

Projectors for Graphics

SIGGRAPH 2008 Course

Ramesh Raskar, MIT Media Lab
Aditi Majumder, University of California, Irvine
Hendrik P. A. Lensch, MPI Informatik
Oliver Bimber, Bauhaus-University Weimar

Course Abstract

Modern digital projectors are a central part of large-format displays, non-intrusive augmented reality systems and computational illumination for 3D image-based modeling. High speed and high framerate projectors also support intriguing applications in optical communication. With a pocket-size form factor, projectors will be widely used for mobile applications. We survey this rapidly evolving landscape and the growing interest in experimenting with projectors. A novel class of applications is emerging, involving illumination+capture of complex 3D shapes as well as dynamic interaction via projection on movable surfaces. We provide a detailed survey of the several approaches for combining real-time computer graphics and computer vision methods for single and multi-projector systems. We cover topics in immersive rendering, projective geometry, reflectance field capture and spatial augmented reality. In addition, we will give practical insights, implementation details and case-studies for a variety of applications in research, art and industry.

Prerequisites

This course is appropriate for beginners in digital art and media. No programming or specific mathematical background is required. General knowledge of basic computer graphics techniques and 3D tools is helpful but not necessary.

Course Syllabus

Course Introduction and Overview (Bimber and Raskar, 10 minutes)

- Goals, Outline, Speakers, Schedule, Opportunities

A. Large Format Displays (Majumder, 40 minutes)

From CAVEs to large visualization centers, single and multi-projector displays are becoming easier to use due to novel camera-based maintenance systems.

- Overview and New opportunities
- Planar, cylindrical, spherical and non-planar displays
- Geometric and color calibration
- Rendering strategies
- Distributed self-calibrating displays

B. Visually Augmenting the real World with Projectors (Bimber, 50 minutes)

Augmenting real world surfaces with projected images can be challenging. This module describes the fundamental concepts.

- Image correction techniques for geometric and radiometric (colored and textured) complex surfaces
- Human vision adapted techniques
- Global illumination compensation (scattering, inter-reflections, caustics)
- Human vision adapted techniques
- Defocus compensation (with multiple and single projectors)
- Imperceptible coded projection
- Superresolution and high-dynamic range projection
- Applications (theatres, museums, historic sites, advertisement, games, television studios, movie sets, on-site visualization, etc.)

Questions and Answers (All, 5 minutes)

Break (15 minutes)

C. Mobile Projectors and Optical Communication (Raskar, 45 minutes)

Pocket projectors allow novel human-computer interaction opportunities. Spatio-temporal modulation of light creates high speed optical communication which can be used in many tracking applications.

- Portable projectors, technology and issues
 - Single-handed interaction, Image stabilization and resizing
 - iLamps: Geometrically aware pocket projectors
- Optical and Radio Frequency Tags
 - RFID for Augmented Reality: Location sensing RFID and automatic authoring
 - Optical communication for space labeling in robotics, games
 - Imperceptible projection for high speed motion capture

D. Computational Illumination for 3D Scene Modeling (Lensch, 45 minutes)

The appearance of real world objects can be represented via a high-dimensional reflectance fields. Projectors are highly suitable for computational illumination to model such objects for rendering and computer vision applications.

- Scene appearance as higher dimensional reflectance fields
- Pattern projection for 3D geometry acquisition
- Measuring appearance parameters
- Capturing (and removing) global versus local illumination effects

Questions and Answers (All, 15 minutes)

Course Presenters

Ramesh Raskar

Associate Professor
Media Lab, MIT
20 Ames Street,
Cambridge, MA 02139
T +1.617.621.7533
F +1.617.621.7550
Email: raskar@merl.com, <http://raskar.info>

Ramesh Raskar will join MIT full time in Spring'2008 and is currently a Senior Research Scientist at MERL. His research interests include projector-based graphics, projective geometry and non-photorealistic rendering. During his doctoral research at U. of North Carolina at Chapel Hill, he developed a framework for projector based 3D graphics, which can simplify the constraints on conventional immersive displays, and enable new projector-assisted applications. He has published several articles on immersive projector-based displays, spatially augmented reality and has introduced Shader Lamps, a new approach for projector-based augmentation. His technical papers have appeared in SIGGRAPH, EuroGraphics, IEEE VR, IEEE Visualization, CVPR and many other graphics and vision conferences. He was a course organizer and speaker for Siggraph 2002 through 2007. He is a member of the ACM and IEEE.

