A Generalized Approach to Radiometric Compensation

GORDON WETZSTEIN and OLIVER BIMBER
Bauhaus-University Weimar

Video projectors have evolved tremendously in the last decade. Reduced costs and increasing capabilities have led to widespread applications in home entertainment and visualization. The rapid development is continuing. Projector-camera systems enable a completely automatic calibration. Novel image compensation techniques that allow seamless projections onto complex everyday surfaces have recently been proposed. 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. Furthermore, the anticipated mobility enabled through laptops and pocket projectors will also imply the possibility to use on-site surfaces for presentations instead of carrying projection screens. So far, existing compensation methods consider only local illumination effects, such as diffuse reflections - which is sufficient for many situations. Global illumination effects, such as inter-reflections, refractions, scattering, etc. are ignored. We propose a novel method that applies the light transport matrix for performing an image-based radiometric compensation which accounts for all possible types of light modulation. For practical application the matrix is decomposed into clusters of mutually influencing projector and camera pixels. The compensation is modeled as a linear system that can be solved with respect to the projector patterns. Precomputing the inverse light transport in combination with an efficient implementation on the GPU makes interactive compensation rates possible. Our generalized method unifies existing approaches that address individual problems. Based on examples, we show that it is possible to project corrected images onto complex surfaces such as an inter-reflecting statuette, glossy wallpaper, or through highly-refractive glass. Furthermore, we illustrate that a side-effect of our approach is an increase in the overall sharpness of defocused projections.

Categories and Subject Descriptors: I.3.3 [**Computer Graphics**]: Picture/Image Generation; I.4.9 [**Image Processing and Computer Vision**]: Applications

General Terms: Theory

Additional Key Words and Phrases: Projector-Camera Systems, Radiometric Compensation, Inverse Light Transport

1. INTRODUCTION

Rapid advances in electro-mechanics and optics have increased the capabilities of projectors in terms of spatial resolution, brightness, dynamic range, throw-ratio, and speed. Cost reductions, availability, and the fact that projectors (in contrast

This work was accomplished at the Augmented Reality Laboratory of the Bauhaus-University Weimar. Some of the results were initially presented as a research poster at SIGGRAPH 2006.
Authors' addresses: G. Wetzstein and O. Bimber, Media Faculty, Bauhaus-University Weimar, Bauhausstrasse 11, 99423 Weimar, Germany: email: {wetzstein,bimber}@medien.uni-weimar.de.

to flat panels) can display images that are much larger than the devices themselves made them a mass-market product. Besides for (home-)entertainment, education and business purposes, high resolution tiled-screens and immersive surround screen projections are used for visualizations of scientific, engineering and other content. Emissive CRT technology is loosing more and more ground to light-valve technology, like LCD, LCOS and especially DLP.

The trend toward a higher flexibility of multi-projector systems with respect to device configuration and screen alignment is well noticeable. Recently, numerous projector-camera approaches that enable a seamless projection onto complex everyday surfaces have been proposed. In general this is referred to as *radiometric compensation*. These techniques correct the projected images for geometrical distortions, color and intensity blending, and defocus 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.

Previously proposed radiometric compensation techniques assume a simple geometric relation between cameras and projectors that can be automatically derived using structured light projections or co-axial projector-camera alignments. This results in a precise mapping between camera and projector pixels. In reality, the light of a projected pixel often bounces back and forth several times at different areas on the surface, before it eventually reaches the imaging sensor of the camera. Due to inter-reflection, refraction, scattering and other global illumination effects, multiple camera pixels at spatially distant regions on the image plane may be affected by a single projector pixel. A direct mapping usually considers only camera pixels with the highest intensity contribution that result from the captured light of corresponding modulated projector pixels. Consequently, all global illumination effects are discarded. In some cases it might not even be possible to acquire a direct mapping at all because global effects are too dominant.

