LiteVis: Integrated Visualization for Simulation-Based Decision Support in Lighting Design

Johannes Sorger, Thomas Ortner, Christian Luksch, Michael Schwärzler, Eduard Gröller, and Harald Piringer

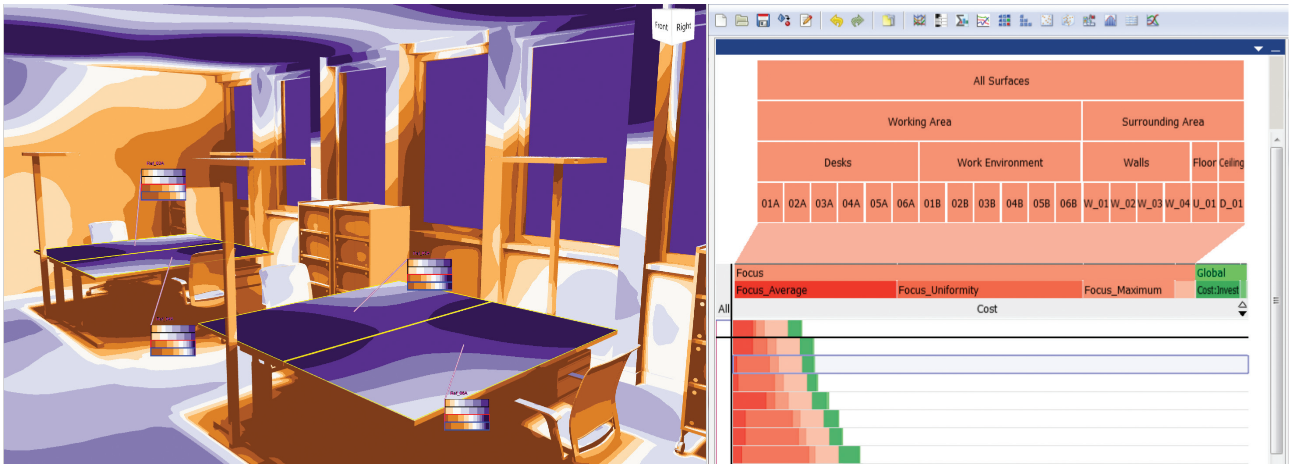

Fig. 1. Comparing lighting parametrizations according to different metrics in the LiteVis workspace. Left: the Simulation View displays a false color rendering encoding illumination quality. For each desk, floating annotations show a binned histogram of illumination values from selected simulation runs. Right: the Simulation Ranking View shows a ranking of simulated lighting parametrizations according to user defined importance values for spatial measurement surfaces and abstracted result indicators.

Abstract— State-of-the-art lighting design is based on physically accurate lighting simulations of scenes such as offices. The simulation results support lighting designers in the creation of lighting configurations, which must meet contradicting customer objectives regarding quality and price while conforming to industry standards. However, current tools for lighting design impede rapid feedback cycles. On the one side, they decouple analysis and simulation specification. On the other side, they lack capabilities for a detailed comparison of multiple configurations. The primary contribution of this paper is a design study of LiteVis, a system for efficient decision support in lighting design. LiteVis tightly integrates global illumination-based lighting simulation, a spatial representation of the scene, and non-spatial visualizations of parameters and result indicators. This enables an efficient iterative cycle of simulation parametrization and analysis. Specifically, a novel visualization supports decision making by ranking simulated lighting configurations with regard to a weight-based prioritization of objectives that considers both spatial and non-spatial characteristics. In the spatial domain, novel concepts support a detailed comparison of illumination scenarios. We demonstrate the applicability of LiteVis by using a real-world use case and discuss qualitative feedback from lighting designers. This feedback indicates that LiteVis successfully supports lighting designers to achieve key tasks more efficiently and with greater certainty.

Index Terms—Integrating Spatial and Non-Spatial Data Visualization, Visualization in Physical Sciences and Engineering, Coordinated and Multiple Views, Visual Knowledge Discovery

1 Introduction

Architectural lighting design has a strong impact on the atmosphere and aesthetics of buildings and open spaces. The illumination also affects how working environments support productivity and creativity. However, finding an acceptable trade-off between often contradictory customer requests regarding the quality and price of a lighting setup while conforming to industry standards is a challenging task. Physically accurate lighting simulations have long been used to support lighting designers in planning and communicating potential solutions to customers. Given a 3D geometric model of the scene as well as the position, type, and additional properties of the involved luminaries, the simulation computes the illuminance for each part of the scene.

Commercial software for lighting design [28, 27, 15] typically takes up to several hours to compute the result of a single setup, and focuses on inspecting a single solution rather than comparing many solutions. As a consequence, the workflow in lighting design has traditionally been restricted to computing a small number of alternative setups in a trial-and-error fashion. As an additional challenge, separate tools have been required for analyzing aesthetic and financial aspects of a given lighting solution [34]. How fast a convincing solution could be found, has therefore been highly dependent on the experience of the lighting designer.

Recent advances in lighting simulation reduced the effort for computing a physically accurate illumination to a few seconds [18]. Technologically, this enables significant improvements in the workflow of lighting designers. First, it is now possible to examine the design space much more systematically and comprehensively by computing hundreds of potential solutions as a starting point of the design process.

- *Johannes Sorger, Thomas Ortner, Christian Luksch, Michael Schwärzler and Harald Piringer are with VRVis Research Center. E-mail: {sorger,ortner,luksch,schwaerzler,piringer}@vrvis.at.*
- *Eduard Gröller is with TU Wien. E-mail: groeller@cg.tuwien.ac.at.*

Manuscript received 31 Mar. 2015; accepted 1 Aug. 2015; date of publication xx Aug. 2015; date of current version 25 Oct. 2015. For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

Second, lighting designers may now define and compute additional setups on-the-fly, e.g., during a workshop with a customer.

However, existing software tools in lighting design do not support these new possibilities well. For this reason, we conducted a design study in collaboration with lighting design experts. The result is LiteVis, a system for efficient decision support in lighting design. In contrast to existing tools, a focus of LiteVis is the comprehensive comparison of multiple solutions. LiteVis tightly integrates global illumination-based lighting simulation, a spatial representation of the 3D scene, and non-spatial visualizations of setup parameters and result indicators. Specifically, LiteVis relies on hierarchically structured measurement surfaces in the 3D scene for defining quantitative quality indicators. Weighting these spatial indicators along with non-spatial indicators, such as monetary costs, supports a holistic view of all relevant decision factors while making the decision maker's preference explicit and reproducible.

The *contributions* of this paper can be summarized as follows:

- A design study of decision making in lighting design resulting in the system LiteVis.
- A problem characterization of the application domain of lighting design, including an abstraction that relates the application problem to other areas of simulation-based decision making.
- A novel visualization for ranking multiple decision options based on weighting hierarchically structured objectives.
- A report on feedback from domain experts and a reflection on the design process.

2 Related Work

2.1 Lighting Design

In the scientific community, several approaches to semi-automatic or user-guided, interactive lighting design have been proposed, e.g., sketch-based methods, where the user "paints" the scene parts to be lit [29, 22, 32, 23, 16], as well as procedural methods [30]. Additionally, the direct specification of lighting-induced features, such as shadows or highlights [26, 24], have been suggested to trigger the generation of an ideal lighting solution. While all these approaches tackle the problem in plausible ways, they have not found their way into today's industry standard lighting design tools. Potential reasons are that some constraints are hard to express mathematically (e.g., regarding aesthetics), or the wish for artistic freedom. Glaser et al. [10] have approached the lighting design problem by developing various 2D visualization prototypes, inspiring our work in terms of analytical aspects, but without linking spatial 3D and non-spatial 2D views in a common tool.

Commercial lighting design applications, on the other hand, leave the task of placing and adjusting the light sources in a scene to the skills of a lighting designer. Tools, such as Relux [28], Dialux [27], Agi32 [15], rely on Radiosity and/or Raytracing-based simulation kernels. They visualize the simulated results on measurement areas placed in the 3D scene using false color visualizations. None of them supports the user in the comparison of different solutions. Separate tools, such as ecoCALC [34], tackle related problem domains like financial aspects. Even though their respective outputs are needed to form an overall decision, the tools used in the workflow of finding an ideal solution are oftentimes isolated from each other, which makes a holistic exploration of the problem space cumbersome.

