A Generic Physically-based Approach to the Opening Design Problem

K. Kalampokis[1], G. Papaioannou[1] and A. Gkaravelis[1]

[1] Department of Informatics, Athens University of Economics and Business,Greece

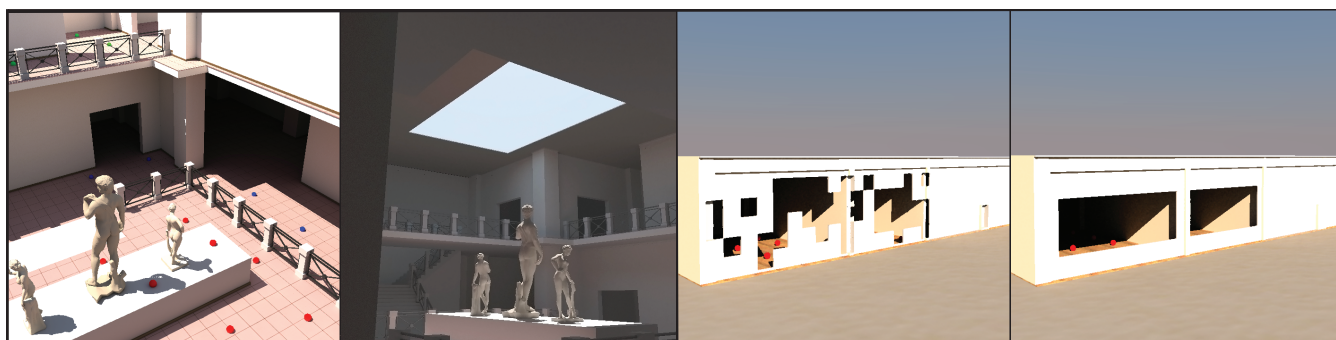Figure 1: From left to right: Definition of lighting goals in a scene; red sampling points correspond to maximum outdoor lightng contribution, blue to low lighting and green to darkness. Next, the computed skylight and illumination result. Third inset: The raw solution in the "apartment" scene and finally, architecturally admissible results after imposing geometric constraints to the opening clusters.

Abstract

Today architectural design harnesses photorealistic rendering to accurately assess energy transport for the design of energy-efficient buildings. In this context, we present an automatic physically-based solution to the opening design problem, i.e. the goal-driven process of defining openings on the input geometry given a set of lighting constraints, to better exploit natural daylight. Based on a hierarchical approach that combines a linear optimization strategy and a genetic algorithm, our method computes the optimal number, position, size and shape of openings, using a path tracing-based estimator to precisely model the light transport for arbitrary materials and geometry. The method quickly converges to an opening configuration that optimally approximates the desired illumination, with no special geometry editing requirements and the ability to trade quality for performance for interactive applications. We validate our results against ground truth experiments for various scenes and time-of-day intervals.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling-Physically based modeling—I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—I.3.8 [Computer Graphics]: Applications—

1. Introduction

Lighting design is of high importance in architectural design and in particular for minimizing energy consumption or achieving desired or comfortable illumination levels in buildings. The *Opening Design Problem* constitutes one of the most important sub-problems of the *Inverse Geometry Problem* (IGP), which studies the modification of the scene's geometry to accomplish the desired illumination effect. It aims at the determination of openings on the geometry, such as windows or skylights, in order to optimize the contribu-

tion of the environment lighting to the illumination of the scene, for a given set of lighting goals (illumination intentions on selected sampling points).

Methods that consider global illumination include the work by Mahdavi et al. [MBK95], who built a radiosity-based simulator for architectural lighting design that embodies for the first time a daylight-driven design of openings, using an interactive forward lighting evaluation. The first complete and formal approach to the inverse simulation of openings was conducted by Tourre et al.

[TMH08]. The authors presented an inverse daylighting framework that creates a geometrical reconstruction of the scene. The surfaces are subdivided into a number of triangular elements that are candidates for opening, solving the problem using a radiosity-based simulation. Their work implements their formal expression of lighting intentions in architectural design, presented in [TMH06]. In a similar manner in [FB12], an inverse lighting algorithm is proposed that models translucent skylights as diffuse light sources. The method is augmented in [FB15] by combining the previously described pinhole modelling technique proposed in [TMH08] with a hemi-cube approach. The above methods lead to a complete simulation considering both artificial and natural light sources, but the limitations in light transport evaluation introduced by the radiosity algorithm and the pinhole approximation to opening elements of non-negligible area, introduce errors in the process.

