Radiometric Compensation through Inverse Light Transport

Gordon Wetzstein[1,2]
[1] University of British Columbia

Oliver Bimber[2]
[2] Bauhaus-Universität Weimar

Abstract

Radiometric compensation techniques allow seamless projections onto complex everyday surfaces. Implemented with projector-camera systems they support the presentation of visual content in situations where projection-optimized screens are not available or not desired - as in museums, historic sites, air-plane cabins, or stage performances. We propose a novel approach that employs the full light transport between a projector and a camera to account for many illumination aspects, such as interreflections, refractions and defocus. Precomputing the inverse light transport in combination with an efficient implementation on the GPU makes the real-time compensation of captured local and global light modulations possible.

1 Introduction

Video projectors have evolved tremendously in the last decade. Reduced costs and increasing capabilities (e.g. spatial resolution, brightness, dynamic range, throw-ratio) have led to widespread applications in home entertainment and visualization. Recently, numerous projector-camera approaches that enable seamless projections onto complex everyday surfaces have been proposed. In general this is referred to as *radiometric compensation*. These methods correct the projected images for geometrical distortions, defocus or color and intensity blending caused by the underlying surface. Eventually, the final images appear as being projected onto a planar white canvas - even though this is not the case.

Projector-camera-based image correction techniques and especially radiometric compensation have proven to be a useful tool for many real-world applications. These include, for instance, on-site architectural visualization and augmentation of artifacts in museums. In several installations we have augmented cultural heritage sites, buildings, stage settings or fossils with additional information (i.e. computer generated 3D reconstructions)

that allowed a broad audience to better understand details about these exhibits. We believe that such approaches help to establish the practical use of computer graphics techniques in many different fields.
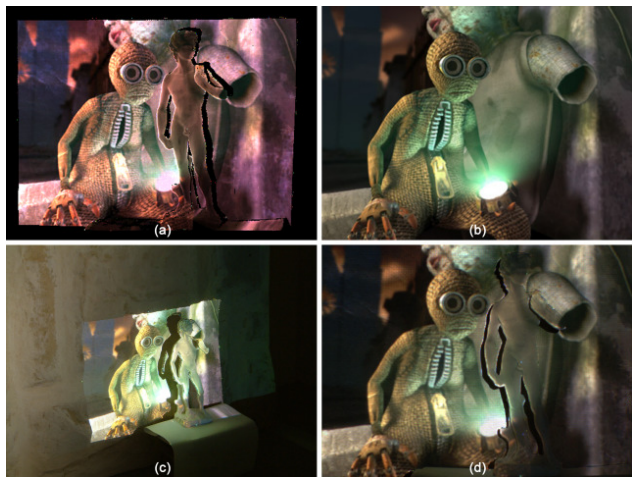

Figure 1. Synthesis of a projector illumination pattern (a) that results in a desired image (b) when projected onto a scene (c) and observed from a predefined viewpoint (d).

In most of these applications one or multiple projectors are mounted at different locations and illuminate complex surfaces. Usually, a calibration camera is placed at a position where the viewers are expected to be situated. This can be different from the projectors' locations and guarantees an optimal geometry correction for a limited sweet-spot or area. Once a geometric correspondence and surface reflectance properties have been estimated automatically, the camera can be removed and a real-time compensation is performed.

Previously proposed radiometric compensation techniques assume a simple geometric relation, represented by a one-to-one mapping between projector and camera pixels that can be automatically derived using structured light projections (range scanning) [18, 3] or coaxial projector-camera alignments [9]. In reality, the

light of an illuminated display element often bounces back and forth several times at different areas on the surface, before it eventually reaches the imaging sensor. Interreflection, refraction, scattering and other global illumination effects are consequently discarded in conventional methods. Often, real-world materials in non-laboratory settings make it very difficult, if not impossible, to use range scanning techniques because global effects are too dominant.

Although radiometric compensation is a well studied subject, our contribution is a novel generalized framework that unifies many of the previously proposed compensation approaches and extends these to account for global illumination. Employing the full light transport to perform an image-based radiometric compensation has not, to our knowledge, been achieved before. We corroborate our theory with many different sample scenes that exemplify various illumination effects. Real-time frame rates for image compensation are achieved with a customized clustering scheme in combination with an optimized GPU implementation. The proposed method does not claim to produce superior quality for simple scenes, but in contrast to previous approaches can handle arbitrarily complex objects. Limitations are set by the physical laws of light transport and the constraints of the applied hardware.