We propose a novel image-based approach to radiometric compensation that accounts for all possible local and global illumination effects. Conventional light transport acquisition schemes are employed to capture these effects with projector-camera systems. The goal of our compensation is to find illumination patterns that, when projected onto a complex surface, result in a desired image from the camera's perspective. We model the compensation as a linear system that can be solved with respect to the projected pattern. Due to the size of the resulting equation systems it is necessary to decompose the light transport into clusters of mutually influencing camera and projector pixels. An interactive compensation by solving the system for each frame is in most cases not possible. This, however, can be achieved by pre-computing the inverse light transport, which makes an efficient implementation on the GPU possible. Depending on the complexity of the scene and the occurring global effects, real-time frame rates can be achieved.

2. RELATED WORK

Seamless Multi-Projections: The research on projector-camera systems has recently gained a lot of interest in the computer graphics and the computer vision community. Traditionally, multi-projector configurations are employed to create large-scale high-resolution displays on planar diffuse screens. To achieve a consistent geometrical alignment, as well as a correct photometric and radiometric appearance can be a challenging problem for such systems. Geometric registration can be obtained using homography matrices for planar screens, via projector calibration and projective texture-mapping for non-trivial screens with known geometry, or through look-up tables and per-pixel displacement mapping for complex surfaces with unknown geometry. Photometric correction involves intensity linearization and fitting as well as color gamut matching and cross fading. A good overview of camera-based projector calibration techniques can be found in [Yang et al. 2005].

For projection screens with spatially varying reflectance, radiometric compensation techniques as presented in [Nayar et al. 2003; Grossberg et al. 2004; Bimber et al. 2005; Fujii et al. 2005] can be applied to minimize the artifacts induced by the light modulation between projection and surface pigments. Content-dependent, adaptive radiometric compensation techniques that are optimized to human perception have been described in [Ashdown et al. 2006; Wang et al. 2005; Grundhoefer and Bimber 2006]. All of these methods presume a well-defined mapping between projector and camera pixels.

Focus Related Projector-Camera Techniques: Yet another interesting aspect of projection systems is image focus and defocus. [Bimber and Emmerling 2006] projected images with a large depth of field that are composed from different contributions of multiple overlapping projectors with varying focal planes. The overall sharpness of an image displayed by a single projector was enhanced in [Zhang and Nayar 2006] and [Brown et al. 2006]. Therefore, the defocus kernels of light samples projected onto complex scenes was analyzed. Based on these results, image sharpening was employed to compensate for optical defocus digitally.

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 [Goesele et al. 2004; Peers et al. 2006], image-based relighting [Debevec et al. 2000; Sen et al. 2005; Masselus et al. 2003; Garg et al. 2006] and environment matting [Zongker et al. 1999].

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 [Bimber et al. 2006]. While the required form factors were precomputed, [Seitz et al. 2005] proposed a technique that estimates global illumination parameters with a camera and a laser pointer for canceling inter-reflections in photographs. The operator that is applied to an image for removing indirect il-

lumination is represented as a matrix. This matrix is a composition of the scene's inverse light transport and a forward light transport matrix containing only the indirect illumination contributions. Latter theoretically exists for arbitrary scenes, but has only been shown to be measurable for lambertian surfaces.

3. RADIOMETRIC COMPENSATION AS INVERSE LIGHT TRANSPORT

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

$$\vec{c}_\lambda = T_\lambda \vec{p}_\lambda + \vec{e}_\lambda. \tag{1}$$

This is a well-known equation, where $\vec{c}_\lambda$ is a single color channel λ of a camera image with resolution mn. It is represented as a column vector of size $mn\,x\,1$. T_λ is the light transport matrix, which can be acquired through structured illumination as described in [Sen et al. 2005]. $\vec{p}_\lambda$ is the projector pattern of resolution pq represented as a column vector of size $pq\,x\,1$, and $\vec{e}_\lambda$ is the environment light including the projector's black level captured from the camera (also represented as a $mn\,x\,1$ column vector). Solving equation 1 for $\vec{p}_\lambda$ is equivalent to radiometric compensation of all global and local light modulations contained in T_λ.

