend with the slower execution times that inevitably follow from recursive assessment.

Tool Availability

SHADERADE is tentatively scheduled to be released as an addition to the DIVA-for-Rhino plug-in (<http://www.diva-for-rhino.com/>).

CONCLUSION

The SHADERADE method reliably matched or exceeded the load reductions and carbon savings of conventional approaches, and provided an efficient means of device size adjustment. Building on previous work - notably Kaftan's cellular method - SHADERADE refined the question of “when” to shade by asking “when and how much?” By linking EnergyPlus with Rhinoceros, the method also established a new workflow between a robust thermal simulation engine and a nimble CAD modeller that enjoys increasing popularity among architectural design students.

Among the challenges to be tackled by workflows of this sort is the proper level of communication between formal exploration and simulated thermal performance. The linkage to EnergyPlus allows for a continual reassessment of form that is more accurate than SHADERADE's simple arithmetic model. Given this power, it is tempting to imagine a shading design process driven entirely by formal parameters, whose arguments evolve in search of an optimal

thermal result. This temptation, however, should be buffered by two important caveats. The first, already alluded to in this paper, is the issue of speed. Simulations themselves are time consuming, and, for designers, there is arguably a difference in kind between feedback that takes minutes and that which takes less than a second. One is appropriate for fine-tuning a device, while the other is appropriate for the broad conceptualizations of early-phase design. The second – and perhaps more fundamental – caveat, is that the realm of possible forms is infinite, and inductive searches within it are therefore inevitably incomplete. A method, such as SHADERADE, that is capable of approximating characteristics *inherent* to potential shading locations, has already taken a giant leap toward productively supervising the boundless space of possibility.

ACKNOWLEDGEMENT

The writing of this manuscript was supported by the “Sustainability in Design” Research Lab at the Graduate School of Design at Harvard University.

REFERENCES

- ANSI/ASHRAE/USGBC/IES. 189.1 Standard for the design of high-performance green buildings. Atlanta, GA, USA: ASHRAE standard; 2010.
- Arumi-Noé, F., 1996: Algorithm for the geometric construction of an optimum shading device, *Automation in Construction*, 5[1], pp. 211-217.
- Etzion Y. (1992). An improved solar shading design tool. *Building and Environment*. 27(3), 297-303.
- Kabre, C., 1999: WINSHADE: A computer design tool for solar control, *Building and Environment*, 34[3], pp. 263-274.
- Kaftan, E., Marsh, A., 2005. Integrating the cellular method for shading design with a thermal simulation. *Passive and Low Energy Cooling Conference (PALENC)*, Santorini, Greece, 2005. 965-970.
- Kaftan, E., 2001. The Cellular Method to Design Energy Efficient Shading Form to Accommodate the Dynamic Characteristics of Climate, *Proceedings, PLEA 2001 - The 18th Conference on Passive and Low Energy Architecture*, Florianopolis, Brazil, November 7-9, V2, 829-833.
- Marsh, A., 2003: Computer-optimised shading design, *Building Simulation 2003*, International Building Performance Simulation Association (IBPSA), Eindhoven, Netherlands, August 11-14, 2003.
- Yeziro A., Shaviv E., 1994: Shading: A design tool for analyzing mutual shading between buildings, *Solar Energy*, 52[1], pp. 27-37.