2 Related Work

Geometric Warping: Geometric registration in projector-camera systems can be obtained using homographies for planar screens either via projector calibration and projective texture-mapping for non-trivial screens with known geometry, or through range imaging for complex surfaces with unknown geometry. Multi-projector techniques usually also apply photometric correction that involves intensity linearization and fitting as well as color gamut matching. Since these topics have been vastly discussed, interested readers may find a good overview of camera-based projector calibration techniques in Brown et al. [6].

Radiometric Compensation: For projection screens with spatially varying reflectance, compensation techniques [18, 12, 3] can be applied to minimize the artifacts induced by the local light modulation between projection and surface. Fujii et al. [9] proposed an approach for compensating dynamic scenes with a co-axial device configuration. For geometrically complex surfaces, however, this requires the viewer to be located on the optical axis of the device - which is impractical for most applications. Content-dependent, adaptive radiometric compensation techniques that are optimized for human perception have been described in [1, 24, 13].

All of these methods presume a well-defined correlation between individual projector and camera pixels.

Focus Related Projector-Camera Techniques: Yet another interesting aspect of projection systems is image focus and defocus. Levoy et al. [14] described how to blur out objects in the camera image, which are not located on the focal plane, and illuminate a scene only on specific slices of the projector's focal plane. Bimber et al. [2] projected images with a large depth of field that are composed of different contributions from multiple overlapping projectors with varying focal planes. The overall sharpness of an image displayed by a single projector was enhanced in [26, 19, 7]. Therefore, the defocus kernels of patterns projected onto complex scenes were captured and employed to perform an image pre-sharpening.

Inverse Illumination: The compensation of scattering for immersive and semi-immersive projection displays with known screen geometry using a reverse radiosity technique was presented in [4]. While the required form factors were pre-computed, Seitz et al. [22] proposed a technique that estimates global illumination parameters with a camera and a laser pointer for canceling indirect illumination in photographs. The cancellation operator that is applied to a picture is a matrix composed of the scene's inverse light transport and a forward light transport matrix containing only the direct illumination contribution. The latter theoretically exists for arbitrary scenes, but has only been shown to be measurable for Lambertian surfaces. Another interesting way to quickly separate direct and indirect illumination for a single camera perspective using structured light projection was introduced in [17]. This, however, does not provide information about the actual light transport itself. A method for removing specular highlights from projections was presented in [20]. It requires the scene geometry to be known and the projectors to be well registered, which is a daunting task.

Forward Light Transport, BRDF Acquisition and Relighting: The forward light transport between a light source and an imaging device implicitly takes all global illumination effects into account. Recently, it has been used for BRDF and BSSRDF acquisition [11, 21], image-based relighting [8, 23, 16, 10] and environment matting [27]. These methods implement various interesting approaches to light transport acquisition.

Our compensation approach combines geometric warping, radiometric and defocus compensation in a single technique by employing the inverse light transport as described in the following sections.

3 Generalizing Radiometric Compensation

The idealized forward light transport between a projector and a camera is given by

$$c_\lambda = T_\lambda p_\lambda + e_\lambda. \tag{1}$$

This is a well-known equation, where c_λ is a single color channel λ of a camera image with resolution $m \times n$. The projector pattern p_λ has a resolution of $p \times q$, and e_λ is the environment light including the projector's blacklevel captured from the camera. A variety of light transport acquisition techniques have been proposed to capture T_λ, we implemented the one described in [23]. For measuring the matrix a set of illumination patterns is projected onto the scene and recorded with the camera. These patterns are designed to refine the light transport representation in a hierarchical manner. Individual matrix entries can then be reconstructed from the captured camera images. The entire acquisition process can take up to several hours, depending on the amount of global illumination effects within the scene. For capturing all lighting effects precisely, high dynamic range (HDR) imaging is a necessity.

In order to account for the device-specific spectral transmission properties of the color filters, equation 1 can easily be extended:

$$\begin{bmatrix} c_R - e_R \\ c_G - e_G \\ c_B - e_B \end{bmatrix} = \begin{bmatrix} T_R^R & T_R^G & T_R^B \\ T_G^R & T_G^G & T_G^B \\ T_B^R & T_B^G & T_B^B \end{bmatrix} \begin{bmatrix} p_R \\ p_G \\ p_B \end{bmatrix}. \tag{2}$$

Each light transport matrix $T_{\lambda_c}^{\lambda_p}$ (size: $mn \times pq$) describes the contribution of a single projector color channel λ_p to an individual camera channel λ_c. Note that in this case each p_λ is the full projector pattern ($p \times q$) and c_λ is the entire camera image ($m \times n$).