Due to the spectral transmission properties of the color filters that are used in cameras and projectors, individual spectral components between both devices cannot be fully separated. Projecting red light only, for instance, leads to non-zero responses in the camera's green and blue color channel. This is known as *color mixing*. Equation 1 is based on the assumption that color mixing between camera and projector is negligible. For deriving a generalized mathematical framework of radiometric compensation it has to be taken into account.

A generalized light transport equation can be formulated that includes color mixing:

$$\begin{bmatrix} \vec{c}_R - \vec{e}_R \\ \vec{c}_G - \vec{e}_G \\ \vec{c}_B - \vec{e}_B \end{bmatrix} = \begin{bmatrix} T_R^R & T_R^G & T_R^B \\ T_G^R & T_G^G & T_G^B \\ T_B^R & T_B^G & T_B^B \end{bmatrix} \begin{bmatrix} \vec{p}_R \\ \vec{p}_G \\ \vec{p}_B \end{bmatrix} \tag{2}$$

Thus, solving equation 2 for $[\vec{p}_R\ \vec{p}_G\ \vec{p}_B]^T$ represents the general radiometric compensation of all light modulations for a single-camera single-projector configuration. It is formulated as a linear equation system of the size $3mn\,x\,3pq$. The subscripts indicate individual color channels. Each single light transport matrix $T_{\lambda 2}^{\lambda 1}$ (size: $mn\,x\,pq$) has three channels for RGB colors, representing the contribution of the projected light to all camera channels $\lambda 2$. The superscripts indicate the light transport matrix that is determined for a specific projector channel $\lambda 1$. Thus T_R^G, for example, is the red color channel of the light transport matrix acquired for the green projector channel.

When describing a direct relation between a single projector and camera pixel, each $T_{\lambda 2}^{\lambda 1}$ is a scalar. In this case, the coefficient matrix is of size $3x3$. This equals the local compensation model used by [Nayar et al. 2003], where the coefficient matrix is referred to as *color mixing matrix* V. However, equation 2 is its generalized form that also takes global light effects into account.

An extension toward general configurations containing an arbitrary number of cameras and projectors can easily be derived from equation 2. Indexing the light transport matrices is done by denoting two additional sub- and superscripts on the left side of T. While the left superscript indicates a projector, the left subscript indicates a camera. Thus, matrix ${}_c^p T_{\lambda 2}^{\lambda 1}$ represents the light transport from projector p to camera c. The radiometric compensation equation for k projectors and r cameras is given by

$$\begin{bmatrix} {}_0\vec{c}_R - {}_0\,\vec{e}_R \\ {}_0\vec{c}_G - {}_0\,\vec{e}_G \\ \vdots \\ {}_{(r-1)}\vec{c}_B - {}_{(r-1)}\,\vec{e}_B \end{bmatrix} = \begin{bmatrix} {}_0^0T_R^R & {}_0^0T_R^G & \cdots & {}_0^{(k-1)}T_R^B \\ {}_0^0T_G^R & {}_0^0T_G^G & \cdots & {}_0^{(k-1)}T_G^B \\ \vdots & \vdots & \ddots & \vdots \\ {}_{(r-1)}^0T_B^R & {}_{(r-1)}^0T_B^G & \cdots & {}_{(r-1)}^{(k-1)}T_B^B \end{bmatrix} \begin{bmatrix} {}^0\vec{p}_R \\ {}^0\vec{p}_G \\ \vdots \\ {}^{(k-1)}\vec{p}_B \end{bmatrix}. \tag{3}$$

It has to be solved for $\left[{}^0\vec{p}_R\ {}^0\vec{p}_G\ \cdots\ {}^{(k-1)}\vec{p}_B\right]^T$. Employing multiple cameras and projectors allows to capture samples of the full 8D reflectance field. While a compensation for multiple cameras can support view-dependent approaches as in [Bimber et al. 2005], a compensation for multiple projectors can remove shadow casts as well as minimize specular highlights (see [Park et al. 2005]), and allows increasing the overall brightness, dynamic range, resolution and focus of projected images, as known from existing multi-projector techniques. However, all possible local and global illumination effects are considered, and existing techniques are unified in our case.

