

Virtual Heliodon: Spatially Augmented Reality for Architectural Daylighting Design

Yu Sheng *

Theodore C. Yapo

Christopher Young

Barbara Cutler

Department of Computer Science, Rensselaer Polytechnic Institute

ABSTRACT

We present an application of interactive global illumination and spatially augmented reality to architectural daylight modeling that allows designers to explore alternative designs and new technologies for improving the sustainability of their buildings. Images of a model in the real world, captured by a camera above the scene, are processed to construct a virtual 3D model. To achieve interactive rendering rates, we use a hybrid rendering technique, leveraging radiosity to simulate the inter-reflectance between diffuse patches and shadow volumes to generate per-pixel direct illumination. The rendered images are then projected on the real model by four calibrated projectors to help users study the daylighting illumination. The virtual heliodon is a physical design environment in which multiple designers, a designer and a client, or a teacher and students can gather to experience animated visualizations of the natural illumination within a proposed design by controlling the time of day, season, and climate. Furthermore, participants may *interactively* redesign the geometry and materials of the space by manipulating physical design elements and see the updated lighting simulation.

Index Terms: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Radiosity, Virtual Reality H.5.1 [Information Interfaces and Presentation (HCI)]: Multimedia Information Systems—Artificial, augmented, and virtual realities

1 INTRODUCTION

The particular area of architectural design we explore is daylighting: the use of windows and reflective surfaces to allow natural light from the sun and sky to provide effective and interesting internal illumination (Figure 2). Appropriate daylighting strategies can reduce energy consumption for electric lighting and create more aesthetically interesting and comfortable architectural spaces. The overall aim of a successful daylighting design is to increase the amount of useful daylight in an architecturally satisfying way while avoiding the problematic aspects of natural illumination including the risk of glare or overheating.

A *heliodon* is a traditional daylighting analysis device in which a small-scale physical model (often $1/4'' = 1'$) is affixed to a platform and rotated relative to a fixed light source that represents the sun (Figure 1a). Alternatively, in more complex devices, the light source may be mechanically moved around a stationary platform. By studying the distribution of light within the model, the designer gains instantaneous and intuitive qualitative feedback on direct sun penetration and the corresponding indirect illumination.

Our new *virtual heliodon*, shown in Figure 1b, shares many of the features of the traditional heliodon and includes a number of important additional advantages. With a heliodon the user must awkwardly peer through windows (possibly blocking light) or place tiny video cameras within the model; in contrast, the virtual heliodon does not require a physical ceiling on the model, allowing




Figure 1: a) The depth of direct light beam penetration is easily ascertained with the traditional *heliodon*. The relative orientation of the model and light source is adjusted to represent different sun positions. b) Our *virtual heliodon* allows interactive redesign and enhanced visualization and exploration of the interior spaces.

easy inspection of the interior spaces. Furthermore, initial construction and revision of models for the virtual heliodon is faster than for heliodon models because the corners of the model need not precisely align. In the virtual heliodon we “fill in” cracks between walls that in a traditional heliodon would allow light to incorrectly leak into the model. Editing surface materials for a heliodon model is labor intensive and impossible if appropriate scale versions of the material are not available (e.g., Venetian blinds), but with our system surface and window materials can be changed digitally. Traditional heliodons only simulate direct illumination from the sun and do not automatically adjust the intensity with altitude. Our system performs a quantitatively-accurate simulation of illumination from both the sun and sky and also models climate variations.

We begin with an efficient, interactive, and accurate global illumination algorithm for daylighting. We draw from the body of literature on virtual and augmented reality to build a novel daylighting tool based on this rendering method. Our interactive immersive



Figure 2: Large, north-facing windows can be effectively used to illuminate spaces such as the classroom shown in the left image. However, glare (when the brightest part of the room is more than seven times brighter than the darkest part of the room) can reduce contrast and visibility. Direct sun far exceeds this range.

*e-mail: shengyu@cs.rpi.edu

daylighting design system can be used by novice or experienced designers who need not be experts in daylighting technology or advanced graphical simulations.

2 BACKGROUND AND RELATED WORK

Static, two-dimensional photographs or digital pre-renderings of a space are often insufficient for full appreciation of a scene. For example, only after visiting concert spaces with spectacular natural lighting did the clients and architects of Schermerhorn Symphony Hall in Nashville choose to incorporate daylighting into their project (a significant expense acoustically) [11]. The virtual heliodon is capable of providing similarly engaging visualizations of the dynamic nature of natural illumination.

2.1 Virtual Environments

For many years, virtual reality environments have been dominated by head-mounted displays and CAVE-style environments [8, 44]. Pertinent examples involving architectural and/or lighting design include navigation through hospital operating room designs [13], interior lighting design [31], and accurate global illumination from daylighting for car interiors [10].