Once the transport matrix is acquired it can be used for relighting and synthesizing images from the projector's point of view under camera illumination (dual photography [23]). The former is done by choosing an illumination pattern p_λ and multiplying it to T_λ in eq. 1 (similarly in eq. 2), which yields an image from the camera's viewpoint under the specified illumination. For dual photography the matrix is simply transposed (not inverted). This allows images from the projector's position under arbitrary camera illumination to be generated by multiplying a desired light pattern to the matrix's transpose.

For a generalized radiometric compensation we replace the camera image c by a desired image o of camera resolution and solve the system for p. This pattern is projected[1] onto the real scene and lets a viewer see an undistorted and color corrected image that is similar to o from a limited area close to the calibration camera's position.

To our knowledge, this is the first time that the full light transport is used for radiometric compensation. In previous work, equation 2 was simplified by modeling the transport of each projector pixel as a simple 3×3 matrix [18, 12] (or 3 scalars by discarding the color mixing [3]). Thus, a direct and uncoupled correspondence of individual projector and camera pixels was assumed. When projecting on flat but textured canvases this yields generally satisfying results, however, for more complex surfaces global illumination effects have to be considered as well.
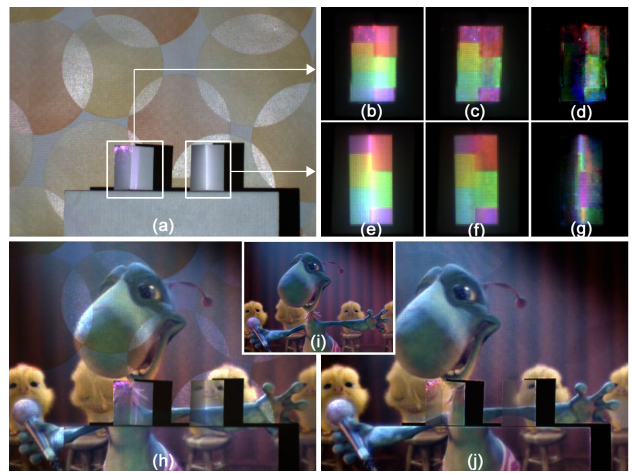

Figure 2. Although diffuse scattering (e+f+g) can be compensated, the physical limits of radiometric compensation become more obvious for specular interreflections (b+c+d). The lower images show a compensation of the entire scene.

The main problem with equations 1 and 2 is their enormous size. It is impossible to store and process a dense matrix that typically has a size of $(640 \times 480) \times (1024 \times 768)$ in the main memory of a personal computer. Thus, when scanning the matrix we insert the captured values directly into a sparse matrix representation. The number of non-zero elements depends on the type and amount of global illumination effects. Clearly, for a scene that comprises a large amount of diffuse scattering between all the visible objects this would lead to a very dense, thus impractical, matrix. However, many of the matrix entries are likely to be very small compared to entries that are due to direct illumination.

---

[1] The projector's response curve needs to be linearized.

These and camera noise make up a large amount of the data. We set a threshold for values that are inserted into the sparse matrix representation. This is chosen by selecting the HDR value of a camera image pixel that exhibits the least amount of illumination that is desired in the matrix. All our computations require a certain sparsity of the matrix.

The sparse equation system can be solved with a simple constrained iterative steepest gradient solver [26]. The constrains ensure that the solution does not contain negative values (projectors cannot display negative light) and that no clamping above 1.0 occurs. In our experiments the solution is found after about 5 iterations in a matter of seconds. We also tested a dense iterative non-negative least squares implementation in a smaller scale that took considerably more time for similar results.

The example presented in figure 2 shows a scene constructed from two V-shaped cardboard pieces in front of a glossy wallpaper (a). The left one is coated with a self-adhesive transparent film. Diffuse scattering and interreflections lead to an increased brightness and to color bleeding in the corner areas (b+e). Performing our radiometric compensation allows these effects to be minimized, as shown in (c+f), the differences (contrast enhanced by 75%) are illustrated in (d+g). Figures 2 (h) and (i) show an uncompensated and a compensated projection of a movie frame ("The Chubb Chubbs", courtesy Pixar). In this case only a single projector is used, shadows could be compensated with multiple displays.