4. CLUSTERING AND HARDWARE ACCELERATION

Due to its enormous size it is impractical to solve the equation system for the entire projector pattern in one step. A possibility of simplifying this computation is to decompose the light transport matrices into clusters of mutually influencing camera and projector pixels. Each of these clusters represents a single, smaller equation system that can be processed independently.

Global and local light modulation will influence the connections within the matrices differently. A flat diffuse surface with an overall focused projector and camera normally produces many small and localized clusters. Capturing a scene with a large depth variance and global light effects will lead to fewer widely connected clusters.

Neighboring projector pixels are likely to overlap in the camera image. This does not necessarily result from global illumination effects, but can also be contributed to camera or projector defocus, differences in resolution, lens imperfections or sensor specific effects such as blooming. Searching for inter-connections in the scanned light transport often leads to a single or to a few large clusters. Since these clusters form equation systems that are too large to be solved efficiently, connections within the matrix must be removed to decompose it into independent sub-clusters. Therefore, the matrix is represented as a weighted bipartite graph that contains camera pixels and projector pixels as nodes. Dependent projector-camera pixel pairs are connected by edges which are weighted with the corresponding luminance transfer contribution. Figure 1 (c) illustrates a sample graph, where all nodes are directly or indirectly connected. Hence, this graph equals a single connected cluster.

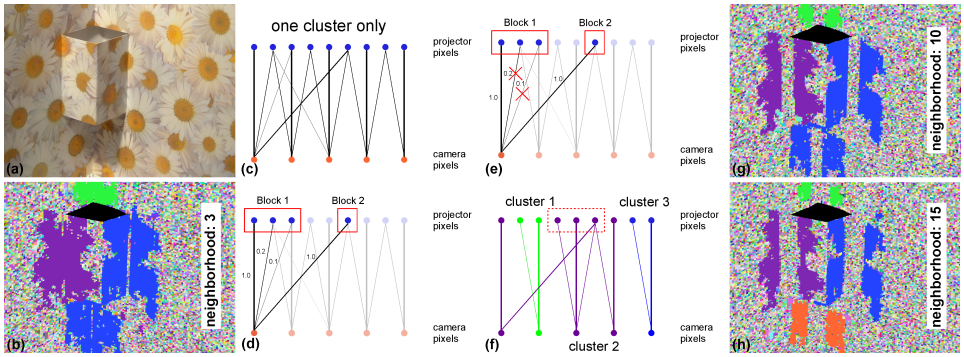

Fig. 1. A camera pixel may be affected by several projector pixels either due to local overlaps or indirect illumination (d). These projector pixels can be grouped into blocks of spatially neighboring pixels in projector space. In order to separate global from local illumination, the contribution for each camera pixel can be restricted per block. A larger neighborhood size leads to smaller cluster sizes (a+b+g+h).

The goal is to split local overlaps while preserving contributions from spatially more distant areas in the projector image which result from global illumination. This can be achieved by grouping all connected projector pixels of a single camera pixel into blocks of spatially neighboring regions in projector space (fig. 1 (d)).

For each camera pixel, blocks are formed by searching the connected projector pixel with the highest edge weight in the graph. This pixel represents the center of a new block. All neighboring pixels in the projector image that are also connected to the same camera pixel are inserted into this block. A constant neighborhood size is used to define the adjacencies. Elements that are assigned to a specific block are not further considered for building new blocks. This is iteratively continued until all connected projector pixels are part of a neighborhood block.

Lower luminance contributions (i.e., edges with lower weights) within each block are then cut out of the graph depending on a given threshold, as it can be seen in figure 1 (e). Searching for new clusters after removing connections leads to a larger

number of smaller clusters (fig. 1 (f)). The process is repeated recursively until each cluster has a predefined size that allows to solve its equations system.

Increasing the neighborhood size is likely to produce smaller and more frequent clusters at the expense of discarding the interaction between close projector pixels as depicted in figures 1 (a+b+g+h)[1].