Research in Spatially Augmented Reality (SAR) [4] addresses some of the physically-immersive criteria that we wish to leverage. Examples include the Office of the Future [37] and Everywhere Displays [47, 34], which use existing wall and desk surfaces to expand the area of the traditional computer interface. With Shader Lamps [38, 39], complex physical models are “animated”, for example, by projecting a video of a spinning wheel onto a stationary car model or a facade texture onto a physical architectural model [21]. The geometry of the physical objects is known *a priori* and the surfaces are assumed to be a uniform diffuse white material. The implementation of most SAR systems involves multiple automatically-calibrated cameras and projectors [36]. The projectors may also be used to generate patterns of structured light to reconstruct the geometry of a multi-planar display [1, 28]. Different mechanisms for interacting with these projected user interfaces have been proposed, prototyped, and evaluated for effectiveness. Examples include using a finger or laser pointer to indicate a position on the projected surface [34, 45] or gesturing with a handheld projector [3]. One or more cameras monitor the projection surface to detect and interpret the user’s actions.

An obvious problem when using projection onto real-world objects are occlusions and shadows cast when the user approaches the projection surface. This problem can be partially alleviated by positioning the projectors on the ceiling at shallow angles to the target surface. By using multiple overlapping projectors and automatic detection of the user and dynamic scene content, many of these shadows can be removed or reduced [2]. The angle and distance between the projector and surface must be taken into account when determining the focal length and relative pixel intensity. Color correction is necessary for multi-projector displays due to slight color mismatches between different projectors and projection onto surfaces with different material properties or textures [20]. Finally, when projecting onto multi-planar scenes or non-planar objects, there will be secondary scattering of the projected light. If the scene geometry and surface reflectance properties can be estimated, Raskar et al. [38] note that inverse global illumination could be applied to handle and undo this scattering (with some important restrictions), which is related to work on “common illumination” between synthetic and real objects [15, 12, 17, 7].

2.2 Interactive Global Illumination for Daylighting

The complexity of daylighting simulation, which accounts for slow rendering speeds and thus a decreased potential for interactive creative exploration of daylighting design, arises from two factors. First, natural illumination in the built environment is provided not




Figure 4: On a sunny day with the blinds up (left image), parts of this conference room table are too bright (> 1000 lux) for comfortable reading while other parts are too dark (< 500 lux). By lowering the blinds we can partially correct the situation by removing excess illumination from the center of the table. Altogether too often in practice, lighting engineers settle for the least creative and least efficient solution, selecting light fixtures for night or when the blinds are fully closed (right image).

only from direct parallel rays of sunlight, but also from omnidirectional illumination from the non-uniform, seasonally-varying sky hemisphere. The hemispherical distribution and relative intensity of the sky for different weather conditions (clear sky \rightarrow overcast sky) are calculated from standard models [6].

Second, surfaces receive illumination not only from direct light from the sun and sky, but also from *indirect* illumination that first reflects off one or more other surfaces in the scene. Real-time rendering systems based on the classic graphics rasterization pipeline (e.g., OpenGL and DirectX) are unable to accurately simulate indirect illumination. Physically-accurate patch-based radiosity [19] or Monte Carlo ray tracing [48] methods can simulate these effects, but require significantly more computational resources.

To facilitate interactive design, daylighting analysis tools must have adequate response time when the user edits the geometry or materials of the model. Decreasing the rendering time for global illumination is an active area of research [43, 33, 16, 9]; unfortunately, without explicit knowledge of these rendering algorithms, most architects are not able to appropriately prepare their models or tune parameters for this rendering software.

In most architectural scenes involving daylighting, light transfer due to diffuse reflection from surfaces dominates the indirect lighting. Additionally, hard shadows from the direct sun provide important visual cues that are necessary to understand the aesthetics of the space. Furthermore, the possibility of glare due to high contrast in the illumination values at the shadow boundaries must be considered. Per-pixel hard shadows greatly improve the perceived visual quality, but are usually not critical for computing accurate indirect illumination in diffuse-dominant scenes. Thus, we use a hybrid technique (Figure 3) that uses radiosity [19] to compute the diffuse reflection between faces on a coarse per-face basis, and replaces the direct illumination from the sun with per-pixel computations using multi-pass stencil shadow volumes [23].

2.3 State of the Art in Daylighting Simulation Software

A wide array of architectural CAD lighting design software tools are currently available. Tools for use in the early stage of design are either quantitative in output and highly restrictive in model complexity [41, 29] or limited to qualitative renderings of direct illumination only [5, 18]. In other words, these tools cannot be used to determine if the daylighting illumination levels in the current design will meet task-specific recommendations [40]. At the other end of the spectrum lie high-end rendering tools (e.g., Radiance [27]), which allow arbitrary model complexity at the expense of compu-




Figure 3: The a) classical radiosity solution does not capture hard-edged shadows. We factor the radiosity solution into b) the first bounce direct illumination and c) the indirect illumination by subtracting b) from a). d) Shadow volumes are used to generate per-pixel hard shadows. e) A hybrid radiosity/shadow volumes rendering is generated by adding c) and d).

tation time. These tools are not suitable for use at the early stage of design because the calculation and rendering times involved (from minutes to hours) restrict exploration of alternative designs and annual variations. By exhaustively pre-computing the lighting solution for simple fixed geometries, the various daylighting analysis programs which use Radiance as the rendering engine are able to allow limited interactive visualization and evaluation.