In this paper, we present a physically-based opening design method that determines the number, location and shape of the openings, given the designer's lighting intentions and a description of the scene, including atmospheric conditions and time intervals of interest as well as the set of candidate surfaces to host openings (see Figure 2). The main contributions of our work with respect to the state of the art are the following:

- No simplification of the discrete elements of the opening domain. Elements are modelled as CSG cutters and properly influence the interior lighting.
- Arbitrarily complex light transport, leading to a physically correct simulation.
- Identification of sources of error, common to all discretization-based opening design approaches and error compensation mechanisms to address them.
- No special geometry representation and modelling, as no editing or subdivision of the geometry is required or performed as a preprocessing step.

2. Method Overview

Our algorithm formulates the opening design problem as a linear zero-one optimization problem, using an arbitrarily fine quantization of the *opening domain* into candidate opening *elements*. The opening domain D_O (Figure 2) is the subset of surfaces on which openings are allowed. This "dicing" phase is performed virtually, by ray-tracing-based constructive solid geometry operations during the estimation of the contribution of each element, without requiring the modification of the geometry itself.

Illumination is measured at user-defined sampling locations s_i, where illumination goals are also provided in the form of irradiance levels E_i. Every elementary opening is then mapped to every sampling point, transforming the problem into a zero-one optimization task. For this, we first employ a hierarchical approach, which uses a non-negative linear system solver to provide a first estimation of the solution. Finally, the results are further optimized by a genetic algorithm, which imposes additional shape constraints for the determination of opening elements that create meaningful openings, such as rectilinear window boundaries for architectural design. A high-level overview of the overall pipeline is shown in Figure 3.
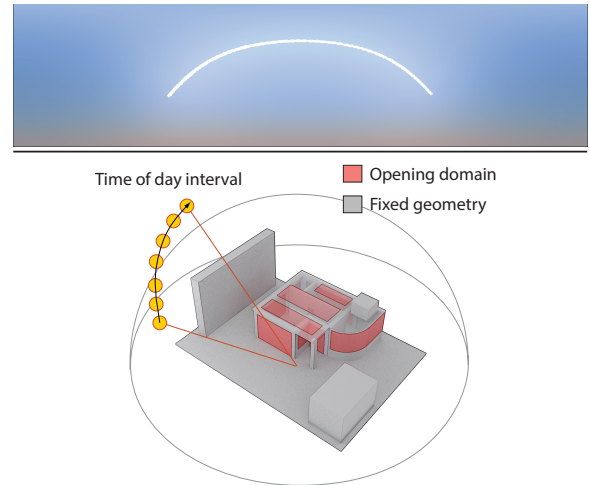

Figure 2: Input of the method: the geometry of the environment, marked surfaces where openings are allowed (opening domain) and daylight parameters that define an average environment map for the corresponding time of day interval and geographical location.

Input. The opening domain is defined by simply tagging certain surfaces during modelling and we implement this mechanism via scripting in Autodesk's 3D Studio Max. The irradiance E_i levels at the sampling locations s_i are also given by directly painting *normalized* irradiance levels on the geometry. The normalized irradiance corresponds to a range between darkness and maximum captured lighting at s_i when *all* elements are open. Alternatively, a uniform level of illumination over a subset of sampling points may be also specified, in which case, the expected irradiance E_i may be additionally constrained within user-specified limits E_{min}, E_{max}. During optimization, values outside this range are further penalized.

Finally, the desired lighting conditions for the period of interest must be also provided, including the parameters for the daylight system and the time of day interval for which we are interested in running the simulation. In particular, for the sky and sun simulation model, we rely on the Preetham model as described in [PSS99] to generate an average HDR environment map for either a single moment in time or a time lapse period, prior to running the simulation.