We base our approach on a lighting modeling system that relies on a simulation kernel recently proposed by Luksch et al. [18]. By providing a faster lighting simulation as compared to Radiosity-based approaches, it enables shorter cycles of scene parametrization and evaluation.

2.2 Visual Parameter Exploration

The field of visual parameter space exploration has shown significant progress in the last years. Beham et al. [3] developed a composite visualization that combines the abstract parameter space of geometry generators with the output space of the resulting shapes. Illustrative parallel coordinates allow the user to study the sensitivity of parameters in a global-to-local drill down fashion. Similarly, in LiteVis, the relation between the input space and the output space can be explored. However, in our case, both have spatial and non-spatial properties. Coffey et al. [7] propose a tool for simulation-based design that provides integration of forward design (input manipulation) and inverse design. Inverse design lets the user query a database of pre-computed samples by specifying the desired simulation output directly in the 3D scene. This specification of goals in the 3D scene shares the same principle as the specification of spatial objectives in LiteVis (see Sec. 4.2). Bruckner and Möller [5] propose a tool for the exploration of clustered time series data of physical fluid simulations for visual effects design. Since their approach is result driven, the input is not of interest to the user. In contrast to the solution by Bruckner and Möller, we do not have to deal with dynamic scenes but our scenario demands an input- as well as an output-driven approach.

2.3 Multi-Objective Decision Making

In lighting design, multiple objectives must be optimized simultaneously. This is a common issue in many application domains. For this reason, multi-objective optimization (MOO) has long been an active field of research (see, e.g., Köksalan et al. [13] for a survey). As a single best solution does not exist for MOO-problems, one approach is to offer the decision maker multiple solutions that have been (semi-) automatically generated. Miettinen and Mäkelä, for instance, introduced an interactive method called NIMBUS [20] that asks the user to repeatedly examine the values of objective functions calculated for a current solution.

There are several approaches for visualizing the set of multi-objective optimal solutions, which is known as the Pareto Frontier. Korhonen and Wallenius [14] stress that visualizing the Pareto Frontier for more than three objectives is difficult. The authors classify several visualization techniques for multi-criterion decision making based on the cardinality of the result set. Lotov et al. employ scatter-plot matrices to show bi-objective slices of the Pareto Frontier [17]. In order to support decision making, Andrienko and Andrienko [1] propose several extensions to parallel coordinates, the prevalent technique for visualizing Pareto Frontiers [2]. More recently, Chen et al. [6] employ self-organizing maps for projecting all Pareto-optimal solutions to a 2D radial visualization. However, most of these approaches do not scale to a very high number of objectives, e.g., up to hundreds in lighting design if spatial aspects are taken into account (see Sec. 3.1). Moreover, inspecting the entire set of Pareto-optimal solutions can still be inefficient.

A common approach is therefore to weight the objectives in order to attain a score, which can then be used for ranking possible solutions for decision making. LineUp [11] is an interactive approach for weighting multiple objectives and visualizing the resulting ranking. While LineUp inspired our approach for decision support in the context of lighting design (Sec. 4.2), it does not scale to a very large number of objectives and does not take the semantics of these options into account, e.g., the underlying parametrization. As an additional problem for most approaches in lighting design, each solution refers to the illumination of an entire scene and thus has complex characteristics, such as aesthetics, which are hard to quantify.

3 Lighting Design Background

This section provides a brief introduction to the field of lighting design as far as necessary for understanding the design decisions that went into LiteVis. This information is based on experiences gained from a tight collaboration with experts in lighting design for nearly five years. After describing the data and the tasks, we motivate the key design goals of LiteVis. In a problem abstraction, we put the domain specific challenges into a broader context.

3.1 The Data: Simulation in Lighting Design

Modern lighting design is based on lighting simulations which compute an output illumination for a particular scene and luminary setup. As for many simulation types, the *input* of a lighting simulation can be classified as control parameters, environmental parameters, and model parameters. We refer to a particular assignment of values to all parameters that are necessary for starting a simulation as *parametrization*. We refer to simulated parametrizations as *solutions* or *simulation runs*.

The *control parameters* in lighting design, are *luminaries* and their position and orientation in the scene. They are placed by designers to achieve certain goals regarding the appearance and budgetary constraints of a solution. Luminaries are placed using controls similar to those in geometric modeling tools. Each luminary has several parameters. The most important ones in the context of this paper are:

- *Type*: Type corresponds to a particular product of a manufacturer which is typically classified by the application as, e.g., floor lamps, ceiling lamps, or wall mounted lamps.
- *Wattage*: This parameter describes the illumination power and energy consumption of a luminary.
- *Dim profile*: A luminary can be dimmed, resulting in lower power consumption and increased longevity, but also lower performance.

Environmental parameters describe aspects of the simulation which are not directly controlled by the user or vary over time:

- *Scene*: Contains the 3D geometry of the scene such as desks, windows, walls, etc. Geometry can absorb and reflect light based on material properties, and also act as an occluder.
- *Environmental conditions*: Such conditions determine external influences, like sunshine that illuminates a scene based on the time of day/year and the weather condition. These factors can optionally be included in the simulation.

Model parameters are the implicit parameters of the lighting simulation algorithm responsible for transforming the simulation input into the simulation output. Model parameters typically define a trade-off between accuracy and speed. An example is the number of times the light bounces in the calculation of the indirect illumination. These parameters are not of direct concern to the lighting designer and therefore play a minor role in scope of this paper.

The output of the lighting simulation comprises an illuminance value (measured in lux) for each texel in the scene. Using this direct output for 3D rendering of the scene is suitable for a qualitative inspection of the results. However, for most tasks, the direct simulation output is too detailed. The standard approach in lighting design for condensing the data and assessing the compliance to industry standards is based on pre-defined *measurement surfaces* on top of the geometry. Each measurement surface corresponds to a semantically meaningful part of the scene such as a particular desk, a certain part of the wall, a door, etc. A scene typically comprises multiple types of surfaces which are structured hierarchically. For example, a top-down classification of a particular measurement surface could be "office – working area – all desks – desk 3".

Measurement surfaces aggregate the illuminance of their geometry to define quantitative *local indicators*. The most important local indicators in lighting design are the minimal, maximal, and the average illumination, and uniformity. Uniformity is defined as the relation of the minimal illumination to the average illumination and describes the evenness of the distribution of illuminance values on a surface.

It is common to define target values for these local indicators per measurement surface and to assess the suitability of a solution in terms of the difference to the respective target values. Different classes of measurement surfaces are subject to varying degrees of constraints regarding the illumination quality, which is reflected in different target values (see Fig. 2). In addition to a customer-driven specification of target values, several classes of measurement areas such as desks must also meet industry norms and standards [8].

This description of lighting simulations in the context of lighting design is generally applicable to commercial systems as well as to the simulation underlying LiteVis [18]. In all cases, the simulation kernels of various software tools involve global illumination techniques. In commercial systems, the effort for computing the illumination of a single parametrization takes up to several hours, depending on the complexity of the scene and the configuration of the luminaries. Using the GPU and clustering virtual point lights into a set of virtual polygon lights, the simulation underlying LiteVis, however, reduces the effort for obtaining comparable results to a few seconds. We refer to Luksch et al. [18] for additional technical details of the lighting simulation.

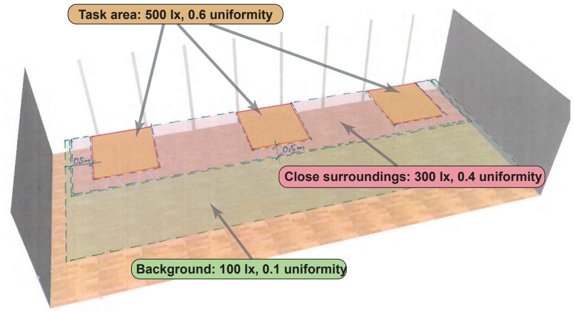

Fig. 2. Conceptional sketch made by a lighting design expert, defining the location of the measurement surfaces and their target values in the office scene according to industry standards [8].