Aditi Majumder

Assistant Professor
University of California, Irvine
Department of Computer Science
3029 Bren Hall, Irvine, CA 92697-3435
T +1 (949) 824-8877
F +1 (949) 824-4056
E-Mail: majumder@ics.uci.edu,
Web: <http://www.ics.uci.edu/~majumder/>

Aditi Majumder is an assistant professor at the Department of Computer Science in University of California, Irvine. She received her BE in Computer Science and Engineering from Jadavpur University, Calcutta, India in 1996 and PhD from Department of Computer Science, University of North Carolina at Chapel Hill in 2003. Her research is in the general area of computer graphics, vision and visualization. Her primary research contribution has been on using computer vision methodologies for easy and automatic assembly of commodity projector(s) and camera(s) to build large tiled multi-projector displays, exploiting human perception limitations for making them perceptually and functionally seamless. and devising human computer interaction modalities for such displays. She is the author of the recent book "Practical Multi-Projector Display Design" released by A.K.Peters in 2007.

Hendrik P. A. Lensch

MPI Informatik
Campus E 1.4
66123 Saarbruecken
Germany
T +49 681 9325 428
F +49 681 9325 499
E-Mail: lensch@mpi-inf.mpg.de
Web: <http://www.mpi-inf.mpg.de/~lensch>

Hendrik P. A. Lensch is the head of an independent research group "General Appearance Acquisition and Computational Photography" at the MPI Informatik in Saarbrucken, Germany. The group is part of the Max Planck Center for Visual Computing and Communication. He received his diploma in computers science from the University of Erlangen in 1999, and after joining the computer graphics group at MPI he received his PhD from Saarland University in 2003. Dr. Lensch spent two years (2004-2006) as a visiting assistant professor at Stanford University, USA. His research interests include 3D appearance acquisition, image-based rendering and computational photography. For his work on reflectance measurement he received the Eurographics Young Researcher Award 2005. He was awarded an Emmy Noether Fellowship by the German Research Foundation in 2007. He has given several lectures and tutorials at various conferences including SIGGRAPH courses on realistic materials in 2002 and 2005.

Oliver Bimber

Associate Professor

Media Faculty

Bauhaus University, Bauhausstr. 11, 99423 Weimar, Germany

T +49-3643-583724

F +49-3643-583709

Email: bimber@uni-weimar.de

<http://www.uni-weimar.de/medien/AR>

Oliver Bimber is a Junior Professor (associate level) for Augmented Reality at the Bauhaus University Weimar, Germany. He received a Ph.D. (2002) in Engineering at the Darmstadt University of Technology, Germany and a Habilitation degree (2007) in Computer Science (Informatik) at the Munich University of Technology. He is co-author of the book "Spatial Augmented Reality" and serves on the editorial board of the IEEE Computer magazine (graphics and multimedia editor). Bimber taught courses on Spatial Augmented Reality and projector-based display techniques at Eurographics 2003 and Eurographics 2004, ICAT 2005, ETD 2006, as well as at Siggraph 2005, Siggraph 2006 and Siggraph 2007. His research interests include real-time rendering, computer vision and human visual perception in the context of next-generation display technologies.