Opening contribution estimation. The surfaces comprising the opening domain are regularly sampled and a box-shaped virtual "cutter" with equal sides a is considered for each location, as shown in Figure 2, although cutters of any shape could be used instead. The cutter position and orientation follows the texture parameterization on the opening domain, but any alternative parameterization can be used instead. For the estimation of the contribution of each elementary opening j to the each sampling point s_i, path tracing is used, inherently including complex light transport and arbitrary materials as well as indirect lighting and shadowing from outdoor elements and other structures. The result of this stage is an $N_o \times N_s$ *opening contribution matrix* $\mathbf{E}$, relating the irradiance (contribution) of each one of the N_o opening elements to the N_s sampling points s_i considering the rest of the opening elements closed. Typically $N_o > N_s$.
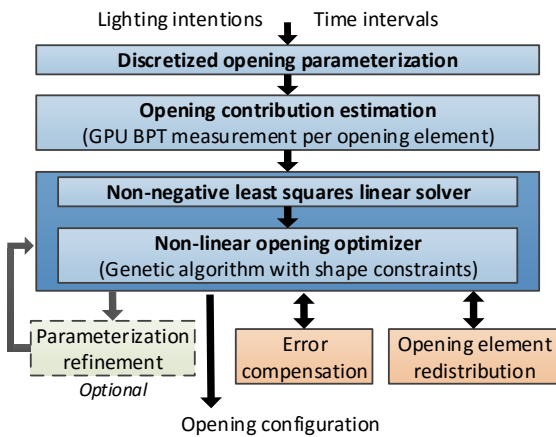
Figure 3: The pipeline of our opening design approach.

Two-stage opening computation. In general, the cumulative lighting contribution of multiple light paths crossing the opening domain can be considered additive up to an error factor $e_{i,j}$ explained and addressed in the following text. This means that for a given sampling location $\mathbf{s}_i$ with an irradiance goal E'_i the measured irradiance E_i is simply:

$$E_i = \sum_{j=1}^{N_o} o_j(E_{i,j} - e_{i,j}) \quad (1)$$

where $E_{i,j}$ is the irradiance contributing to point $\mathbf{s}_i$ through the opening j and o_j reflects the transparency of the opening. Eq. 1 allows the estimation of $E_{i,j}$ independently for every element - sampling location combination *before* entering the optimization loop, significantly reducing its complexity and providing reusable data. Eq. 1 forms an under-determined system of N_s equations and N_o unknowns. To solve this, we exploit the speed of a non-negative linear systems optimizer to bootstrap a non-linear optimization stage closer to the desired solution. The solver is based on the algorithm presented in [LH74] by C.Lawson and R.Hanson. It solves the linear mean squares optimization problem using non-negative values in the solution vector:

$$argmin = \left\| \mathbf{E} - \mathbf{E}' \right\|_2, \ \text{SUBJECT TO} \ o_j \geq 0, j = 1...N_o \quad (2)$$

The output of this stage contains non-binary values. To this end, we perform a redistribution of superfluous contributions, i.e. $o_j > 1$, to the N-8 neighbouring opening elements. Finally, all values are rounded to the either 0 or 1 and the result vector is used as an initial state of a genetic algorithm, which is convenient for the incorporation of structural constraints in the opening patterns formed, such as the compactness and boundary shape of the openings. As it is illustrated in Figure 3, these two steps are optionally combined in a hierarchical manner to quickly obtain a relevant starting solution using a coarser discretization of the opening domain. This solution is further fine-tuned by a genetic algorithm that stochastically mutates the solution favouring both error minimization and shape coherence.

Shape constraints In general, the ODP solution generated is not architecturally admissible, as it may contain irregular opening configurations. In architecture design, windows and skylights are usually expected to have rectangular shapes and comply with size specifications. Thus, the user can impose constraints on the shape, minimum and maximum covered area and aspect ratio of the openings. Our experiments showed that favourable results emerged by allowing for the unbiased evolution of the two-stage procedure before imposing shape restrictions as a separate, final refinement step. This step first clusters open elements with N-8 proximity into regions. Next all line/column combinations within the bounding box of the region are tested, excluding region-splitting and oversized/undersized combinations, and the configuration that best matches the contribution of the original region is retained. Regions that do not meet aspect ratio constraints are further split along their major axis. Resulting openings are re-clustered and the process is repeated until a stable number of openings is achieved.