A critical gap in daylight modeling remains: for moderately complex geometries such as those found in schematic design, no program today enables climate-based, time-varying performance analysis capabilities in combination with an interactive, highly visual, and creativity-promoting design exploration process. Thus, daylighting simulation software is *very seldom used in academic or professional practice to inform design*. As a result, we are left with poorly designed spaces that exhibit uneven natural illumination conditions. These conditions are often corrected by blocking out nearly all of the daylight and relying primarily on electric lighting (Figure 4). Our project aims to fill this gap in the available daylighting design tools by providing an immersive and interactive system that allows iterative analysis and modification of the geometry during the design process.

3 OVERVIEW

We have built a spatially augmented reality heliodon system to complement modern desktop architectural daylighting design software tools. The user positions a set of small-scale physical walls within the workspace to “sketch” the 3D geometry of their design. Tangible walls are chosen to project daylighting simulation on instead of pure virtual reality solutions because they are low-cost, convenient, and easy for users to rebuild. Images captured by a camera mounted above the scene are processed to detect the wall positions. Gaps between the wall are filled to construct a closed 3D mesh. The daylighting solution in the virtual 3D building is computed with the hybrid radiosity/shadow volume global illumination rendering method (shown in Figure 3), accounting for illumination from both the sun and sky. This illumination solution is then displayed on the physical walls by four calibrated projectors.

The efficiency of the hybrid rendering method makes it possible for users to interact with daylighting simulations in our system. The user can explore the high-dimensional configuration space of the design by adjusting the position of wall modules to manipulate the geometry of the design. Through a wireless remote mouse, the user can vary the external conditions such as the sun position (time of day and day of the year) and weather conditions. The simulation results are recomputed and updated at interactive rates. In addition to visualizing the environment for a single point in time, the system allows time-lapse animations as the sun moves across the sky during the course of a day, throughout the season, or under different weather conditions. In addition, a simple 4” blue “north” arrow may be placed anywhere on the table in the camera’s field of view to indicate the overall orientation of the building on the site. When the orientation is changed, the daylighting condition in the building will be appropriately updated.




Figure 5: The user *sketches* a design with physical walls (left images), and the system interprets the sketch as a closed watertight architectural model (right images). Windows are automatically placed on red walls and cracks between the exterior walls are filled in. The ceiling polygons are omitted for visualization.

The novel contributions of our project include:

- Dynamic and immersive visualization of daylighting variation for different moments (time of day and day of the year),
- Spatially augmented reality for visualization and exploration of natural illumination by multiple simultaneous users, and
- Interactive redesign of the scene geometry with online recomputation of the resulting lighting simulation.

4 TECHNICAL ALGORITHMS & IMPLEMENTATION DETAILS

4.1 Physical Setup

Our table-top SAR system prototype (Figure 1b) centers around a standard 30” high, 42” diameter table. The surface of the table is covered with matte white museum board.

We have constructed a set of small-scale (1” = 1’) planar wall modules of three heights. All walls are made from lightweight, 1/2” thick matte white “foam core” and have small strips of museum board glued to the bottom of the wall for stability so they are easily balanced perpendicular to the table surface. The top of each wall is colored to identify its height and characteristics. The 10” wall modules are colored red and represent exterior walls that

contain windows. The 8" wall modules are colored green and represent exterior walls without windows. The 4" wall modules are colored yellow and are used for interior partitions, similar to the cubicle walls in many modern offices. We have found that these wall modules are sufficiently complex and expressive to represent a wide variety of office and residential programs and styles.

A simple frame of 4" and 2" diameter PVC pipe with aluminum and plywood bracing for rigidity is built around the table. The frame is 6' square at the base and 8' tall allowing for easy movement around all sides of the table. A camera is mounted approximately 5' above the center of the table and oriented downward to capture the table and walls. Four projectors are mounted at the upper corners of the frame and oriented to project to the center of the table at an approximately 45° downward angle. The distance from the projectors to the table center is approximately 5'. The frame is not necessary if the cameras and projectors are mounted from the ceiling instead.

Six halogen lights designed for "under-cabinet" use are mounted in a ring surrounding the camera for illuminating the scene during image capture. While this arrangement provides nearly ideal images with uniform illumination and no visible shadows, we found in practice that our image processing algorithms (Section 4.2) are sufficiently robust to allow use of available room lighting.

4.2 Image Processing: Wall Detection

We have colored the tops of the walls to indicate their height, allowing them to be tracked by a stationary, calibrated camera [47]. Additional labeling conventions will be added in the future to allow fine-grained control of window dimensions, position, height, and glazing material. We considered using ARToolkit [25] for tracking, but found the required large marker size to be problematic. Furthermore, since our walls are constrained with respect to distance from the camera, the total error can be reduced by essentially tracking the slender colored rectangles in 2D.