Besides the simulation output and the indicators derived from this output, global indicators regarding budgetary values of the scene parametrization are crucial for the final decision. Specifically, each type of luminary has an investment cost as well as a run-time cost in terms of energy consumption.

Summarizing, the key entities of LiteVis are the simulated parametrizations. Regarding the data model, each parametrization comprises a multivariate part and a geometric part. The multivariate part includes (1) the parameters per luminary such as wattage, (2) environmental variables like the simulated time of day, (3) the local indicators per measurement surface and their deviation from the target values, and (4) the global indicators regarding monetary cost. The geometric part describes the spatial extent of the illumination in terms of an illuminance value per texel.

3.2 Task Analysis

Our collaboration with experts in lighting design enabled us to identify the following recurring tasks (Fig. 3). The subsequent list of tasks is based on insights gained from semi-structured interviews and contextual inquiries [12].

- **T0 - Scene setup**: After the scene is modeled by a 3D artist, measurement surfaces are defined on top of the geometry, e.g., for desks and walls. The designer defines target values and constraints for local indicators, such as uniformity per measurement surface, in accordance with customer requirements and industry norms.
- **T1 - Scene parametrization**: The lighting designer defines one or more parametrizations for a subsequent simulation. This particularly includes the placement and parametrization of a specific set of luminaries in the scene according to customer requirements, budget constraints, industry standards, and the designer's experience in order to achieve specific aesthetic effects. The designer may also specify certain environmental conditions, such as the position of the sun. The initial definition of parametrizations is mainly based on the experience and skill of the designer while later iterations are based on insights from previous solutions. In all cases, the generation of parametrizations may be partly automated based on parameter variations, e.g., for varying values of luminary wattage.
- **T2 - Assessment of single solutions**: Once a simulation has been completed, its qualitative and quantitative output is typically evaluated individually in order to quickly reject inappropriate solutions. The qualitative evaluation assesses the aesthetic appearance of the illuminated scene in false color renderings, e.g., regarding the strength of shadow gradients. The quantitative evaluation is based on the indicators of measurement surfaces and, for example, rejects solutions that violate industry norms.
- **T3 - Comparison of multiple solutions**: Multiple candidate solutions are compared in detail in order to understand in which regard one is superior. One goal is to further reduce the set of possible candidates by excluding those which are not superior to others according to the involved criteria. Another goal is to

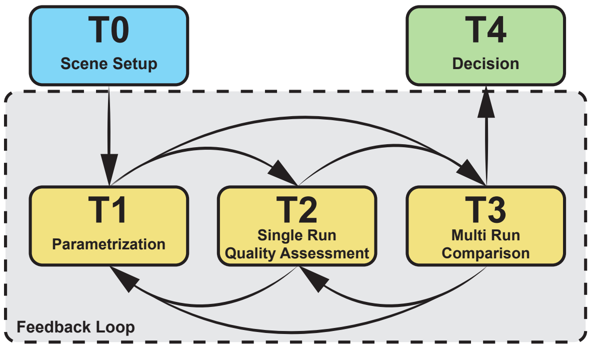

Fig. 3. The five steps of the lighting designer's workflow. After initialization of the scene (T0), the lighting designer works in an iterative fashion (T1-T3) until he or she finds a satisfactory solution (T4).

identify key trade-offs, e.g., between a certain aesthetic aspect and monetary cost. Traditionally, this task is performed using side-by-side comparisons of false color renderings and the corresponding data sheets.

- **T4 - Decision**: In order to arrive at a final decision, it is common to weigh the various result indicators. Qualitative aspects and a detailed comparison of the illumination of multiple solutions still play an important role at this stage. In many cases, this task is performed by the lighting designer together with the customer.

The traditional strategy for decision making in lighting design is based on trial&error. As computing a single solution took up to several hours, only very few alternatives were typically considered in a rather sequential manner. In some cases, an initial guess that appeared to be "good enough" was used without further comparison. In other cases, it could take up to several work-days to identify an acceptable solution, especially for less experienced designers or very complex scenes.

In LiteVis, the description of these tasks remains valid. However, the lighting simulation underlying LiteVis enables two important improvements to this workflow. First, it enables to initially specify and compute a set of numerous samples in one step within T1. For example, hundreds of variations regarding luminary placements, wattages, and environmental parameters such as external lighting conditions can be simulated automatically over night. This provides a more global understanding of the space of possible solutions at a very early stage of the design process. The understanding supports the fast identification of acceptable designs while avoiding to miss interesting solutions.

Second, the simulation is fast enough to enable the exploration of additional solutions in real-time. The improved feedback loop makes it possible to reach a decision during a single meeting with a customer. Previously, multiple consecutive meetings could have been necessary.

3.3 Design Goals of LiteVis

The design goals of LiteVis are motivated by the gap between the state-of-the-art software in lighting design and the new possibilities for improving the workflow as enabled by the fast simulation.

- **G1 - Holistic overview of all relevant aspects**: Currently, separate tools are necessary for a qualitative inspection of the simulated illumination, a quantitative analysis of result indicators, and an examination of budgetary constraints. This separation makes comparing multiple solutions cumbersome. A goal of LiteVis is thus to provide a convenient and integrated access to the simulated illumination, the local surface indicators, budgetary values, and the parameter values of a specific solution.
- **G2 - Integration of scene parametrization and evaluation**: The lighting simulation (T0, T1) and the evaluation of the solution (T2, T3, T4) are currently isolated from each other since they are performed with separate tools. This lack of integration impedes an efficient feedback loop between scene parametrization and evaluation. Consequently, a goal of LiteVis is to allow the user to easily trigger new simulation runs based on the inspection of prior results, and to display new results immediately when they become available.
- **G3 - Effective comparison of multiple solutions**: The comparison of solutions is essential for defining additional iterations and reaching a final decision. However, current simulation tools in lighting design do not offer the designer appropriate means to understand, compare, and make decisions based on alternatives. The status quo for comparing multiple solutions is to assess the corresponding data sheets side-by-side, which is ineffective for decision making. A goal of LiteVis is to enable a direct comparison of simulated illuminations as well as of indicators regarding illumination quality and monetary costs. This includes enabling a sensitivity analysis to assess the impact of parameter variations.
- **G4 - Explicit and reproducible decision support**: Different measurement surfaces typically have different target values for their illumination indicators as well as different importances in general. For example, the illumination quality on a table surface is more important than on a wall. Current tools consider neither distances from local target values, nor semantic importance factors. Moreover, different stakeholders may put more emphasis either on qualitative or on financial aspects. A key goal of LiteVis is therefore to support reproducible decision making based on an explicit preference specification for illumination-related as well as cost-related indicators. A related goal and a prerequisite is to enable the specification of target values for indicators.

3.4 Problem Abstraction

For problem abstraction, the domain of simulation-based lighting design can be characterized in terms of the conceptual framework for visual parameter space analysis by Sedlmair et al. [31]. The lighting simulation is a deterministic computational *input-output model*. It enables an integrated sampling of control parameters, environmental parameters, and model parameters as input (see Sec. 3.1) to generate the illumination per texel as *direct output*. In relation to the results of computational models in other application contexts, the direct simulation output can be considered a *complex object*. Apart from a qualitative assessment, this complex object requires a derivation step for subsequent decision making. In our context, the derivation is an aggregation of local indicators based on measurement surfaces.

The navigation strategy has traditionally been *informed trial and error* whereas LiteVis supports a shift towards *global-to-local*. While a sensitivity analysis can be interesting to lighting designers, the primary task is the *optimization* of multiple competing objectives corresponding to the indicators for illumination quality and monetary cost. Due to the hierarchical structure of the measurement surfaces, the corresponding objectives can also be considered hierarchical, which has motivated key design decisions of LiteVis.

4 Design Study of LiteVis

4.1 System Overview

The design of LiteVis comprises two tightly integrated parts, i.e., the spatial *Simulation View*, and the non-spatial *Analysis Views*.

The Simulation View (Fig. 1, left) provides a spatial display of the scene. While the 3D geometry of the scene is modeled externally in a 3D modeling tool, the Simulation View supports an interactive definition of measurement surfaces for scene setup (T0). It also supports the parametrization of new simulation runs (T1). To accomplish this, the user can create luminaries and interactively place their 3D representations in the scene via dragging. The user can inspect and modify luminary parameters, such as the value for wattage. Optionally, the user can specify a value range to define and simulate multiple parametrizations.