For radiometric compensation, the camera pixels of each cluster are replaced by pixels of the input image. While $\vec{c}$ and T are known, $\vec{p}$ has to be computed and finally be projected onto the scene. The environment light $\vec{e}$ is a camera image captured under black projector illumination. To avoid negative values in the compensation image $\vec{p}$, the cluster-individual equation systems can be solved separately using iterative non-negative least squares (NNLS) methods. Figure 2 shows an example[2].

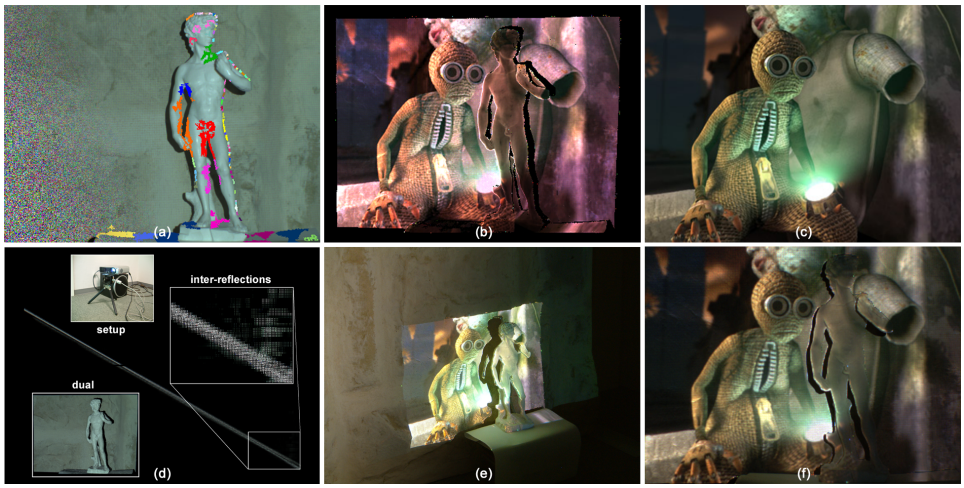

Fig. 2. A floodlight image of a scene containing inter-reflections between the statue's arms and legs partially superimposed with decomposed clusters (a). Applying the light transport matrix visualized in (d) enables to synthesize a compensation image (b) for an input image (c). Projecting this onto the objects (e) results in a corrected view from the camera's perspective (f). Since only one projector was used, shadow casts cannot be filled in the camera view.

Computing a single compensation image by solving multiple equation systems takes several seconds to minutes - depending on the size of the clusters and available computing resources. This makes a real-time compensation of dynamic content, such as movies or real-time graphics impossible.

[1]The luminance filter-threshold remained constant for the visualized decompositions.
[2]The compensation was performed using the idealized light transport (eqn. 1) and took approximately 3 minutes on a P4, 3GHz, 2 GB RAM. The displayed content is a screen shot from the short film "9", Focus Features and 9, LLC.

However, equation 1 can be reformulated by applying the inverse of the light transport matrix to both sides, yielding $T_\lambda^+ (\vec{c}_\lambda - \vec{e}_\lambda) = \vec{p}_\lambda$. Similarly, equations 2 and 3 can be converted. A matrix inverse exists only for regular square matrices. Since this cannot be guaranteed for the light transport matrix of a general projector-camera system its pseudo-inverse has to be determined. This is calculated using SVD, which minimizes the 2-norm of difference between the input image and the result of the radiometric compensation. Each of the pseudo-inverse's columns represents the projector pattern that would have to be projected onto the scene for illuminating only a single camera pixel at a time.

Computing a pseudo-inverse matrix is numerically expensive compared to solving the corresponding equation system explicitly. However, for our case this allows the problem to be split into a computational expensive preprocessing step (computing T^+ for each cluster) and a simple vector dot product of each T^+'s rows and the input image. The latter step is carried out for multiple projector pixels in parallel on the GPU during runtime. A pseudo-inverse is computed for each cluster and inserted into the appropriate locations of the matrix. All of its non-zero elements are packed into floating point textures for direct access on the GPU. A look-up table provides information about the appropriate matrix elements for each projector pixel.