The camera provides an overhead view of the model. As shown in Figure 7a, the colored tops of the walls are clearly visible from this viewpoint. Wall models are assumed to be vertical, planar, and of the exact, fixed height specified by their color. To detect the walls within the 2D image, pixel RGB values are transformed into an HSV representation, and a nearest neighbor classifier is applied to the hue and saturation components. The resulting image of color class labels is enhanced using a series of morphological open/close operations to remove outlying noise pixels and close small holes in detected walls. We use connected components analysis to label collections of pixels as individual walls. The colored wall top stops 1/2" from each end of the wall, allowing for easy separation of physically touching walls during the connected components step; the length of the virtual walls are subsequently extended to compensate. The edge pixels of each wall are then found, and robust line fits for each edge are computed using RANSAC [14]. An analytic solution to the intersections of the fit boundary lines of each wall allows for an estimate of wall corners to sub-pixel precision. Using the calibrated camera projection matrix, P , these 2D points are back-projected to rays in 3D world coordinates [22], and the known height of the wall (inferred from its color) allows calculation of the 3D positions of the top wall corners. Dropping virtual "plumb lines" from these corners to the known table surface height allows full 3D models of the walls to be reconstructed (Figure 6).

4.3 Model Construction and Remeshing

The 3D wall geometry sketched by the user and detected in the previous image processing stage is not a "watertight" model. The walls may have gaps between them and the corners of the physical walls may overlap and stick beyond their natural intersection point. Furthermore, neither ceiling nor windows is explicitly included in the sketch and must be inferred. Thus, the geometry must be pro-




Figure 6: Wall detection geometry. 2D image processing algorithms are used to detect the tops of each wall, then the camera calibration parameters are used to back-project rays from the camera to the known wall heights, yielding 3D positions of the walls.

cessed to build a closed 3D triangle mesh model that is appropriate for computing a radiosity global illumination solution. For our current implementation, we limit the geometry to *star shaped* exterior wall room footprints, which means that there exists a point within the volume of the room that can "see" every face. This geometry is more general than a simple convex hull that is not always a suitable assumption for interior architectural design.

First, the geometry is separated into exterior and interior wall modules. We orient the exterior walls by determining which side faces "inward". This is done by first computing the center of mass of the scene geometry (alternatively, we could use the zero point of the world coordinate system, the center of the table) and assume that this point meets the star-shaped property. The normal of each wall with respect to this point is computed and the interior surface is marked. Next, the walls are sorted in clockwise order by the angle the center of mass of each wall makes with the center of mass of the entire wall arrangement. Then, the interior face of each wall is intersected with its neighbors. If the intersection falls within the wall segment, the wall is clipped to the intersection. If the intersection is beyond the edge of the wall, extra *fill-in* polygons are created to fill the gap. A ceiling is inferred for the model at the height of 8' and a triangle fan is added to construct a manifold closed model. In the future, we would like to explore ways for the user to express variations in ceiling height which can have an important impact on daylighting design. Finally, interior cubicle walls are added to the model. It is important that the floor polygons correctly and manifoldly mesh at the edges of the interior partitions, rather than allow the walls to float above the floor. If the joint is not handled properly, indirect illumination from the radiosity solution might incorrectly leak under the partitions.

Once the closed model is constructed, it is remeshed with a combination of edge split, edge flip, edge collapse, and move vertex operations [24] to arrive at a triangle mesh with approximately 1500 polygons, an appropriate number for interactive radiosity simulations.

4.4 Projection of Daylighting Solution

The daylighting solution is computed on the closed model from the previous section, but must be displayed on the physical non-closed model. Therefore, these gaps and unused portions of the walls are specifically labeled (Figure 7). The patches corresponding to the *projection surfaces*, the interpreted *fill-in surfaces* necessary to make a closed model, and the additional unused *physical occlusions* are identified and processed to ensure that the model is correctly rendered. The fill-in surfaces include the ceiling and gaps between the walls. Although not displayed in the final rendering, these patches are part of the closed building model and thus included in the computation of the radiosity solution (Figure 7d). Portions of the physical walls that are not included as projection surfaces because they are taller than the ceiling or because they are




Figure 7: a) The image captured by the camera is processed to determine the height, length, and position of b) the 3D physical walls. A closed 3D model of the intended room geometry is constructed and meshed with polygons that are roughly uniform in area. These polygons are classified as c) physical projection surfaces and d) fill-in polygons. Additionally, e) the full extent of each physical wall is also represented in order to correctly account for occlusions between the projectors and physical projection surfaces.




Figure 8: To create a smooth transition between the coverage of three projectors we use intensity blending.

overlapping and thus clipped from the closed model are labeled as physical occlusions (Figure 7e). They will not be taken into consideration for the radiosity computation, but they must be drawn black in the final rendering because they may cause real physical occlusions between the projectors and projection surfaces.