Due to their large optical apertures and resultant narrow depth of field, conventional video projectors can display focused images on single fronto-parallel planes only. Projecting onto surfaces with large depth variance leads to regionally defocused images. Previously described techniques [26, 19, 7] employed the measured defocus kernels of a specific setup to estimate an optimal sharpening of the projection image. This is done by encoding the measured kernel in a matrix and minimizing the sum-of-squared pixel differences between original image and the product of the matrix with the compensation image [26]. Brown et al. [7] applied Wiener Deconvolution. Since the projector defocus is also included in the light transport matrix, blurred projections are implicitly sharpened by our approach. This is demonstrated in figure 3. This simple scene exemplifies projections onto objects with high depth variance, which do not allow the projection to be focused on a single plane. Note, that the limits of all such sharpening approaches are set by the actual defocus of the projector and the amount of soft contours in the original image.

Figure 3. The presentation of an image (a) with a defocused projector (b) can be enhanced by our compensation approach that implicitly sharpens the compensation image (c) that, projected onto the scene, results in a visually more focused projection (d).

4   Towards Real-Time Compensation

Although a compensation based on equation 1 or 2 can be performed very quickly with iterative solvers, this is not fast enough for real-time applications. When projecting movies or interactive 3D content compensation rates of at least 30 frames per second are desired. However, we can reformulate equation 1 (and similarly eq. 2) by applying the light transport matrix's pseudo-inverse $T_\lambda^+ (c_\lambda - e_\lambda) = p_\lambda$. Usually, a matrix's pseudo-inverse is calculated using a singular value decomposition (SVD) $T_\lambda = U\Sigma V^T$, where U and V contain the set of orthonormal left-singular and right-singular vectors and Σ the singular values. The pseudo-inverse is defined as $T_\lambda^+ = V\Sigma^+ U^T$, where the scalars on the diagonal in Σ are individually inverted and zeros not modified to get Σ^+. However, performing an SVD on the light transport matrix is due to its size not feasible. Besides U and V being dense, T^+ is not necessarily sparse, even if T_λ is. In this section we describe a customized clustering scheme that approximates the light transport matrix's pseudo inverse. Using T^+, radiometric compensation reduces to a sparse matrix-vector multiplication, which can be implemented on the GPU and yields real-time framerates. Solving the system with the pseudo-inverse provides a least-squared error solution (also in the degenerate case). Although it violates the non-negative and maximal color value constraints (such values are clamped on the GPU), visual results of our experiments

are similar to the solution of the equation systems.
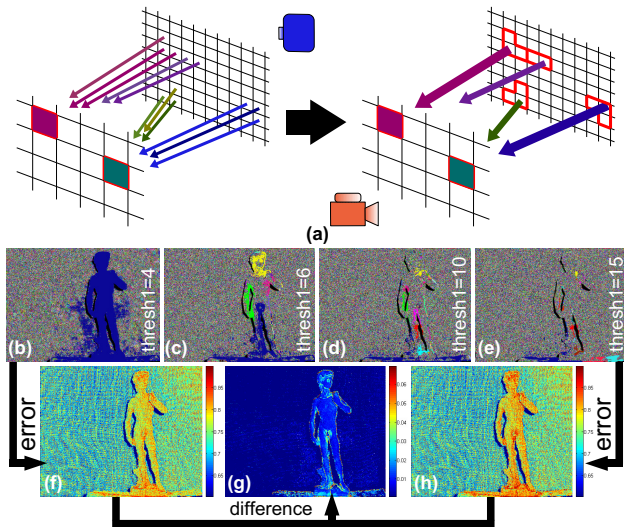

Figure 4. A customized clustering scheme approximates the illumination of spatially close regions in the projector image for each camera pixel. The effect of varying patch radii on the sizes of the clusters ("speckled noise") is depicted in (b-e). RMS errors between original and approximated matrices are visualized in camera space (f-h).