After computing simulation results, the Simulation View supports a qualitative assessment of the direct output of a single simulation run (T2) as well as a comparison of multiple runs (T3, see Sec. 4.3 and Sec. 4.4). As an alternative to rendering the illumination in a realistic way, a false color mode conveys the negative or positive distances to the respective target values (Fig. 1). While the comparison of parametrizations in the Simulation View is novel in the lighting designer's workflow, the scene parametrization and false color rendering represent familiar aspects of their routine.

The Analysis Views (Fig. 1, right) support the comparison and selection of simulated solutions in terms of local and global result indicators, and the underlying input parameters. Most importantly, the Simulation Ranking View enables a direct comparison of multiple solutions regarding their overall superiority as defined by a weighting of local and global indicators (Sec. 4.2). Additionally, various well-known multivariate visualizations such as parallel coordinates, bar charts, and spreadsheets may optionally be used to inspect, e.g., parameter values of the solutions. All analysis views support the selection of simulated solutions via brushing and are linked to highlight selections synchronously.

Integration of the Simulation View and the Analysis Views is implemented as follows: First, selecting one or more simulated solutions in the Analysis Views loads the corresponding parametrizations and illuminations into the Simulation View for a detailed inspection and spatial comparison. Additionally, the user may specify and simulate a new variation of the loaded parametrization by modifying parameter values. Second, selecting a representation of a measurement surface in either view, highlights this surface in both spatial and non-spatial domains. Third, creating and simulating a new parametrization in the Simulation View automatically extends the set of solutions in the Analysis Views. These types of integration enable an efficient feedback loop (G2) between scene parametrization (T1), the assessment of the result (T2), and a comparison of multiple simulation runs in a spatial and a non-spatial context (T3). The integration also helps the user in maintaining a mental connection between individual scene parts and their respective indicators.

4.2 Simulation Ranking View

As one contribution of this paper, we designed the Simulation Ranking View to match the requirements for efficient decision support (G4). The key idea is to rank multiple simulated solutions based on a notion of optimality that is defined as a weighted scoring of local and global indicators. The user should be able to define an emphasis on, e.g., certain parts of the scene, certain illumination indicators, such as uniformity, or on global indicators, such as investment cost. Furthermore, the user should be able to explore different preference settings by interactively changing the weights in order to see the effect on the ranking of the solutions. While fundamental aspects of the design were inspired by LineUp [11], the number and structure of involved indicators required some significant extensions.

4.2.1 Score Computation

Before explaining the design, it is necessary to understand the computation of the scores per solution. In this context, we face three main issues: (1) The score must provide a combined assessment of global indicators, such as investment cost, as well as local indicators, such as uniformity. (2) Each local indicator (e.g., uniformity) applies to a potentially large number of measurement surfaces, where each surface may be weighted differently. (3) Per measurement surface, each local indicator may have different target values.

The overall score of a solution is defined as the weighted sum of the score per objective. Objectives can be classified as spatial and non-spatial. *Non-spatial objectives* refer to global indicators, e.g., the investment cost and the cost per month. For each global indicator, a scoring function maps the indicator value to a score which ranges from zero to infinity. Our current implementation simply employs linear functions with a user-defined slope that map a cost of zero to a score value of zero. Smaller score values are therefore considered better. In other words, the score value of a particular objective represents the distance of the indicator from a user-defined target value. In the case of monetary costs, this target value is zero.

We note that this definition of score values is different from the one in LineUp [11], which maps each indicator to a bounded score ranging from zero to one, with one being considered as the best. The main motivation for inverting the scale of the scoring function and for avoiding an upper bound is to enable the expression of constraints. For example, a scoring function for the indicator "investment cost" may exclude

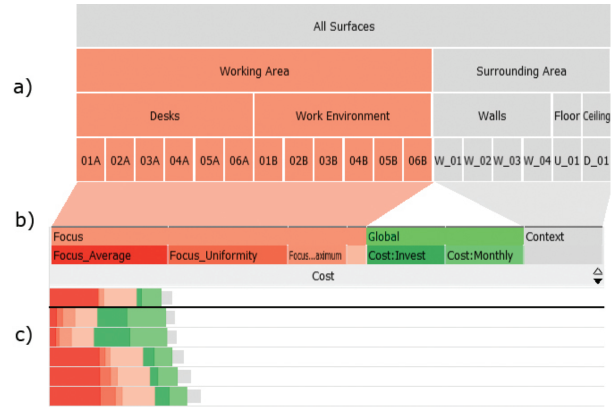

Fig. 4. The Simulation Ranking View: The Spatial Hierarchy (a) allows the user to specify the individual importance per measurement surface and to select a subset of surfaces (colored in red) for detailed inspection. The Indicator Bar (b) displays the weight per indicator. It also discriminates between global indicators (green), spatial indicators in focus (red), and the accumulated weight of objectives outside the spatial focus (gray). The Ranked Table (c) displays the individual solutions as stacked bars corresponding to the sum of weighted scores per indicator. Smaller bars correspond to indicator values which are close to a user-defined target and thus reflect superiority.

solutions where the indicator value exceeds a certain threshold by assigning very high score values to investment costs above the threshold.

Spatial objectives refer to the local indicators of measurement surfaces. Specifically, each combination of local indicator type (i.e., min, max, average, uniformity) and measurement surface defines a separate objective. For example, 18 measurement surfaces and four local indicators define a set of 72 spatial objectives. Also for spatial objectives a scoring function maps the indicator value for a particular measurement surface to the range between zero and infinity, where zero is considered as the best value. For each spatial objective, the user may define a separate scoring function. In this case, the user can define linear as well as non-linear scoring functions. For example, a function could assign score values of zero for a certain target range while increasing exponentially with growing distance from that range.

Given the score values for all objectives (i.e., spatial and non-spatial), *weighting* enables the user to explicitly specify a personal preference. Internally, a separate weight is maintained for each objective and all weights sum up to one. However, exposing several dozens of objectives and corresponding weights to the user is neither intuitive nor effective. We thus define the per-objective weights implicitly as the product of a *per-indicator weight* and a *per-measurement surface weight* in the case of spatial objectives. For global objectives, the per-objective weight is equivalent to the per-indicator weight.

Regarding the per-indicator weights, the user may specify a weight for each indicator type. In our scenario, there are six indicator types consisting of the two global indicators and the four local indicators. For example, the user may specify a weight of 0.3 for objectives regarding uniformity. All per-indicator weights sum up to one. Regarding the per-measurement surface weights, the user may independently define a weight for each measurement surface. The weights for all surfaces also sum up to one. For example, a per-indicator weight of 0.3 for uniformity and a per-measurement surface weight of 0.1 for surface "01A" would yield a per-objective weight of 0.03 for the objective referring to maximizing the uniformity of surface "01A".

4.2.2 Visual Encoding and Interaction

The design of the Simulation Ranking View consists of three visual components. The *Spatial Hierarchy* (Fig. 4 a) provides a hierarchically structured representation of measurement surfaces for specifying preferences in terms of per-measurement surface weights. The hierarchy is represented as an icicle plot [19] where each hierarchy level refines the degree of detail of the spatial subdivision. The root node corresponds to the entire scene and the leaf nodes represent the individual measurement surfaces. For all levels, the lengths of the nodes reflect the size of the underlying per-measurement surface weights. The user can modify these weights by dragging the boundaries between nodes on all levels. For example, reducing the size of the node "Desks" in Figure 4 would decrease the weights for the underlying measurement surfaces 01A to 06A while increasing the weights for the surfaces of "Work Environment". Moreover, the user may specify a *spatial focus* by clicking on nodes at any hierarchy level. Besides highlighting the geometric representations of the corresponding measurement surfaces in the linked Simulation View, the spatial focus enables the inspection of weighted scores for individual scene parts, as explained below.