Comparing the NNLS solutions to the result of the multiplication with T^+ did not reveal visible differences. Slight intensity variations in both solutions are due to numerical instabilities. Computing the pseudo-inverses for the example shown in figure 2 took approximately 15 minutes (P4, 3 GHz, 2 GB RAM), while the compensation was performed with 7 fps on a GeForce 7900 GTX, 512 MB.

5. COMPENSATING LOCAL AND GLOBAL ILLUMINATION EFFECTS

Refraction and other complex light modulations represent a challenging problem for structured light scanning techniques. It is often not possible to determine a precise mapping of individual projector pixels and corresponding camera pixels. Figure 3 (a), for instance, shows a glass in front of a wallpapered surface. The projection of text[3] through the glass reveals image distortions that can be attributed to refraction. Multiple characters are visible at different locations within the camera image. These distortions increase near the glass' bottom area.

Geometric image distortions, intensity variations as shadows and caustics, as well as color artifacts from refraction and blending with the background make a radiometric compensation difficult in this example. The off-diagonal branches in the acquired light transport matrix (fig. 3 (b)) 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 3 (b).

[3]The poem "Jabberwocky" by Lewis Carroll from the book "Through the Looking-Glass" (1872).

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. In general, portions of the camera image that are not lit by the projector (such as the darker parts of the glass' shadows on the background) cannot be compensated. Employing multiple projectors can account for such shadow regions.

Fig. 3. A glass in front of a colored wallpaper (a). The acquired light transport matrix (b) is decomposed and used to compute a compensation image (c) for an input image (d). The result captured from the camera's point of view for the uncompensated (e) and the compensated (f) projection. Applying the inverse light transport matrix (g) allows a real-time compensation for displaying interactive content and movies - uncompensated (h) and compensated (i).

Figure 3 illustrates real-time[4] compensation examples for a static photograph and a movie sequence (g)[5]. It also demonstrates several physical limitations of radiometric compensation techniques in general. 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

[4]The matrix's pseudo-inverse was computed in app. 13 minutes (P4, 3 GHz, 2GB RAM). About 30 fps were achieved for compensation with the GPU implementation (GeForce 7900 GTX, 512 MB).
[5]From the short film "Mike's New Car",courtesy Pixar).

generally not possible to compute a single compensation value for strongly different input responses. These cases are implicitly optimized by solving the equation systems in a least squared error sense (e.g., via with NNLS).

The second example, presented in figure 4, shows a scene containing two V-shaped cardboard pieces in front of a glossy wallpaper (fig. 4 (a)). The left one is coated with a self-adhesive transparent film. Diffuse scattering and inter-reflections lead to an increased brightness and to color bleeding in the corner areas (fig. 4 (b+e)). Performing a radiometric compensation to compensate these effects, as shown in figures 4 (c+f), the differences[6] are depicted in figures (d+g). Figures 4 (h) and (j) show an uncompensated and a compensated projection of a movie frame (i)[7].