Large Format Displays


Aditi Majumder
University of California, Irvine

Multi-projector Displays

- Tile multiple projectors
 - Covers a much larger viewing area
- Logical abstraction of a single display
- Seamless imagery across projectors

Commercial Systems

- CAVE



Design of a CAVE

CAVE at University of Illinois, Chicago



Commercial Systems

- CAVE
- Powerwall




First Generation Displays

- Cost Prohibitive
 - Projectors (\$75,000)
 - SGI Infinite Reality (\$1,000,000)
- Manual Registration
 - Expensive 6 DOF mounts
 - Fresnel lens
 - Manual manipulation
 - Projector and mount controls



Courtesy: ANL

First Generation Displays



- Precise *abutting* construction
- Hardwired in rendering software

Problems

- Rigid permanent structures in dedicated rooms
- Not scalable
- Not easily deployable
- Not reconfigurable

Current Generation Displays

- Affordable
 - Portable projectors, PC Cluster Rendering
 - 10 projector wall < \$50,000
- Casually aligned
- No expensive optics
- Allowing *overlaps* between projectors

Geometric & Photometric Mismatch



Registration for Seamless Display



Camera Based Registration

- Camera feedback detects misregistration
- Encoded in a mathematical function
 - Both geometric and photometric
- Change the projected image digitally
 - Apply the inverse function
 - In real-time via GPU

Overview

- Geometric Registration
- Photometric Registration
- PC Cluster Based Rendering
- Distributed Rendering

Overview

- Geometric Registration
- Photometric Registration
- PC Cluster Based Rendering
- Distributed Rendering

Classification

- Based on nature of display surface
 - Parametric
 - Planar
 - Non-planar
 - Non-parametric

Classification

- Based on nature of display surface
 - Parametric
 - Planar
 - Non-planar
 - Non-parametric


Notation

- Display $D (s', t')$
- Image $I (s, t)$
- Projector $P_i (x_i, y_i)$
- Camera $C (u, v)$
- Display parameterization is directly related to image parameterization
 - $(s', t') = (s, t)$


Basic Methodology

- Goal: Find G mapping (x_i, y_i) to (s, t)
- Method: Introduce a camera C
 - Find H mapping projector (x_i, y_i) to camera (u, v)
 - Camera observes projected features
 - Find F mapping camera (u, v) to image (s, t)
 - Camera observes display surface
 - $G = F \bullet H$
- Result: Image seamlessly wallpapered on D

Basic Methodology



Illustration



Different Methods for Planar Display

- Representation of G , H and F
 - Type of display surface
 - Geometric imperfections in projectors
 - Desired accuracy
- A centralized computer controls
 - One/multiple camera(s)
 - Multiple projectors
 - Planar or non-planar

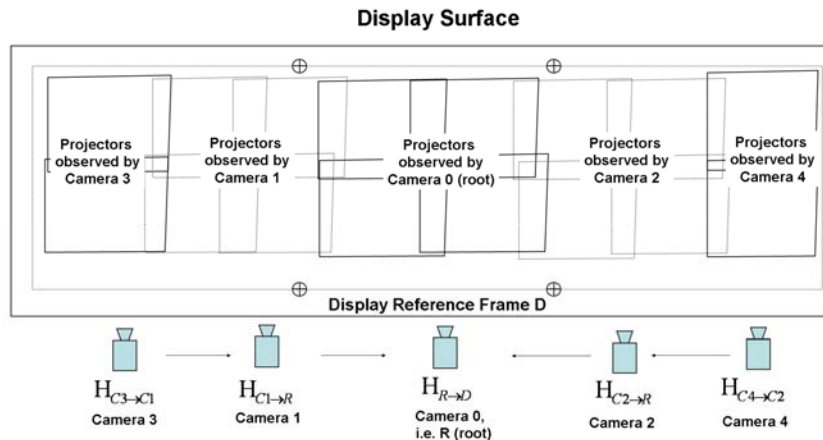
Linear Method

- Assumptions
 - Perfect linear devices (No radial distortion)
- H and F are both linear 3x3 matrices
 - Commonly called *homography*
- $G = F \times H$
 - Matrix multiplication
- G^{-1} applied to I to generate image for each P_i
 - Easy to find the inverse

Using Multiple Cameras

- Homographies can be concatenated
- Scalability not limited by camera resolution
- Cheaper cameras with smaller FOV

Set Up



Method

- FOV of adjacent cameras C_j and C_k overlaps
- C_j and C_k are related by homography
 - $H_{C_j \rightarrow C_k}$
 - Observing projected points in overlapping FOV
- Choose a root camera R
- R is related to D by a homography
 - $H_{R \rightarrow D}$