As the second visual component, the *Indicator Bar* (Fig. 4 b) displays the sizes of per-indicator weights as a stacked bar. Vertically, the bar is subdivided into two levels. The upper level shows the overall sizes of weights for spatial indicators and non-spatial indicators. Hue is used to discriminate spatial (red) and non-spatial (green) indicators. The lower level displays the proportion of the specific indicators such as uniformity, average illumination, etc. The length of a bar on the lower level thus encodes the corresponding per-indicator weight. As for the Spatial Hierarchy, the user may modify weights by dragging the boundaries between two bars on either level.

If the user has specified a spatial focus, the lengths of the red bars correspond to the subset of spatial objective weights inside the focus. In this case, the Indicator Bar displays a separate monochrome gray area representing the sum of all spatial objective weights outside the spatial focus, referred to as *spatial context* (see Fig. 4). The definition of weights and thus the resulting overall scores are independent from a discrimination of spatial focus and context. A visual link connects the representations of the spatial focus and context in the Spatial Hierarchy and the Indicator Bar to illustrate the relation between these two components. The design choice to include the distinction between spatial focus and context in the Indicator Bar was motivated by user feedback. They considered this visual component as a suitable legend for explaining the color-coding of the Ranked Table.

The *Ranked Table* (Fig. 4 c) is the third visual component. It displays the individual solutions as stacked bars. The total length of each stacked bar corresponds to the overall weighted score of the respective solution. Solutions are ordered by their weighted scores. The topmost bar thus shows the best solution in the candidate list given the current weighting. The best solution is the one with the smallest weighted score. The individual bars that a stacked bar is composed of represent the weighted scores per indicator type for the particular solution. The color coding is identical to the one of the Indicator Bar.

The subdivision of the stacked bars provides a visual decomposition of the overall score. For each solution, this immediately reveals the indicators contributing most to the overall score. For example, it may show that the quality of a particular solution is affected by cost-related rather than by illumination-related indicators. Analogous to the Indicator Bar, the specification of a spatial focus restricts the scores that are encoded in the lengths of the red bars to the objectives inside the focus. The weighted scores of the spatial objectives outside the focus are accumulated and appended as a gray bar. This supports the user in efficiently comparing the impact of a certain spatial region on the total scores of the solutions.

Clicking on rows in the Ranked Table selects the corresponding solutions in all views of LiteVis. Linking highlights the corresponding data in other abstract views such as Parallel Coordinate Plots or a bar chart displaying the distribution of simulated parameter values (see Fig. 6 d for an example). Selecting a solution furthermore loads the illumination and parameters for display in the spatial Simulation View. Selecting multiple solutions enables the comparison of their illumination (see Sec. 4.3 and Sec. 4.4). In the Ranked Table, selected solutions are marked by a red border.

In conclusion, the Simulation Ranking View is a key element of LiteVis for achieving all four design goals. Most importantly, the ability to express preferences in terms of weights separately for measurement surfaces as well as result indicators supports an explicit and reproducible decision making process (G4). Moreover, the Ranked Table enables an effective comparison of solutions regarding their weight-based desirability (G3) and enables an efficient selection of preferable solutions for a detailed inspection in the Simulation View (G2). Combining global indicators and local indicators in one view furthermore supports a holistic overview of relevant aspects (G1).

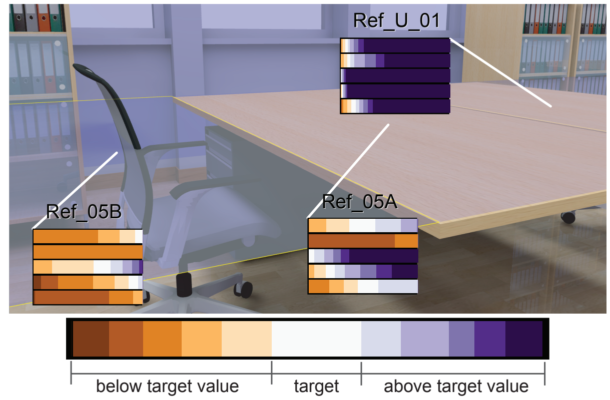

Fig. 5. Measurement Surface Annotations displaying illumination distribution histograms for five selected simulation runs. The annotations enable the comparison of abstracted illumination values in a spatial context across multiple measurement surfaces and parametrizations.

4.3 Measurement Surface Annotations

While the abstraction of a scene's illumination quality is necessary to enable the comparison of multiple solutions, the analysis and comparison of abstract information alone is not sufficient to form an educated opinion. For an in-depth analysis and comparison, the spatial context of local indicators is also important. The measurement surface annotations (MSA) in LiteVis form a compromise between abstraction and spatial context. The MSA enable the user to compare the distribution of illumination values on measurement surfaces from multiple simulation runs directly in the Simulation View (Fig. 5).

For a simulation run, this is achieved by encoding the distribution of illumination values of each surface in a binned histogram bar. For each run that is currently selected in an Analysis View, a histogram bar is created and positioned in a floating overlay that is visually linked to the corresponding surface. Rows of MSA therefore correspond to the current selection of simulated solutions. A histogram bar consists of eleven bins that are encoded with a brown/white/purple color scale. Depending on its illumination value, each texel on a surface is assigned to one of these bins. The central bin is colored white and represents the number of texels with illumination values that correspond to the respective target value of a surface. Values below/above the target are assigned to one of the brown/purple bins, depending on their distance from the target. This is the same color coding that we employ in the false color render mode (as displayed in Fig. 1). The histogram visually encodes several indicators for each surface and each selected run. Apart from the distribution of illumination values around the target, the colors of the bins inform the user about the minimum and maximum illumination values on each surface. The uniformity of illumination is encoded in the size of the bins. An entirely white bar would be most desirable for each surface. Depending on the use case, values above the target value might also be acceptable. Values below the target are never desired.

The display of MSA is triggered as soon as more than two parametrizations are selected in an Analysis View. Mouse selection and hover states on the histogram bars are linked with the corresponding solutions in the Analysis Views. In order to avoid visual clutter, MSA are displayed only for surfaces that are part of the current spatial focus in the Spatial Hierarchy.

4.4 False Color Comparison

Even though the MSA give the user more detailed information about the spatial distribution of illumination values in a scene and allow a comparison of solutions in a spatial context, they are still abstract representations. If the local objectives of two simulation runs have the same score, they might still differ in qualitative illumination aspects. For a complete assessment of a solution, the analysis of illumination values on the actual scene geometry is therefore indispensable. For a single run, this is typically done with false color renderings that facilitate the estimation of illumination values by emphasizing the negative or positive distance to a specified target value of each point in a scene. For the comparison of multiple solutions, this had to be done in a side-by-side manner. The False Color Comparison (FCC) mode (Fig. 6 h) gives the lighting designer the means for a detailed comparison of the illumination quality between two simulation runs directly within the simulation environment. The areas where the illumination values of each run correspond to the respective target value of a surface are rendered in a false color. Red areas indicate ideal illumination in the currently selected solution (conforming to the red selection color), while blue areas indicate ideal illumination in the currently hovered simulation run (conforming to the blue hover color). If both runs achieve ideal illumination on a texel, it is rendered in white. In compliance with the lighting designers' way of thinking, texels are considered as ideally illuminated, when they achieve at least the defined industry standard for the respective surface type. However, the acceptable range above the target value can be adjusted to adhere to special scenario requirements (e.g., when high illumination values should be avoided due to energy consumption reasons). The comparison is triggered when a single run is selected and a second run is hovered in an Analysis View or MSA.

The direct comparison that the FCC enables, has two advantages for the user. On the one side, users do not have to switch their focus between multiple views anymore when comparing two runs. On the other side, the FCC allows for more efficient feedback on newly created parametrizations by enabling an instant comparison of a new simulation run to a reference parametrization.

5 IMPLEMENTATION

LiteVis is built as an extension of multiple frameworks that communicate with each other via a network protocol. The Simulation View is based on an interactive global illumination lighting software that was developed with our collaborating domain experts in the scope of another project [18]. Its simulation kernel is based on a novel many-light simulation that allows the resulting illumination to be computed and visualized within a few seconds. This is achieved by using a GPU-based shadow mapping algorithm [33] for the visibility calculation concerning both direct and indirect light sources. The results are collected and stored adaptively in so-called light map textures that are mapped onto the scene geometry. After a scene modification, designers get immediate visual feedback, which continuously improves by converging to the physically accurate solution.