A hybrid technique that combines per-patch radiosity computation and per-pixel shadow volumes (Figure 3) is used to interactively render images that simulate the daylighting conditions in the room. A perspective projection matrix, which represents the projector’s intrinsic and extrinsic parameters, is used to render an image of the simulation results for each of the four projectors. Since the radiosity global illumination rendering method is view-independent, the lighting solution is calculated once and the illumination is correct when viewed from any angle [38].

The projectors have been positioned at an approximately 45° downward angle, which gives good coverage and visibility of the scene; however, there are still significant occlusions between the walls and projection surfaces that must be considered to produce correct and consistent projections. The display of each triangle is assigned to one or more projectors for display. When multiple projectors are available to cover a certain patch, they are appropriately blended to avoid bright overlap regions. This is done by intensity blending with attenuating intensities in these regions [4]. We adopt the following strategy to smoothly illuminate the physical walls and floor of the model. For each vertex, a projector is set as a candidate if it has a clear line of sight to the vertex (i.e., not occluded by either projection surfaces or physical occlusions) and the dot product of the vertex normal and the vector from the vertex to the projector is greater than zero. All candidate projectors are sorted by decreasing order of the dot product. If a projector is also a candidate among all of its one-ring vertices, it may be used for intensity blending across all triangles at that vertex. Thus we select the first candidate projector that is also a candidate at all one-ring vertices as the best projector for display at that vertex. However, some vertices may fail to find such a projector. This problem can be resolved by locally refining the mesh. In practice this occurs infrequently, and for simplicity in our prototype implementation, we select the projector that is a candidate at the largest number of its neighboring vertices. Figure 8 illustrates this blending algorithm for a set of floor patches illuminated by three projectors.

Treating the projector lamp as a point light source, the radiance reflected from the physical surface is related to the diffuse reflectance ρ of the surface, distance r between the surface and projector and the angle θ subtended by the surface normal and the vector from the surface to the projector. In order to achieve the same radiance in physical world as in the virtual environment, radiance adjustment [38] is used. We apply the equation

$$I = E \frac{r^2}{\rho \cos \theta}$$

to each vertex that is rendered (I is the radiance value after adjustment, and E is the value from radiosity solution).

Our radiosity solution produces linear radiance values. Since the projectors apply gamma correction ($\gamma = 2.2$), we must apply the inverse gamma transformation to pixel intensities before display. We note that although we have used 8-bit pixel values throughout, approximately 14 bits would be required to accurately encode the range of pixel intensities in a linear space. As a result, our current implementation shows banding artifacts in the darkest regions.

4.5 Camera and Projector Calibration

Calibration of our system comprises two individual components. The camera must be geometrically calibrated in order to produce accurate 3D models of the scene, and each projector must be geometrically calibrated to project correctly onto the wall surfaces.

The overhead camera is calibrated using Zhang’s algorithm [49]. We created a calibration target consisting of 212 black and white corner marks on a white background (Figure 9a). The target is printed directly on a 40” square sheet of “Gatorboard”, which provides a flat and rigid surface. Extraction of the corners from camera images is facilitated by the addition of four colored disks to the outer edge of the target. These disks are easily detected using a color-based classifier, and the centroids of the colored regions are used to estimate a planar homography between the known target model and each captured image. Using this homography, we are able to efficiently scan the image for the grid of calibration corners. Isolated corners are used for calibration points since they are immune to bias induced by perspective and radial lens distortion [32]. To extract the corner positions with sub-pixel accuracy, we fit a quadratic function to a blurred version of the image region surrounding each corner, and compute the location of the extremal points of the resulting polynomial [30]. Approximately 40 images of the calibration target in various poses are taken, roughly filling the active 3D volume of the model space. Zhang’s algorithm is used to estimate the intrinsic camera parameter matrix, K , from these images. Finally, a single image is taken of the calibration target lying flat on the table surface. The extrinsic parameters estimated for this image are used to establish the world coordinates for the system. This fixes the table surface as the x - z plane, with the y -axis normal to the table top. Once the images are captured, the calibration process is automatic and takes less than a minute.

The four projectors are also calibrated with a traditional camera model, using Tsai’s method [46]. Tsai’s algorithm requires a set of




Figure 9: Calibration targets for camera and projector. On the left, a processed image of the camera calibration target; on the right, a projector calibration image. Detected calibration points are marked with yellow crosses.




Figure 10: Visualizing the active model volume. Rays are shown from the camera (white) and four projectors (red, green, cyan, magenta) projected into the modeling volume. Calibration points are drawn as small spheres.

correspondences between projector image pixels and 3D points in our world coordinate system. We generate these correspondences by projecting the target pattern shown in Figure 9b onto a set of eight horizontal planes spaced at 24.6mm intervals within the model space, approximately spanning the active volume (Figure 10). The matte white reverse side of the camera calibration target is used as a projection screen. We extract the 2D position of the center of each corner in the camera image with the same sub-pixel corner extraction algorithm used for camera calibration. Back-projection of the 2D camera location of the corner onto the known plane height above the table (Section 4.2) provides the necessary 3D coordinates. This process produces a set of 1,504 point correspondences for each projector. After calibration, we achieve a reprojection error of $0.4 \pm 0.4\text{mm}$ (1σ).