We propose to cluster regions in the light transport matrix to get many independent and smaller equation systems that can be solved individually. Searching for independent subgraphs within the matrix is generally unsuccessful, because neighboring projector pixels are likely to overlap in the camera image. This does not necessarily result from global illumination effects, but is often due to camera and projector defocus, sampling issues, lens imperfections or sensor specific effects such as blooming. Our clustering scheme is designed to approximate the contribution of spatially close projector pixels by a single or by a few points (cf. fig. 4 (a)). This will remove local projector pixel overlaps for individual camera pixels, while approximately preserving contributions from spatially more distant portions in the projector image.

We compute the clusters by grouping all connected projector pixels of a single camera pixel into patches of spatially neighboring regions in the projector space. Patches are formed by searching the projector pixel with the highest luminance contribution for a specific camera pixel and assigning it as the center of a new patch (pseudo-code in lines 2 to 11). All neighboring pixels in the projector image that are also connected

---

Algorithm 1 Light Transport Matrix Clustering

1: **for** all camera pixels c_i **do**
2:   **while not** all connected projector pixels (CPP) of c_i labeled **do**
3:     newPatch.center=findBrightestUnlabeledCPP(c_i)
4:     label(newPatch.center)
5:     **for** all unlabeled CPP p_j **do**
6:       **if** distance(newPatch.center,p_j) $\leq$ thresh$_1$ **then**
7:         newPatch.insert(p_j)
8:         label(p_j)
9:       **end if**
10:     **end for**
11:   **end while**
12:   **for all** patches pa_j **do**
13:     **for all** projector pixels p_k in pa_j **do**
14:       **if** luminance(c_i,p_k) $\leq$ thresh$_2$ **then**
15:         removeConnection(c_i,p_k)
16:       **end if**
17:     **end for**
18:   **end for**
19: **end for**
20: reinsertCutOffProjectorPixels()
21: scaleMatrixRowsToFitFloodlight()

---

to the same camera pixel and within a certain distance (thresh$_1$ in line 6 of alg. 1) to the patch center are inserted into this patch. Elements that are assigned to a specific patch are not further considered for building new patches. This is continued iteratively until all connected projector pixels are part of a patch. We chose thresh$_1$ to be 5-12 pixels for our examples, the effect of this parameter is depicted in fig. 4 (b-e).

Lower luminance contributions within each patch are cut off the graph depending on a given threshold (thresh$_2$ in line 14). Typically we chose a threshold of 30% to 70% (relative to the highest luminance contribution in each patch respectively). Thus, no camera pixels are removed from the graph and projector pixels that are completely cut off are reconnected to the camera pixel that they originally illuminated the most (alg. 1 line 20). The contributions of all remaining connections are scaled to match the original contribution of the entire patch (pseudo-code line 21).

Once the clusters are computed they are individually inverted and recombined to our approximation of T's pseudo-inverse, which is then packed into floating point textures and uploaded on the GPU for real-time compensation. Details about the GPU implementation can be found in [25]. The smaller the clusters are, the sparser is the approximated pseudo-inverse and the higher is the framerate.

The clustering basically removes all defocus information from the transport matrix, which cannot be compensated anymore. For global illumination effects such as scattering or caustics the illumination contribution is approximated within the patches in projector space. Generally, the clustering reduces the light interaction on flat diffuse surfaces that receive no indirect illumination to one-to-one or at least few-to-few correspondences as seen on the background in fig. 4 (b-e) and fig. 5 (a). This approximation is applied to the entire scene by previous compensation approaches. Our clustering method still preserves indirect light interaction. An error analysis for different parameters of the clustering is visualized in figure 4 (f-h). The relative RMS error is the norm of the difference between original light transport matrix and clustered approximation for each matrix row (camera pixel) divided by the norm of the original row. As seen in fig. 4 (f) the relative error, even for a small $thresh_1$, due to clustering is relatively high (about 70% to 80%). This is because some of the direct illumination (defocus) has been filtered out. However, increasing the threshold further mainly modifies the error for surface portions that receive indirect illumination by approx. 3% to 5% (cf. fig. 4 (g,h)) while further reducing the cluster sizes.
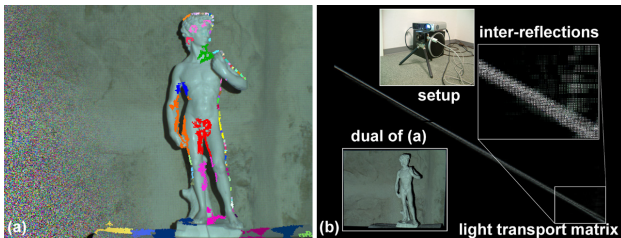

Figure 5. Floodlight image of the scene in fig. 1 superimposed with color-encoded clusters in camera space (a). A visualization of the light transport matrix reveals indirect light contributions (b).