We extended the simulation with means to store, load, and export simulation data and added a web-overlay that handles rendering of and interactions with the MSA (implemented using the d3 toolkit [4]). The Analysis Views are part of a versatile visual analysis framework [25] in which we implemented the Simulation Ranking View. In order to enable the integration of the simulation and the analysis frameworks, both were extended with a network communication interface that manages the exchange of selection states, commands and simulation data.

6 USE CASE SCENARIO

In this section, we demonstrate the applicability of LiteVis based on a real-life use case scenario from our collaborating lighting designers.

6.1 Scenario Description

The simulation setup consists of a 3D office scene with 18 measurement surfaces. These surfaces are arranged into three semantic groups with different target values according to industry standards (Fig. 2):

- Task area (desks): 500 lx, 0.6 uniformity
- Close surroundings (work environment): 300 lx, 0.4 uniformity
- Background (walls, floor, ceiling): 100 lx, 0.1 uniformity

Further, the lighting designers supplied us with a set of 107 parametrizations that refer to variations of four different lighting scenarios. A scenario is characterized by the deployed luminary types that are grouped into *light groups* (LGs). LG1 is positioned over the desks. LG2 is positioned over the hallway of the office.

- Scenario A: LG1: pendulum, LG2: downlight
- Scenario B: LG1: double floor lamp, LG2: downlight
- Scenario C: LG1: single floor lamp, LG2: downlight
- Scenario D: built-in ceiling luminaries

We demonstrate how LiteVis supports the decision making process based on a global-to-local exploration of the supplied sample data.

6.2 Scenario Workflow

As starting point, the sample data base has been loaded in the Simulation Ranking View. The measurement surfaces in the Spatial Hierarchy are grouped by type and all surfaces are selected as spatial focus. Four local and two global objectives are specified (Fig. 6 a).

A typical first step in the workflow is the specification of per-measurement surface weights in the Spatial Hierarchy. Since the desks and work environments are of primary importance, the analyst decides to adjust the weights in the following way: Desks 60%, Work Environment 20%, Floor 10%, Walls 5%, and Ceiling 5% (Fig. 6 b). As a next step, the objectives for global and local indicators are adjusted in the Indicator Bar. To focus on runs with the best illumination quality regardless of budgetary constraints first, the analyst sets the weights of the non-spatial (financial) objectives to zero. Among the four local indicators of illumination quality, the analyst considers average illumination and uniformity on a surface as the most important metrics. Thus, she adjusts the weights of the four indicators as follows: Average 50%, Uniformity 30%, Maximum 15%, and Minimum 5% (Fig. 6 c). The Ranked Table now shows a ranking of the simulation runs with the best illumination quality according to the specified weighting of objectives (Fig. 6 c). As a next question, the analyst investigates which input parameters are involved in the highest ranking simulation runs. She selects the four top-ranking simulation runs. Then she parametrizes a linked bar chart (see Sec. 4.1) for each light group to inspect the luminary types used in the top-ranking runs (Fig. 6 d). This reveals that the LG1 luminaries are distributed across four different types, while all four runs are parametrized with the same luminary in LG2 (highlighted in red in Fig. 6 d).

The score in the ranking also shows that the local indicators are very distinctly distributed in each of the four simulation runs (Fig. 6 c). Run R4, for instance, has the best average illumination but a high cost regarding the minimal illumination value, which the analyst finds remarkable given the small weight assigned to this objective. She therefore decides to assess the illumination quality of these four runs directly in the Simulation View. First, she investigates the distribution of illumination values for R1 using the false color render mode (Fig. 6 e). While R1 exhibits a high overall score, the rendering exposes illumination values below the target value on the desk corners.

To compare the illumination distribution to the other top-ranking runs, the analyst simultaneously selects the first four runs in the Ranked Table. This triggers the display of MSA for the selected runs (Fig. 6 f). For R1, the histogram bar in the MSA confirms that parts of the illumination are below the target value, while R2, R3 and R4 lie entirely on or above the target on the desks. In fact, their distributions are very similar. Therefore, the analyst decides to assess their quality in a spatial context as well. By selecting the histogram bars in the MSA, R2-R4 are shown individually in the Simulation View (Fig. 6 g). R4 exhibits the most uniform illumination on the desks. However, as the analyst knows about the customer's preference for pendulum lights in this scenario, she does not want to discard R2 and R3 yet. To gain insight about eventual trade-offs between R2 and R3, she compares them directly in the 3D scene (Fig. 6 h). The FCC reveals that the ideally illuminated area of R3 (blue) lies in the center between two desks, while the ideal illumination in R2 (red) is more evenly distributed across the desk surface. She therefore considers R2 as the preferred solution in terms of illumination quality.

A satisfactory parametrization has been found. Additionally, the analyst sees multiple directions for further investigations. For instance, Scenarios C and D were not present in the top four runs. She could investigate, where the runs are situated in the ranking by selecting all associated runs based on their input parametrization using linked views (see Fig. 6 d). She could even create an individual ranking for

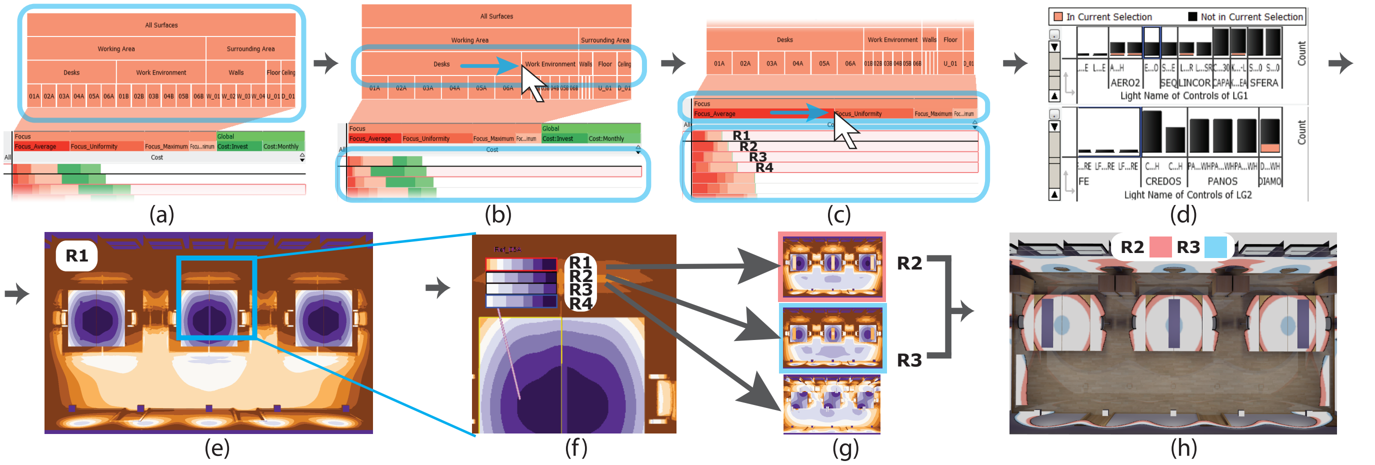

Fig. 6. The schematic overview of the workflow in our use case scenario showing an exemplary path for decision making. a) Initial state of spatial and non-spatial objective weights. b) Defining per-measurement surface weights in the Spatial Hierarchy. c) Changing the weights of global and local objectives in the Indicator Bar results in a ranking update. d) Highlighting the input parameters of the top four ranks (R1-R4). e) Assessing the spatial illumination distribution of R1 using false color rendering. f) Comparing the top four ranks in MSA. g) Assessing R2-R4 in the Simulation View. h) Triggering the FCC for a detailed comparison of R2 and R3 in a spatial context (white encodes ideal illumination in both R2 and R3).

each scenario by filtering runs, e.g., by the type of applied luminaries. Similarly, she could eliminate R1 from the ranking by filtering all parametrizations that have a minimum illumination value on desk surfaces below the specified target value. She could also investigate how the chosen parametrizations compare against others with respect to financial goals. For this, she would mark the top-ranked runs as selected in order to track their new ranks after adjusting the weights of the financial objectives. She could also directly compare multiple objective prioritizations by creating and comparing multiple Simulation Ranking Views, each one with a different objective weight distribution. She could check, how the cost of local indicators is distributed among the surfaces in the scene by setting, e.g., the working area as the spatial focus. Further, she could manipulate a run, e.g., by dimming the wattage and energy consumption, and see in which respect the resulting lower illumination values and runtime costs would influence the ranking.