3. Error Estimation and Compensation

Discretising the opening domain and considering one element at a time introduces two sources of error. The first one is related to the wall's thickness d. When opening each element individually instead of simultaneously enabling a set of them, some light will be either blocked or deflected by the side faces of the closed elements. As demonstrated in Figure 4, the severity of the error depends on the ratio between d and the opening element's side a. Thinning down the elements to near-zero thickness eliminates the deflection and shadowing problem, but introduces an over-estimation of direct lighting and scattering on the rim of the final opening. Both of these effects however are far less pronounced, especially when using small elements. Therefore in our method we modify opening domain walls to near-zero-thickness, as even openings with $d << a$ can cause significant shadowing for oblique direct lighting directions.

The other source of error, e_i in Eq. 1, emerges during the irradiance evaluation for each individual element $E_{i,j}$, due to the assumption that the other elements $k \neq j$ are *closed*. If a k element is closed, while in the final solution it should be open, light is reflected on its internal surface instead of leaving the interior space through the opening, as illustrated in Figure 5. This results in an increase of illumination levels. To rectify this, during the lighting evaluation step we separately accumulate a) an $N_o \times N_s$ matrix with elements $A_{i,k}$ storing contributions of paths starting from sample $\mathbf{s}_i$ and hitting the inside surface of opening k, b) an $N_o \times N_o$ symmetric matrix whose elements $T_{m,n}$ store the average (monochromatic) light transport operator between elements m and n, when paths connect both. We stochastically define the superfluous energy (error) $e_{i,j}$ as the light that crossed opening j and contributed to point $\mathbf{s}_i$ after bouncing off each other element k that is currently open ($o_k = 1$). Since we store the cumulative energy from a closed element $A_{i,k}$ and we are only interested in opening interactions $j \rightarrow k$, the result is weighted by the relative average transport weight:

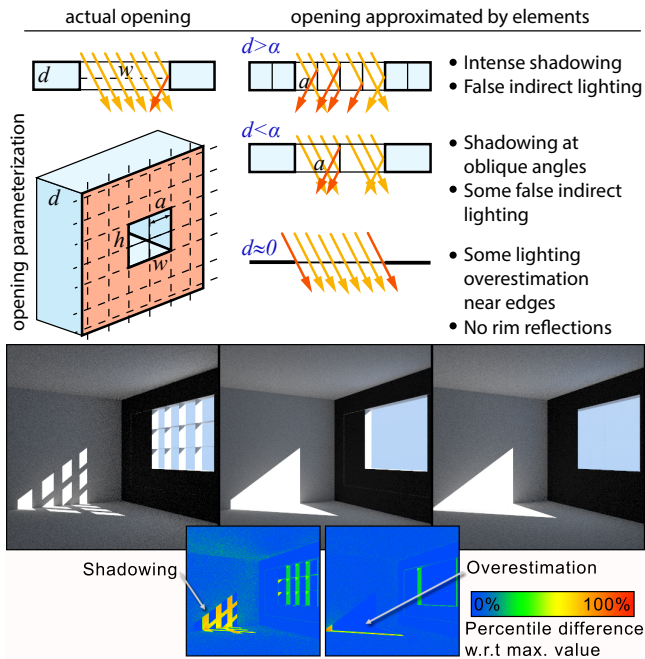$$e_{i,j} = \sum_{k \neq j} o_k A_{i,k} \frac{T_{j,k}}{\sum_m T_{m,k}} \quad (3)$$

Figure 4: The impact of the opening element size and wall thickness on the lighting evaluation, when individual elements are considered instead of the aggregate opening. Bottom: Stacked contribution of individual elements versus the actual opening for a thick and zero-thickness opening domain. The inside of the opening domain wall is purposefully black to factor out other sources of error.

Figure 5: Measuring the contribution from each element individually causes in-scattering from closed elements.