4.6 System Integration and Implementation Details

In our setup, we employ a Point Grey Flea2 1394b camera coupled to an 8mm $f/1.4$ lens to capture the images at a resolution of 1392×1032 pixels. The projectors are DLP-based Optoma EP727 units, with an image resolution of 1024×768 . We run the projectors on two dual-head video cards, and run a standard monitor for the user interface on a third card. The computer runs Linux on an Intel Core 2 Quad Q9450 processor.

5 RESULTS

The current system runs sufficiently fast for interactive daylighting visualization during iterative architectural design. Using unoptimized code, it takes approximately 11 seconds to generate the projection images for a model with 6 walls and 2 windows. Of this, approximately 3 seconds are used to process the captured image with our prototype code, and then additional 2.5 seconds are used to remesh the wall coordinates into a closed triangle mesh. After remeshing, it takes approximately 3 seconds to calculate form factors for the mesh with about 1500 triangles and compute an initial lighting solution. Because the image capture, remeshing, and form factor calculations need only be performed once for a particular wall geometry, additional lighting solutions for this geometry (changing the time of day, day of the year, weather conditions, or site orientation) can be done in about 0.6 seconds. This facilitates interactive time-lapse animations of the daylighting solution which is shown in our companion video. Figure 11 shows still frames from an animation of the daily variation in lighting conditions for a fixed model geometry and Figure 12 shows several images from an iterative design session.

Preliminary feedback from architecture students about the system has been positive. We plan to conduct formal user studies with our virtual heliodon to compare it with both the traditional heliodon and state-of-the-art daylighting analysis software packages. We believe our system will result in an improved understanding of the dynamic nature of daylighting and increased awareness of and attention to sustainable architecture design.

6 LIMITATIONS AND FUTURE WORK

Although our current system with four overhead projectors can be used to model many room configurations, tall, narrow room designs can result in areas occluded from all four projectors. Although this results in a presentation that does not reflect the true lighting solution, such shadow areas are easily distinguished visually and we have not found them to be overly distracting.

Accurate representation of the full dynamic range of daylight is problematic. Projectors have a limited range of intensity: our four projectors together can produce a maximum 12,000 lux on the table surface; direct sunlight is $10\times$ brighter. Even using the brightest projector commercially available (15,000 lumens [42]), projection onto a $3' \times 4'$ screen would result in only one ninth the brightness of direct sunlight. To compensate for the limited dynamic range we will extend tone mapping research [26] to immersive displays. As an alternative, we plan to explore false-color representations which can indicate overall lighting levels, or highlight areas that experience significant glare or inadequate illumination during some parts of the day.

While our current wall detection and tracking method has proven relatively robust, the algorithm may become less reliable as many more colors are used to code different wall heights or styles. Accordingly, we plan to explore other options for wall labels: bar-code sequences on wall tops, magnetic position tracking system [35], visible or infrared LEDs with unique colors or on/off sequences, and active systems where our projectors display structured light patterns captured with the camera or synchronized sensors that uniquely determine their 3D position [28]. Cost, robustness, ease-of-use, accuracy, speed, and how to minimize visual distractions for the designer will all be considered.

To correct for unintentional scattering of projected illumination between physical wall surfaces [38], we will perform inverse global illumination and color/intensity compensation [20] to determine the color that should be projected onto each wall so that the resulting intensity matches the software simulation. A solution to the general problem does not necessarily exist – we cannot emulate a scene with a white wall adjacent to a fully-absorptive black floor because we cannot prevent scattered light from striking and reflect-




Figure 11: Top row: lighting simulation for a room design for June 21 (summer solstice) at 2 hour intervals between 7AM and 5PM. Middle row: same simulation for March/September 21 (spring & autumnal equinox). Bottom row: same simulation for December 21 (winter solstice).




Figure 12: Images from an interactive design session beginning with a) the initial design for a simple room configuration of 6 walls. b) The user adjusts the orientation of the entire building on the site by rotating the North indicator arrow. c) The placement of windows within the design is changed by swapping a pair of walls. d) The geometry is adjusted to fine tune the shadow placement. e) An interior wall is added to the design. All images in this figure show the daylighting solution for March 21st at 1pm.

ing off the real-world floor. However, several characteristics of our application will allow us to compensate for most of the scattering; the diffuse surface assumption holds for most architectural scenes; neutral mid-range tones are used in most designs; gaps between the wall panels reduce surface inter-reflection; and the use of tone mapping will effectively reduce the desired contrast of the projections.

The mechanisms for user interaction can be further enhanced. We envision extending the system with physical control elements or gestures similar to the north arrow to adjust other aspects of the geometry or simulation. For example a cloud shape could be used to indicate an overcast sky condition. Or a simple icon placed next to a wall could be used to change the material reflectance properties of the surface or window.