Method

- C_j can be related to D by a concatenation of camera homographies
 - $H_{C_j \rightarrow D} = H_{R \rightarrow D} \times H_{C_k \rightarrow R} \times \dots \times H_{C_j \rightarrow C_k}$
- More than one path from C_j to R
 - Minimum spanning homography tree
 - Hence, unique path

Method

- Projector P_i can be related to C_j
 - $H_{P_i \rightarrow C_j}$
- Hence, P_i can be related to D by concatenation of homographies
 - $H_{P_i \rightarrow D} = H_{C_j \rightarrow D} \times H_{P_i \rightarrow C_j}$
- Errors can accumulate along a path of tree
 - Global error diffusion

Non-Linear Method for Planar Display

- Projectors can have non-linearities
- Rear projection systems
- H is non-linear
- Issues
 - Not easily invertible
 - Cannot be concatenated
 - Cannot scale to multiple cameras

Non-Linear Method for Planar Display

- H is a cubic polynomial
 - Linear regression for polynomial fitting
- Issues
 - Not perspective projection invariant
 - Assumes near rectangular array

Non-Linear Method for Planar Display

- H is a rational Bezier function
 - Perspective projection invariant
 - Can tolerate large non-linearities
 - Uses iterative procedure (Levenberg-Marquadt) for Bezier fitting
 - Assures global smoothness of lines
 - Requires sparse sampling

Results



Results



Piecewise Linear Method

- H is a piecewise linear function
 - Not as compact
 - Reduces local errors
 - Requires dense sampling

Planar Displays

References:

- R. Raskar, *Immersive Planar Display using Roughly Aligned Projectors*, IEEE VR, 2000.
- H. Chen, R. Sukthankar, G. Wallace, *Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees*, IEEE Visualization, 2002.
- R. Yang, D. Gotz, J. Henseley, H. Towles, M. S. Brown, *PixelFlex: A Reconfigurable Multi-Projector Display System*, IEEE Visualization, 2001.
- M. Hereld, I. Judson, R. Stevens, *DottyToto: A Measurement Engine for Aligning Multi-Projector Display Systems*, Argonne National Laboratory preprint ANL/MCS-P958-0502, 2002.
- E. Bhasker, R. Juang, A. Majumder, *Registration Techniques for Using Imperfect and Partially Calibrated Devices in Planar Multi-Projector Displays*, IEEE Visualization, 2007.

Classification

- Based on nature of display surface
 - Parametric
 - Planar
 - Non-planar
 - Non-parametric

Parametric Non-planar Display

- Cylindrical display
- Display parameterization
 - Equally placed physical markers
 - Top and bottom rim of the surface
- H and F cannot be closed form
 - Piece-wise linear functions
 - Sample densely

Set-Up





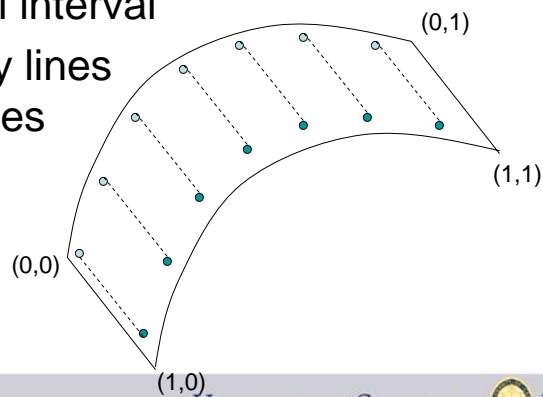
Image and Display Parameterization

- $D(s', t') = I(s, t)$



Fiducials only on the rims

- Equal distance from rim and equal interval
- Connected by lines parallel to sides



Camera to Display - F

- Find camera to display correspondence
- Distances may not be equal due to perspective
- Piece-wise linear triangulation