Decomposing the light transport of the sample scene in figure 5 results in smaller clusters within diffuse regions (brick wall) while approximately preserving specular interreflections between the limbs of the statuette (fig. 5 (a)). They make the acquisition of direct pixel correlations between projector and camera very difficult. The larger clusters on the statue's base are due to the steep angle of incident light rays. Indirect illumination is clearly visible in a visualization of the light transport matrix (fig. 5 (b)), which spans the entire background of the image. The matrix's transpose can be used to compute a dual image of the scene as
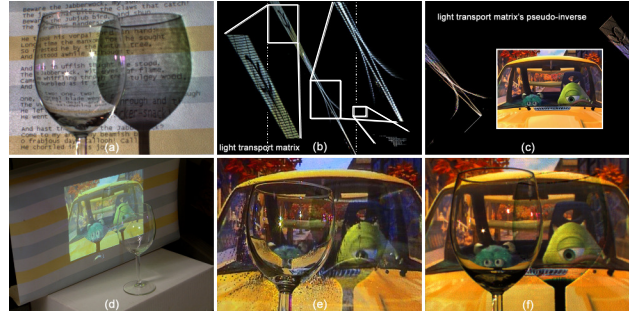

Figure 6. A wine glass in front of a colored wallpaper (a). The light transport matrix's (b) pseudo-inverse (c, background) is approximated with our clustering scheme and allows a real-time compensation for displaying interactive content and movies (c) - from an angle (d), compensated with a conventional method [3] (e) and with our approach (f).

described in [23]. Approximating the pseudo-inverse for this example took about 15 minutes (P4, 3 GHz, 2 GB RAM), while the compensation was then performed with 7 fps on a GeForce 7900 GTX, 512 MB. The displayed content is a screenshot from the short film "9" (courtesy Focus Features and 9, LLC).

Refraction and other complex light modulations represent a challenging problem for structured light scanning techniques, because it is often impossible to determine a one-to-one correspondence between projector and camera pixels. Figure 6 (a), for instance, shows a wine glass in front of a wallpapered surface, which cannot be compensated properly with previous methods (e). The projection of text through the glass exhibits image distortions that can be attributed to refraction. The off-diagonal branches in the acquired light transport matrix (fig. 6 (b) - area bounded by dashed lines) clearly indicate the existence of global illumination effects. While the twisted narrow bands (upper right close-up) are due to refractions, the blank portions on the matrix's diagonal (left magnification) represent the thicker parts of the glass' rim that do not reflect light toward the camera. An interesting effect is highlighted in the magnified lower right part of figure 6 (b). These matrix entries belong to the image of the glass' base that is visible to the camera only because of reflections at the bottom of the glass. Note that this area cannot be fully compensated because it also reflects other parts of the scene that are not illuminated by the projector. Some darker parts of the glass' shadows on the background cannot be corrected with a single projector.

Figure 6 illustrates real-time compensation examples

for a movie sequence (c,d,f, "Mike's New Car", courtesy Pixar). The matrix's pseudo-inverse was computed in approx. 13 minutes (P4, 3 GHz, 2GB RAM). About 30 fps were achieved with the GPU accelerated compensation (GeForce 7900 GTX, 512 MB). This example also demonstrates several physical limitations of radiometric compensation techniques in general. Portions of the camera image that are not lit by the projector or that lie in shadowed regions cannot be compensated. If the same surface is visible multiple times in the camera image (e.g., due to a refracted/reflected and a direct view) it is generally not possible to compute a single compensation value for strongly different input responses. These cases are solved in a least-squared error sense.

5   Discussion and Outlook

Acquiring the forward light transport between a projector and a camera allows capturing the entire light modulation of a projected pattern within an arbitrarily complex scene - including all local and global illumination effects. It represents a 4D slice of the 8D reflectance field and can be described as a linear equation system. Replacing the camera image in the forward light transport equation with a desired (original) picture allows radiometric compensation to be performed by solving these equations for the corresponding illumination pattern and projecting it onto the scene.