7 CONCLUSIONS

We have presented the *virtual heliodon*, a novel application of spatially augmented reality for architectural daylighting design. This small-scale, low-cost physical design environment is practical for academic architectural studios or professional design firms and suitable for installation in a multi-purpose space such as a conference room. Our system has a number of advantages over the traditional heliodon including ease of visualizing the interior lighting conditions and ease of editing the model geometry. This new design tool supports enhanced communication between client and architect and provides a platform for education in sustainable architectural design practice. The system can improve the architect's ability to make effective use of daylighting and reduce the need for supplemental electric lighting, reducing consumption of non-renewable energy resources. The same framework can be applied to other architec-

tural design problems that require complex physical simulation and visualization, including passive solar heating and cooling, acoustics, aerodynamic building envelopes, and structural analysis.

ACKNOWLEDGEMENTS

The authors wish to thank Marilyne Andersen, Melissa Schroyer and the students in the computer vision and computer graphics research groups at RPI. We also thank the anonymous reviewers for their helpful comments. This work was supported in part by grants from IBM and NSF CMMI-0841319.

REFERENCES

- [1] M. Ashdown, M. Flagg, R. Sukthankar, and J. M. Rehg. A flexible projector-camera system for multi-planar displays. In *2004 Conference on Computer Vision and Pattern Recognition (CVPR 2004)*, pages 165–172, June 2004.
- [2] S. Audet and J. R. Cooperstock. Shadow removal in front projection environments using object tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, June 2007.
- [3] P. Beardsley, J. V. Baar, R. Raskar, and C. Forlines. Interaction using a handheld projector. *IEEE Computer Graphics and Applications*, 25(1):39–43, January 2005.
- [4] O. Bimber and R. Raskar. *Spatial Augmented Reality: Merging Real and Virtual Worlds*. A K Peters, Ltd., 2005.
- [5] S. Bund and E. Y.-L. Do. Spot! fetch light interactive navigable 3d visualization of direct sunlight. *Automation in Construction*, 14:181–188, 2005.
- [6] CIE International Commission on Illumination. Spatial Distribution of Daylight - Luminance Distributions of Various Reference Skies, 1994. Technical Report CIE 110 - 1994.