7 Evaluation

7.1 Design Process

The design process of LiteVis can be described as an iterative cycle characterized by the three phases of (1) domain problem characterization, (2) data/operation abstraction, and (3) encoding/interaction technique design [21].

For characterizing our collaborators' domain problem and refining our understanding of their workflow, we performed semi-structured interviews during small meetings. As we observed the lighting designers' accustomed workflow, contextual inquiries [12] helped us to clarify tasks and challenges. The study of lighting design literature [8] helped to increase our knowledge of the associated data. We presented the collective findings from this phase in Sections 3.1 and 3.2. While we were able to understand the overall goal of the lighting designers well, the most challenging part of this phase was to get a deep understanding of the nuances of each task, such as: what exactly are the independent degrees of freedom for parametrizing a scene, which particular aspects and metrics in the evaluation of a scene are considered relevant? The resulting data abstraction can be found in the end of Section 3.1, the problem abstraction is the subject of Section 3.4.

Parallel prototyping [9] of hand-drawn sketches enabled us to refine our understanding of the requirements and tasks, and to discuss and validate our abstractions with the experts. The Simulation Ranking View was subject to the most design iterations. The basic requirement of providing an overview of local result indicators on different aggregation levels was addressed already in the first visual prototype. However, the approach for ranking a large number of simulation runs by these indicators changed throughout the design process. An early version featured a transposable grid, but turned out to be cluttered and did not provide a sufficient overview. In the next iteration, we investigated an icicle plot layout where each leaf corresponded to a parametrization. The hierarchy levels represented groupings of parametrizations by input or output parameters, and encoded aggregates of the contained result indicators. This structure was still too rigid due to the inherent hierarchical grouping by input or output parameters. Finally, we decoupled the leaves from the hierarchy. The leaves, i.e., parametrizations, were listed as stacked bars in the Ranked Table. The hierarchy was used to represent groups of measurement surfaces instead of parameters as a first version of the Spatial Hierarchy.

Once we settled on the design, we iteratively refined a software prototype and discussed the progress in monthly intervals with two to three domain experts. The overall design and implementation process took approximately 15 months. Important feedback during the design process concerned the automatic grouping of surfaces by type. Depending on the task, the experts considered other groupings more effective, e.g., by spatial proximity. We thus added the flexibility of grouping surfaces manually. Other feedback concerned the Spatial Hierarchy. The lighting designers had difficulties in understanding the relation between their interaction with surfaces in the Spatial Hierarchy and the resulting effect on the local objectives. The visual linking that we subsequently added alleviated this issue. Concerning the FCC, an initial version displayed a range around the target value for the involved parametrizations. However, the lighting designers stated that a range that includes values below the target was not helpful to them. The FCC therefore now only considers a range of values above the target by default. The range is adjustable as described in Section 4.4.

7.2 User Feedback

This section reports on user feedback collected during an evaluation workshop with three lighting design experts. Based on a protocol, we first gave an overview of the system and explained the individual components as well as their interplay. The demonstration was based on a dataset provided by our collaborators (Sec. 6). After two hours of introduction including questions and discussion, the domain experts used the system for another two hours including time for feedback and discussion. To document the feedback, participants completed a questionnaire with qualitative questions regarding the individual features.

The overall feedback of the interactive trial session was very positive. All experts agreed that analysis and comparison of parametrizations are key problems in the lighting design workflow. All of them considered LiteVis as an effective approach to increase efficiency and confidence in identifying relevant lighting parametrizations. The most appreciated component of the visual design included the Simulation Ranking View. The experts approved the design decisions of summarizing parametrizations by a combined score of aggregated result indicators as well as the interactive specification of measurement surface importance. Specifically, the ranking gives them "a never before envisioned overview of the quality of their simulation data". The specification of objectives concerning result indicators and measurement surface weights gives them "the necessary means to control and explore the massive amount of information". The head of lighting solutions support of our collaborator pointed out: "The system is very powerful and comprehensive. In particular the possibilities of an intuitive comparison of lighting solutions are a significant improvement as compared to other solutions that are currently deployed. It is exactly what we were looking for."

More critical feedback concerned the MSA in the spatial view. The experts commented that developing an intuition of the spatial illumination distribution based on the illumination histograms requires some time for familiarization. However, the lighting designers were generally able to understand the encoded illumination metrics and appreciated the feature for comparing the illumination distribution on individual measurement surfaces. In this regard, the FCC is considered as a useful complement to the abstract comparison of parametrizations in the Analysis Views and the MSA.

In summary, the discussions with the domain experts confirmed that we were able to meet our design goals. In fact, the lighting designers intend to incorporate LiteVis into their daily workflow. Discussions for a deployment of our prototype are already underway.

8 Discussion and Future Work

8.1 Goals

Summarizing the key components of LiteVis explains how our design corresponds to the goals stated in Section 3.3. We enable a holistic overview of all relevant features of the parameter space (G1), by linking the visualizations of the qualitative (Simulation View) and quantitative (Analysis Views) characteristics of input parametrizations and result indicators. Through this linked overview, it is not only easier for a designer to understand the results, but also to incorporate additional semantic information, customer requirements or prioritization concerns that are not included in the simulation itself. This linking also enables the tight integration of the scene parametrization with the evaluation of simulation results (G2) and thereby a more efficient feedback loop between these two tasks. The effective comparison of multiple solutions (G3) is enabled through the Analysis Views that can be flexibly adjusted to depict all sampled simulation runs based on their input and output values. This detailed comparison also enables a sensitivity analysis to assess the impact of parameter variations. Especially our novel Simulation Ranking View provides explicit and reproducible decision support (G4) by clearly indicating the distributions of spatial and non-spatial objective weights and the resulting ranking scores.

8.2 Scalability

The scalability of our tool can be analyzed according to different aspects. The number of global and local indicators in the Simulation Ranking View's Indicator Bar is limited by the horizontal screen space. Due to the aggregation of local indicators, their number is independent from the number of measurement surfaces. Local indicators could be further aggregated into categories if they would occupy more than the available horizontal space. The number of surfaces in the Spatial Hierarchy is limited by the available screen space, as well. However, since they are hierarchically groupable, they could be managed in a sort of zoomable icicle plot. In large office scenes, the surfaces could go into the hundreds. However, these scenes are typically evenly structured, i.e., they consist of repeating patterns of furniture and lighting constellations. Measurement surfaces can therefore be grouped without losing valuable information. The number of solutions displayed in the Simulation Ranking View is not an issue concerning scalability as the most interesting ones (according to the specified objectives) are always ranked at the top. If an overview of hundreds or thousands of simulation runs is desired, a table lens could be used to gain a synopsis of all runs. Hierarchical grouping by parameters could be used to divide a large sample base of simulation runs. One part of LiteVis that does not scale well are the MSA. The annotations become cluttered when displaying more than 20 simulation runs. However, feedback from lighting designers indicated that an in-depth comparison typically happens on a smaller set of runs.

8.3 General Applicability

The general applicability of concepts from LiteVis is given especially in our Simulation Ranking approach. It allows the specification of objectives on two orthogonal levels. This concept has general applicability in decision making, i.e., in cases where objectives can be shared by hierarchically superordinate entities. In disaster management, for instance, a simulation run can represent the results of a parametrizable flooding simulation. Instead of measurement surfaces, the Spatial Hierarchy could represent the evacuation prioritization of hierarchically groupable areas. Local indicators could represent different properties of these areas, such as the degree of damage, or repair cost.