4. Method Evaluation and Conclusions

To compute in parallel the opening contribution matrix and the auxiliary error compensation buffers, we implemented a lighting measurement function based on GPU path tracing, using the NVIDIA Optix framework. We typically allow for at least 100 paths per sampling point-opening element combination,100 sampling points and 200 opening elements, which corresponds to about 1 minute on a GeForce GTX TITAN Black graphics card for the scenes in Figure 1. Also,the genetic algorithm is allowed to run for at least 500 to converge to an admissible solution. Note that increasing the number of paths during the evaluation decreases the genetic algorithm's running duration.

We validated our method against ground truth experiments (e.g. Figure 1 left), where a user-provided set of openings determines the illumination level at the sampling points. We subsequently seek to compute openings as the solution to the ODP solver that match the original, user-provided ones. In such cases, the resulting solution converges to the ground truth with error levels of up to 2%, depending on the number of paths, the actual wall thickness and the opening domain quantization size *a*. The apartment scene example in Figure 1 - right demonstrates an actual goal-driven ODP case, where the lighting intentions were manually defined as the irradiance on the sampling points in the two leftmost rooms of the scene. Our method achieves an initial error of 1.9% after the genetic algo-
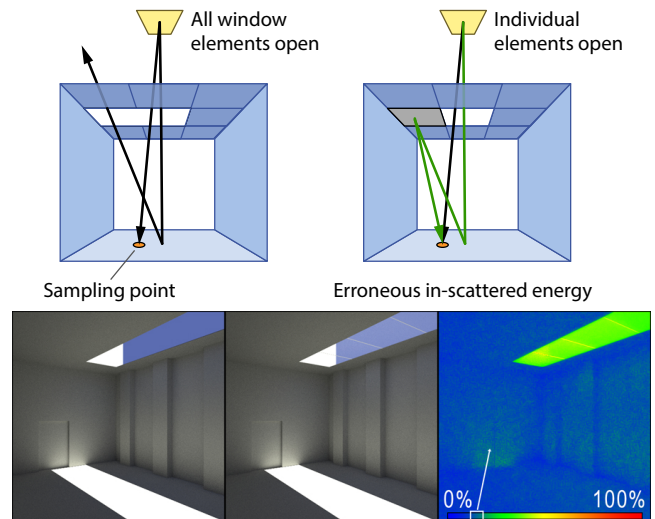
rithm stage (third inset from the left). The result is then processed to fulfil geometric constraints (rightmost inset).

Concerning the error compensation procedure, we experimentally confirmed that when factoring out the potentially overexposed user goals by ground truth experiments, the method managed to halve the error producing a more accurate evaluation of the proposed solution. It is not surprising that user-defined lighting intentions with arbitrary values produce significant error levels. For example, in the museum scene test case, sampling points around the statues are assigned the maximum possible value, while the rest of the ground floor is assigned low target illumination levels and the first floor is intended to be dark. The proposed solution produced an error of 9.3% creating a skylight that illuminates the atrium, which is a good solution, despite the increased error value compared to ground truth experiments.

References

[FB12]   FERNÁNDEZ E., BESUIEVSKY G.: Inverse lighting design for interior buildings integrating natural and artificial sources. *Computers & Graphics 36*, 8 (2012), 1096–1108. 2

[FB15]   FERNÁNDEZ E., BESUIEVSKY G.: Inverse opening design with anisotropic lighting incidence. *Computers & Graphics* (2015). 2

[LH74]   LAWSON C. L., HANSON R. J.: *Solving least squares problems*, vol. 161. SIAM, 1974. 3

[MBK95]   MAHDAVI A., BERBERIDOU-KALLIVOKA L.: A generative simulation tool for architectural lighting. *Evaluation 2*, 3 (1995), 4. 1

[PSS99]   PREETHAM A. J., SHIRLEY P., SMITS B.: A practical analytic model for daylight. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 91–100. 2

[TMH06]   TOURRE V., MARTIN J.-Y., HÉGRON G.: Architectural lighting ambience intentions for inverse rendering. *PLEMENOS, D., éditeur: 3IA 6* (2006), 211–218. 2

[TMH08]   TOURRE V., MARTIN J.-Y., HÉGRON G.: An inverse daylighting model for caad. In *Proceedings of the 24th Spring Conference on Computer Graphics* (2008), ACM, pp. 83–90. 2