- [7] O. Cossairt, S. Nayar, and R. Ramamoorthi. Light-field transfer: Global illumination between real and synthetic objects. *ACM Transactions on Graphics*, 27(3), July 2008.
- [8] C. Cruz-Neira, D. Sandin, T. DeFanti, R. Kenyon, and J. Hart. The cave: Audio visual experience automatic virtual environment. *Communications of the ACM*, 35(6):65–72, June 1992.
- [9] C. Dachsbacher, M. Stamminger, G. Drettakis, and F. Durand. Implicit visibility and antiradiance for interactive global illumination. *ACM Transactions on Graphics*, 26(3):61:1–61:10, July 2007.
- [10] K. Dmitriev, T. Annen, G. Krawczyk, K. Myszkowski, and H. Seidel. A cave system for interactive modeling of global illumination in car interior. In *ACM Symposium on Virtual Reality Software and Technology*, pages 137–145, 2004.
- [11] E. Donoff. Schermerhorn symphony center. *Architectural Lighting Magazine*, March 2007. <http://www.archlighting.com/industry-news.asp?articleID=460542§ionID=1331>.
- [12] G. Drettakis, L. Robert, and S. Bougnoux. Interactive common illumination for computer augmented reality. In *Eurographics Rendering Workshop 1997*, pages 45–56, June 1997.
- [13] P. Dunston, J. McGlothlin, and L. Arns. An immersive virtual reality mock-up for design review of hospital patient rooms. In *Conference on Construction Applications of Virtual Reality*, 2007.
- [14] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, 1981.
- [15] A. Fournier, A. S. Gunawan, and C. Romanzin. Common illumination between real and computer generated scenes. In *Graphics Interface*, pages 254–263, 1993.
- [16] P. Gautron, J. Krivánek, K. Bouatouch, and S. Pattanaik. Radiance cache splatting: A gpu-friendly global illumination algorithm. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, pages 55–64, June 2005.
- [17] S. Gibson and A. Murta. Interactive rendering with real-world illumination. In *Eurographics Workshop on Rendering (Rendering Techniques)*, June 2000.
- [18] Google. SketchUp: 3D modeling software, 2008. <http://www.sketchup.com>.
- [19] C. M. Goral, K. E. Torrance, D. P. Greenberg, and B. Battaile. Modelling the interaction of light between diffuse surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 84)*, volume 18(3), pages 213–222, July 1984.
- [20] A. Grundhofer and O. Bimber. Real-time adaptive radiometric compensation. *IEEE Transactions Visualization and Computer Graphics*, 14(1):97–108, January-February 2008.
- [21] S. H. Han, J. H. Kim, T. S. Yun, and D. H. Lee. Extensible interface using projector-based augmentation. In *International Conference on Computer Graphics & Virtual Reality (CGVR)*, pages 132–137, 2006.
- [22] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.
- [23] T. Heidmann. Real shadows, real time. *Iris Universe*, 18:28–31, 1991. Silicon Graphics, Inc.
- [24] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. In *Proceedings of SIGGRAPH 93*, Computer Graphics Proceedings, Annual Conference Series, pages 19–26, Aug. 1993.
- [25] H. Kato and M. Billinghurst. Marker tracking and hmd calibration for a video-based augmented reality conferencing system. In *Proceedings of the 2nd International Workshop on Augmented Reality (IWAR 99)*, San Francisco, USA, Oct. 1999.
- [26] G. W. Larson, H. Rushmeier, and C. Piatko. A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, 1997.
- [27] G. W. Larson and R. Shakespeare. *Rendering with Radiance - The Art and Science of Lighting Visualization*. Morgan Kaufmann Publishers, Inc., 1998.
- [28] J. C. Lee, P. H. Dietz, D. Maynes-aminzade, and S. E. Hudson. Automatic projector calibration with embedded light sensors. In *ACM Symposium on User Interface Software and Technology (UIST)*, 2004.
- [29] M. A. Lehar and L. R. Glicksman. Rapid algorithm for modeling daylight distributions in office buildings. *Building and Environment*, 42:2908–2919, 2007.
- [30] L. Lucchese and S. Mitra. Using saddle points for subpixel feature detection in camera calibration targets. *Circuits and Systems, 2002. APCCAS '02. 2002 Asia-Pacific Conference on*, 2:191–195 vol.2, 2002.
- [31] C. Mackie, J. Cowden, D. Bowman, and W. Thabet. Desktop and immersive tools for residential home design. In *Conference on Construction Applications of Virtual Reality*, 2004.
- [32] J. Mallon and P. F. Whelan. Which pattern? biasing aspects of planar calibration patterns and detection methods. *Pattern Recognition Letters*, 28(8):921–930, June 2007.
- [33] M. Nijasure, S. Pattanaik, and V. Goel. Interactive global illumination in dynamic environments using commodity graphics hardware. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, page 450, Washington, DC, USA, 2003. IEEE Computer Society.
- [34] C. Pinhanez. The everywhere displays projector: A device to create ubiquitous graphical interfaces. In *Ubiquitous Computing (Ubicomp)*, September 2001.
- [35] F. Raab, E. Blood, T. Steiner, and H. Jones. Magnetic position and orientation tracking system. *Aerospace and Electronic Systems, IEEE Transactions on*, AES-15(5):709–718, Sept. 1979.
- [36] R. Raskar and P. Beardsley. A self-correcting projector. In *2001 Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, pages 504–508, Dec. 2001.
- [37] R. Raskar, G. Welch, M. Cutts, A. Lake, L. Stesin, and H. Fuchs. The office of the future: A unified approach to image-based modeling and spatially immersive displays. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 179–188, July 1998.
- [38] R. Raskar, G. Welch, K.-L. Low, and D. Bandyopadhyay. Shader lamps: Animating real objects with image-based illumination. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, pages 89–102, June 2001.
- [39] R. Raskar, R. Ziegler, and T. Willwacher. Cartoon dioramas in motion: Apparent motion effects on real objects with image-based illumination. In *International Symposium on Nonphotorealistic Animation and Rendering (NPAR)*, June 2002.
- [40] M. Rea, editor. *The IESNA Lighting handbook - Reference & Application*. Illuminating Engineering Society of North America, 9th edition, 2000.
- [41] C. Reinhart, D. Bourgeois, F. Dubrous, A. Laouadi, P. Lopez, and O. Stelescu. Daylight 1-2-3 - A state-of-the-art daylighting/energy analysis software for initial design investigations. *Building Simulation*, 2007.
- [42] Sanyo. XGA Multimedia Projector PLC-XF47, 2007. <http://www.sanyo-lcdp.com/english/product/xf47/xf47.html>.
- [43] P.-P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3):527–536, July 2002.
- [44] O. Staadt, M. Gross, A. Kunz, and M. Meier. The Blue-C: Integrating real humans into a networked immersive environment. In *Proc. ACM Collaborative Virtual Environments 2000*, pages 201–202, 2000.
- [45] R. Sukthankar, R. G. Stockton, and M. D. Mullin. Smarter presentations: Exploiting homography in camera-projector systems. In *Proceedings of International Conference on Computer Vision*, 2001.
- [46] R. Y. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision metrology using off-the-shelf tv cameras and lenses. *IEEE Journal of Robotics and Automation*, RA-3(4):323–344, August 1987.
- [47] J. Underkoffler, B. Ullmer, and H. Ishii. Emancipated pixels: Real-world graphics in the luminous room. In *ACM SIGGRAPH*, August 1999.
- [48] G. J. Ward, F. M. Rubinstein, and R. D. Clear. A ray tracing solution for diffuse interreflection. In *SIGGRAPH '88: Proceedings of the 15th annual conference on Computer graphics and interactive techniques*, pages 85–92, New York, NY, USA, 1988. ACM Press.
- [49] Z. Zhang. A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22:1330–1334, 2000.