In order to achieve real-time compensation rates, the transport matrix's pseudo-inverse is approximated with a customized clustering scheme. This is designed to approximate light contributions locally in the projector space while preserving contributions from spatially more distant portions. It allows a variety of global illumination effects such as diffuse scattering, refractions and interreflections to be compensated. However, captured defocus information is generally lost.

Our generalized approach unifies a set of existing methods that address individual problems with specialized techniques (i.e. [18, 26, 7, 2, 3, 4, 20]). The light transport acquisition time, which is -depending on the scene- between 15 minutes up to several hours, makes many of the previously presented techniques more practical for specific tasks (i.e. for frequently changing setups or for dynamic scenes). However, our method is more general and not limited to specific geometric or radiometric properties of the projection surface. Furthermore, for common static configurations and sceneries a one-time calibration or infrequent adjustments are sufficient.

The presented method reveals several general limitations of radiometric compensation. Shadows and view-dependent effects such as specular reflections may be impossible to be compensated with a single projector. Employing multiple projectors, however, could account for these situations, but have not been implemented for our experiments. Overall focal depth, brightness and resolution could also be increased with a larger number of projectors. Therefore, an extension to our method toward the full 8D reflectance field, which accounts for an arbitrary number of projectors and cameras is described in the appendix. However, in order to validate our theory with respect to multi-device configurations these will have to be implemented and evaluated in the future.

Finally, the limited brightness, resolution, contrast and in particular the relatively high black level contribution of conventional LCD or DLP projectors represent technical limitations that currently prevent making all surfaces types disappear completely when applying radiometric compensation.

As described in [5], multiple sample cameras can be used along with image-based rendering and interpolation techniques to support a view-dependent radiometric compensation of local illumination effects for moving observers. Combining this with the presented technique eventually results in a light field-like rendering technique (see [15]) that performs the compensation of projected images. If the light transport is known for multiple sample cameras, rays in between can be synthesized and rendered in real-time. In order to compensate view-dependent local and global illumination effects such as specular reflections or refractions, many sample cameras are necessary. However, the light transport from a single projector to multiple cameras can be acquired simultaneously. Hence, the overall acquisition time does not increase significantly.

The bottleneck of our current implementation is the light transport acquisition time. The next step will be an accelerated measuring scheme. In case of multiple co-axially mounted projector-camera pairs, for instance, the method introduced in [10] can increase the performance. A more general approach would quickly estimate direct light contributions and their source pixels, for instance using [17], and then progressively refining global illumination contributions during run-time. This would improve the general applicability of our generalized radiometric compensation.

Appendix

Equation 2 can be extended to account for general device configurations with an arbitrary number of cameras and projectors. Distinguishing between different projectors and cameras is done by denoting two additional sub- and superscripts on the left side of the

variables. The radiometric compensation equation for k projectors and r cameras is

$$\begin{bmatrix} {}_1c_R - {}_1e_R \\ {}_1c_G - {}_1e_G \\ \vdots \\ {}_rc_B - {}_re_B \end{bmatrix} = \begin{bmatrix} {}_1^1T_R^R & {}_1^1T_R^G & \cdots & {}_1^kT_R^B \\ {}_1^1T_R^R & {}_1^1T_R^G & \cdots & {}_1^kT_R^B \\ {}_1^1T_G^R & {}_1^1T_G^G & \cdots & {}_1^kT_G^B \\ \vdots & \vdots & \ddots & \vdots \\ {}_r^1T_B^R & {}_r^1T_B^G & \cdots & {}_r^kT_B^B \end{bmatrix} \begin{bmatrix} {}^1p_R \\ {}^1p_G \\ \vdots \\ {}^kp_B \end{bmatrix}$$

References

[1] M. Ashdown, T. Okabe, I. Sato, and Y. Sato. Robust Content-Dependent Photometric Projector Compensation. In *IEEE International Workshop on Projector-Camera Systems (PROCAMS)*, page 6, 2006.

[2] O. Bimber and A. Emmerling. Multi-Focal Projection: A Multi-Projector Technique for Increasing Focal Depth. In *IEEE Transactions on Visualization and Graphics (TVCG)*, volume 12, pages 658–667, 2006.

[3] O. Bimber, A. Emmerling, and T. Klemmer. Embedded Entertainment with Smart Projectors. In *IEEE Computer*, pages 56–63, 2005.