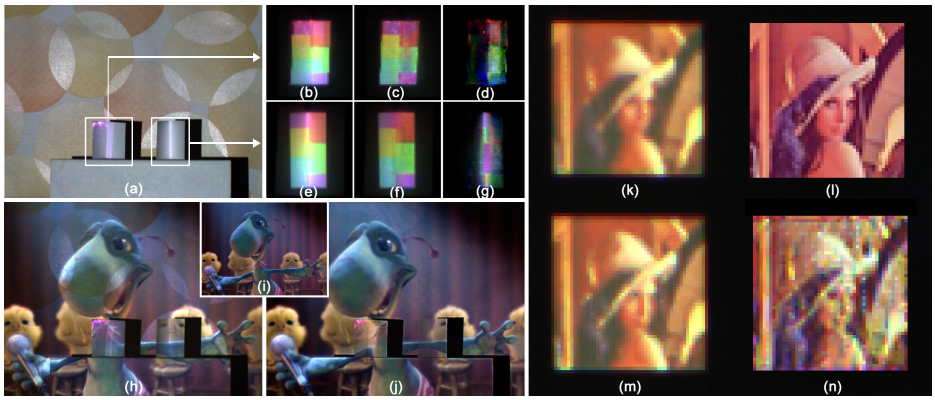

Fig. 4. The inverse light transport enables to compensate diffuse scattering (e+f+g), glossy reflections and inter-reflections (b+c+d), and to increase the perceived focus (n). Uncompensated projections of the original images (i+l) are shown in (h+k), while (m) is a manually sharpened projection of (l).

Because of 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. Since the projector defocus is also included in the light transport matrix, blurred projections are implicitly sharpened by the radiometric compensation. This is demonstrated in figures 4 (k-n)[8]. An uncompensated projection of a picture (l) with a defocused projector on a planar uniformly colored surface is shown in (k). The desired image is very small (48x48 pixels), however, the resulting equation system contains 2304x2700 elements. About 25 minutes were required for solving this with a non-negative least squared error solution (P4, 3GHz, 4 GB

[6]The images are contrast enhanced by 75%.
[7]From the short film "The Chubb Chubbs", courtesy Pixar.
[8]Note that the light transport matrix is not decomposed for this experiment, however, only affected projector and camera pixels are included in the equation system.

RAM) on the CPU.

Previously described techniques such as [Zhang and Nayar 2006] and [Brown et al. 2006] employed the measured defocus kernels of the specific setup to estimate an optimal sharpening of the projection image. Due to the optical defocus of the projector, the captured image appeared more similar to the original image than its unmodified projection. The idea of focus optimization through pre-sharpening is illustrated in figure 4 (m) while figure 4 (n) is the result of our radiometric compensation. Although only a standard sharpening operator (as found in most image processing applications) was used, the captured projection appears more focused. The limits of sharpening are set by the actual defocus of the projector and of the original image. Thus, it is not possible to compensate optical defocus if the original image does not contain a minimal amount of digital blur.

6. 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 picture enables to perform radiometric compensation by solving this equation system for the corresponding projector pattern.

The single-projector single-camera forward light transport equation can be extended to support multi-projector multi-camera setups, as well as color mixing between the devices. This leads to a large equation system that addresses the general case of the full 8D reflectance field. As in the 4D case, this equation system has to be solved for the projected radiance based on the expected irradiance on the camera sensor for radiometric compensation.

To achieve this in practice, the transport matrix is decomposed into a set of independent clusters that form smaller equation systems which can finally be processed. Solving the set of equation systems explicitly on the CPU, however, does not support a compensation in real-time. Pre-computing the inverse light transport matrix on the CPU and performing a simple matrix-vector multiplication during run-time on the GPU does lead to interactive frame-rates.

With this general approach it becomes possible to compensate a variety of local and global illumination effects with a single technique. This unifies a pallet of existing methods that all address individual problems with specialized techniques (e.g., [Nayar et al. 2003; Zhang and Nayar 2006; Brown et al. 2006; Bimber and Emmerling 2006; Bimber et al. 2005; Bimber et al. 2006; Park et al. 2005; Seitz et al. 2005]). Examples for inter-reflections, refractions, diffuse scattering, and the sharpening of defocused projections have been provided in this article. However, it can be expected that other modulation types, such as diffractions, sub-surface

scattering, or specular highlights will follow the same scheme. All experiments have been carried out with single-camera single-projector setups. In order to validate our theory, multi-projector multi-camera configurations have to be implemented and evaluated in the future.

As described in [Bimber et al. 2005], 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 generalized approach that is described in this article eventually results in a light field like rendering technique that performs radiometric compensation. If the light transport is known for multiple sample cameras, rays in between can be synthesized and rendered in real-time (see [Levoy and Hanrahan 1996]). 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 much.