8.4 Future Work

Our evaluation indicates that the current version of LiteVis already has clear benefits compared to current decision-making support-tools in lighting design. In our discussions with domain experts, we identified additional topics to be addressed in the scope of LiteVis:

- More user guidance: the Spatial Hierarchy could be used to supply the user with more information about the parameter space. The cells in the Spatial Hierarchy could be colored to encode how much they contribute to the aggregate of a local indicator. Another idea would be to use the cells to encode how many places in the ranking the current selection would change if the respective cell would be assigned an importance value of 100%.
- More control over spatial objectives: subdivision of measurement surfaces into sub-areas with individual spatial importances could help the lighting designer to fine-tune the specification of spatial optimality. On a desk, for instance, good illumination in the center is more important than at the corners. This is actually already possible but only in the setup phase (T0), i.e., when the measurement surfaces are defined.
- Analysis of illumination over time: especially in out-door illumination scenarios, the sun plays an important role. Allowing users to analyze parametrizations at different times of day and to incorporate day-time dependent dimming profiles into this process, would be an interesting application area to explore.

9 Conclusion

In this paper, we describe the LiteVis design study. LiteVis is a system aimed at assisting lighting designers in their decision making process in regard to finding an ideal lighting solution that adheres to competing qualitative and quantitative objectives. By integrating the spatial input domain of the parameter space with the abstract output domain of resulting simulation runs, we enable exploration and comparison of lighting parametrizations. We propose a novel ranking visualization that was inspired by multi-objective decision making solutions and takes into account the special intricacies of the domain's complex parameter space. A use case illustrates how LiteVis supports decision making in the context of a database of lighting solutions that would have overwhelmed lighting designers when using current state of the art lighting design solutions. We further gave a problem characterization of the application domain and reflected on our design process. Qualitative feedback from domain experts confirmed that we met our design goals and that lighting designers want to incorporate LiteVis into their daily workflow.

Acknowledgments

This work has been supported in the scope of the FWF-funded project P24597-N23 (VISAR) as well as by the Austrian Funding Agency (FFG) within the scope of the COMET K1 program. Thanks go to all project participants of Zumtobel Lighting, and to G. Hesina, S. Pajer, C. Arbesser, and T. Mühlbacher for valuable comments.

REFERENCES

[1] G. Andrienko and N. Andrienko. Constructing parallel coordinates plot for problem solving. In *1st International Symposium on Smart Graphics*, pages 9–14, 2001.

[2] M. Bagajewicz and E. Cabrera. Pareto optimal solutions visualization techniques for multiobjective design and upgrade of instrumentation networks. *Industrial & engineering chemistry research*, 42(21):5195–5203, 2003.

[3] M. Beham, W. Herzner, M. E. Gröller, and J. Kehrer. Cupid: Cluster-based exploration of geometry generators with parallel coordinates and radial trees. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):1693–1702, Nov. 2014.

[4] M. Bostock, V. Ogievetsky, and J. Heer. D^3 data-driven documents. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2301–2309, 2011.

[5] S. Bruckner and T. Möller. Result-driven exploration of simulation parameter spaces for visual effects design. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1468–1476, 2010.

[6] S. Chen, D. Amid, O. M. Shir, L. Limonad, D. Boaz, A. Anaby-Tavor, and T. Schreck. Self-organizing maps for multi-objective pareto frontiers. In *Visualization Symposium (PacificVis), 2013 IEEE Pacific*, pages 153–160. IEEE, 2013.

[7] D. Coffey, C.-L. Lin, A. G. Erdman, and D. F. Keefe. Design by dragging: An interface for creative forward and inverse design with simulation ensembles. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2783–2791, 2013.

[8] DIN Standard EN 12464-1: Light and lighting - Lighting of work places - Part 1: Indoor work places, Aug. 2011.

[9] S. P. Dow, A. Glassco, J. Kass, M. Schwarz, D. L. Schwartz, and S. R. Klemmer. Parallel Prototyping Leads to Better Design Results, More Divergence, and Increased Self-efficacy. *ACM Trans. Comput.-Hum. Interact.*, 17(4):pages 18:1–18:24, Dec. 2010.

[10] D. C. Glaser, R. Tan, J. Canny, and E. Y.-L. Do. Developing architectural lighting representations. In *Proceedings of the Ninth Annual IEEE Conference on Information Visualization*, INFOVIS'03, pages 241–248, Washington, DC, USA, 2003. IEEE Computer Society.

[11] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. Lineup: Visual analysis of multi-attribute rankings. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2277–2286, 2013.

[12] K. Holtzblatt and S. Jones. *Contextual Inquiry: A Participatory Technique for System Design*, pages 177–210. Lawrence Erlbaum Associates, Hillsdale, 1993.

[13] M. Koksalan, J. Wallenius, and S. Zionts. *Multiple criteria decision making: from early history to the 21st century*. World Scientific, 2011.

[14] P. Korhonen and J. Wallenius. Visualization in the multiple objective decision-making framework. In *Multiobjective optimization*, pages 195–212. Springer, 2008.

[15] Lighting Analysts, Inc. AGi32. http://www.agi32.com/. Accessed: 2015-03-30.

[16] W.-C. Lin, T.-S. Huang, T.-C. Ho, Y.-T. Chen, and J.-H. Chuang. Interactive lighting design with hierarchical light representation. *Comput. Graph. Forum*, 32(4):133–142, 2013.

[17] A. V. Lotov, V. A. Bushenkov, and G. K. Kamenev. *Interactive decision maps: Approximation and visualization of Pareto frontier*, volume 89. Springer Science & Business Media, 2004.

[18] C. Luksch, R. F. Tobler, R. Habel, M. Schwärzler, and M. Wimmer. Fast light-map computation with virtual polygon lights. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '13, pages 87–94, New York, NY, USA, 2013. ACM.

[19] M. J. McGuffin and J.-M. Robert. Quantifying the space-efficiency of 2d graphical representations of trees. *Information Visualization*, 9(2):115–140, 2010.

[20] K. Miettinen and M. M. Mäkelä. Synchronous approach in interactive multiobjective optimization. *European Journal of Operational Research*, 170(3):909–922, 2006.

[21] T. Munzner. A Nested Model for Visualization Design and Validation. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):921–928, 2009.

[22] M. Okabe, Y. Matsushita, L. Shen, and T. Igarashi. Illumination brush: Interactive design of all-frequency lighting. In *Proceedings of the 15th Pacific Conference on Computer Graphics and Applications*, PG '07, pages 171–180, Washington, DC, USA, 2007. IEEE Computer Society.

[23] F. Pellacini, F. Battaglia, R. K. Morley, and A. Finkelstein. Lighting with paint. *ACM Trans. Graph.*, 26(2), June 2007.

[24] F. Pellacini, P. Tole, and D. P. Greenberg. A user interface for interactive cinematic shadow design. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '02, pages 563–566, New York, NY, USA, 2002. ACM.

[25] H. Piringer, C. Tominski, P. Muigg, and W. Berger. A multi-threading architecture to support interactive visual exploration. *Visualization and Computer Graphics, IEEE Transactions on*, 15(6):1113–1120, 2009.

[26] P. Poulin and A. Fournier. Lights from highlights and shadows. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, I3D '92, pages 31–38, New York, NY, USA, 1992. ACM.

[27] Relux Informatik AG. DIAL GmbH. http://www.dial.de/DIAL/en/dialux/. Accessed: 2015-03-30.

[28] Relux Informatik AG. ReluxSuite. http://www.relux.info/. Accessed: 2015-03-30.

[29] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenberg. Painting with light. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '93, pages 143–146, New York, NY, USA, 1993. ACM.

[30] M. Schwarz and P. Wonka. Procedural design of exterior lighting for buildings with complex constraints. *ACM Trans. Graph.*, 33(5):166:1–166:16, Sept. 2014.

[31] M. Sedlmair, C. Heinzl, S. Bruckner, H. Piringer, and T. Möller. Visual parameter space analysis: A conceptual framework. *Visualization and Computer Graphics, IEEE Transactions on*, 20(12):2161–2170, 2014.

[32] A. Shesh and B. Chen. Crayon lighting: Sketch-guided illumination of models. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia*, GRAPHITE '07, pages 95–102, New York, NY, USA, 2007. ACM.

[33] L. Williams. Casting curved shadows on curved surfaces. In *Proceedings of the 5th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '78, pages 270–274, New York, NY, USA, 1978. ACM.

[34] Zumtobel Lighting GmbH. ecoCALC. http://www.zumtobel.com/com-en/ecoCALC.html. Accessed: 2015-03-30.