[4] O. Bimber, A. Grundhfer, T. Zeidler, D. Danch, and P. Kapakos. Compensating Indirect Scattering for Immersive and Semi-Immersive Projection Displays. In *IEEE Virtual Reality*, pages 151–158, 2006.

[5] O. Bimber, G. Wetzstein, A. Emmerling, and C. Nitschke. Enabling View-Dependent Stereoscopic Projection in Real Environments. In *International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 14–23, 2005.

[6] M. Brown, A. Majumder, and R. Yang. Camera Based Calibration Techniques for Seamless Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):193–206, 2005.

[7] M. S. Brown, P. Song, and T.-J. Cham. Image preconditioning for out-of-focus projector blur. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1956–1963, 2006.

[8] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *Proc. of ACM SIGGRAPH*, pages 145–156, 2000.

[9] K. Fujii, M. Grossberg, and S. Nayar. A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 814–821, Jun 2005.

[10] G. Garg, E.-V. Tavala, M. Levoy, and H. P. A. Lensch. Symmetric Photography: Exploiting Data-sparseness in Reflectance Fields. *Eurographics Symposium on Rendering (ESGR)*, pages 251–262, 2006.

[11] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel. DISCO: Acquisition of Translucent Objects. In *Proc. of SIGGRAPH*, pages 835–844, 2004.

[12] M. Grossberg, H. Peri, S. Nayar, and P. Belhumeur. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In

[13] A. Grundhöfer and O. Bimber. Real-Time Adaptive Radiometric Compensation. In *To appear in IEEE Transactions on Visualization and Graphics (TVCG)*, 2007.

[14] M. Levoy, B. Chen, V. Vaish, M. Horowitz, I. McDowall, and M. Bolas. Synthetic Aperture Confocal Imaging. In *Proc. of SIGGRAPH*, pages 825–834, 2004.

[15] M. Levoy and P. Hanrahan. Light Field Rendering. In *Proc. of ACM SIGGRAPH*, pages 31–42, 1996.

[16] V. Masselus, P. Peers, P. Dutr;, and Y. D. Willems. Relighting with 4D Incident Light Fields. *ACM Trans. Graph.*, 22(3):613–620, 2003.

[17] S. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination. In *Proc. of ACM SIGGRAPH*, pages 935–944, 2006.

[18] S. K. Nayar, H. Peri, M. D. Grossberg, and P. N. Belhumeur. A Projection System with Radiometric Compensation for Screen Imperfections. In *International Workshop on Projector-Camera Systems (ICCV)*, 2003.

[19] Y. Oyamada and H. Saito. Focal Pre-Correction of Projected Image for Deblurring Screen Image. In *IEEE International Workshop on Projector-Camera Systems (PROCAMS)*, 2007.

[20] H. Park, M.-H. Lee, S.-J. Kim, and J.-I. Park. Specularity-free projection on nonplanar surface. In *Proc. of Pacific-Rim Conference on Multimedia (PCM)*, pages 606–616, 2005.

[21] P. Peers, K. vom Berge, W. Matusik, R. Ramamoorthi, J. Lawrence, S. Rusinkiewicz, and P. Dutré. A Compact Factored Representation of Heterogeneous Subsurface Scattering. In *Proc. of ACM SIGGRAPH*, pages 746–753, 2006.

[22] S. M. Seitz, Y. Matsushita, and K. N. Kutulakos. A Theory of Inverse Light Transport. In *International Conference on Computer Vision (ICCV)*, pages 1440–1447, 2005.

[23] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual Photography. In *Proc. of ACM SIGGRAPH*, pages 745–755, 2005.

[24] D. Wang, I. Sato, T. Okabe, and Y. Sato. Radiometric Compensation in a Projector-Camera System Based on the Properties of Human Vision System. In *IEEE International Workshop on Projector-Camera Systems (PROCAMS)*, 2005.

[25] G. Wetzstein. Radiometric Compensation of Global Illumination Effects with Projector-Camera Systems. Master's thesis, Bauhaus-Universität Weimar, Technical Report #860, 2007.

[26] L. Zhang and S. K. Nayar. Projection Defocus Analysis for Scene Capture and Image Display. In *Proc. of ACM SIGGRAPH*, pages 907–915, 2006.

[27] D. E. Zongker, D. M. Werner, B. Curless, and D. H. Salesin. Environment Matting and Compositing. In *Proc. of ACM SIGGRAPH*, pages 205–214, 1999.