The discussed generalized approach reveals several general limitations of radiometric compensation. Shadows and view-dependent effects such as specular reflections may not be compensatable with a single-camera single-projector configuration. Employing multiple projectors and cameras, however, can account for these situations. The overall focal depth, brightness and resolution are increased as well. However, it is impossible to find an exact solution when a single projector pixel is mapped to multiple camera pixels with different values in the input image (i.e., if the same surface portion is visible multiple times in the camera image). This can be the result of reflections, refractions, or other global effects.

Applying the inverse light transport, either by solving the equation system with NNLS or by precomputing the transport matrix's pseudo-inverse with SVD and multiplying it with the input image, minimizes the squared error between input image and provided solution in the first, case and the 2-norm of difference between both images in the second case.

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 from making all surfaces types disappear completely when applying radiometric compensation.

REFERENCES

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

Bimber, O. and Emmerling, A. 2006. Multi-Focal Projection: A Multi-Projector Technique for Increasing Focal Depth. In *IEEE Transactions on Visualization and Computer Graphics (TVCG)*. Vol. 12.

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

Bimber, O., Grundhoefer, A., Zeidler, T., Danch, D., and Kapakos, P. 2006. Compensating Indirect Scattering for Immersive and Semi-Immersive Projection Displays. In *IEEE Virtual Reality (IEEE VR'06)*.

Bimber, O., Wetzstein, G., Emmerling, A., and Nitschke, C. 2005. Enabling View-Dependent Stereoscopic Projection in Real Environments. In *International Symposium on Mixed and Augmented Reality (ISMAR'05)*. 14–23.

Brown, M., Song, P., and Cham, T.-J. 2006. Image Pre-Conditioning for Out-of-Focus Projector Blur. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'06)*.

Debevec, P., Hawkins, T., Tchou, C., Duiker, H.-P., Sarokin, W., and Sagar, M. 2000. Acquiring the Reflectance Field of a Human Face. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 145–156.

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

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

Goesele, M., Lensch, H. P. A., Lang, J., Fuchs, C., and Seidel, H.-P. 2004. DISCO: Acquisition of Translucent Objects. *ACM Trans. Graph. 23,* 3, 835–844.

Grossberg, M., Peri, H., Nayar, S., and Belhumeur, P. 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Vol. I. 452–459.

Grundhoefer, A. and Bimber, O. 2006. Real-Time Adaptive Radiometric Compensation. In *SIGGRAPH Poster*.

Levoy, M. and Hanrahan, P. 1996. Light Field Rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. ACM Press, New York, NY, USA, 31–42.

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

Nayar, S. K., Peri, H., Grossberg, M. D., and Belhumeur, P. N. 2003. A Projection System with Radiometric Compensation for Screen Imperfections. In *International Workshop on Projector-Camera Systems (PROCAMS 2006)*.

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

Peers, P., vom Berge, K., Matusik, W., Ramamoorthi, R., Lawrence, J., Rusinkiewicz, S., and Dutré, P. 2006. A Compact Factored Representation of Heterogeneous Subsurface Scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH) 25,* 3 (July).

Seitz, S. M., Matsushita, Y., and Kutulakos, K. N. 2005. A Theory of Inverse Light Transport. In *ICCV*. 1440–1447.

Sen, P., Chen, B., Garg, G., Marschner, S. R., Horowitz, M., Levoy, M., and Lensch, H. P. A. 2005. Dual Photography. *ACM Trans. Graph. 24,* 3, 745–755.

Wang, D., Sato, I., Okabe, T., and Sato, Y. 2005. 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)*.

Yang, R., Majumder, A., and Brown, M. 2005. Camera Based Calibration Techniques for Seamless Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics 11,* 2.

Zhang, L. and Nayar, S. K. 2006. Projection Defocus Analysis for Scene Capture and Image Display. *ACM Trans. on Graphics (also Proc. of ACM SIGGRAPH)*.

Zongker, D. E., Werner, D. M., Curless, B., and Salesin, D. H. 1999. Environment Matting and Compositing. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 205–214.