Image Triangulation


- Do a similar triangulation in image space



Projector to Camera Mapping: H

- Projector Pattern
 - Equally placed blobs


Pattern to be projected



Projector to Camera Mapping: H


- Projector pattern
 - Equally spaced blobs
- Piece-wise linear triangulation

Pattern to be projected



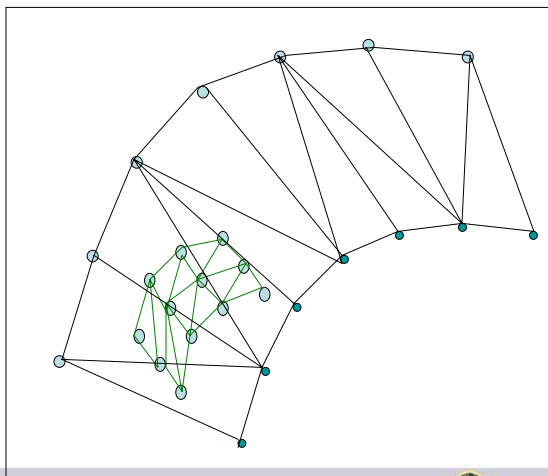
Projector to Camera Mapping: H

- Find corresponding points in camera image
- Do a similar triangulation



Two meshes in camera space

- Camera to projector mesh: H
- Camera to display mesh: F
- Use $H \bullet F$ to generate projector image



Results



Classification

- Based on nature of display surface
 - Parametric
 - Planar
 - Non-planar
 - Non-parametric

Non-parametric Displays

- $G = H \bullet F$
- F is ignored
 - Assume camera and image space are same
 - $(s, t) = (u, v)$
 - View dependent
 - Viewer is located where the camera is
- H is a piecewise linear function

Non-parametric Display





Illustration: Corner of Two Walls

Original projector input



Projected image is distorted



Warped projector input




Projected image is undistorted



Structured Light

- Projector Pattern
 - Equally placed blobs


Pattern to be projected



Structured Light

- Use multiple frames
- Binary encode blob index
- Each bit represents a frame
- Each blob with the bit on is present in that frame


Pattern to be projected



Structured Light

- Use multiple frames
- Binary encode blob index
- Each bit represents a frame
- Each blob with the bit on is present in that frame


Frame 1 (LSB)



Structured Light

- Use multiple frames
- Binary encode blob index
- Each bit represents a frame
- Each blob with the bit on is present in that frame


Frame 2 (2nd LSB)



Structured Light

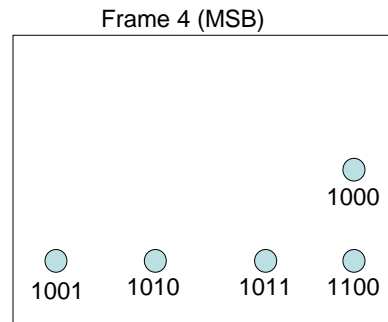
- Use multiple frames
- Binary encode blob index
- Each bit represents a frame
- Each blob with the bit on is present in that frame

Frame 3 (3rd LSB)



Structured Light

- Use multiple frames
- Binary encode blob index
- Each bit represents a frame
- Each blob with the bit on is present in that frame



Non-planar Displays

References:

- M. Harville, B. Culbertson, I. Sobel, D. Gelb, A. Futzhugh, D. Tanguay, [Practical Methods for Geometric and Photometric Correction of Tiled Projector Displays on Curved Screens](#), IEEE CVPR Workshop on Projector-Camera Systems, 2006.
- R. Raskar, J. VanBaar, T. Willwacher, S. Rao, [Quadric Image Transfer for Immersive Curved Screen Displays](#), Eurographics 2004.
- R. Yang, A. Majumder, M. S. Brown, [Camera-Based Calibration Techniques for Seamless Multi-Projector Displays](#), IEEE Transactions on Visualization and Computer Graphics 11(2), Mar-Apr 2005
- R. Raskar and G. Welch and M. Cutts and A. Lake, L. Stesin, H. Fuchs, [The Office of the Future: A Unified Approach to Image Based Modeling and Spatially Immersive Display](#), ACM Siggraph, 1998.
- R. Raskar, M.S. Brown, R. Yang, W. Chen, H. Towles, B. Seales, H. Fuchs, [Multi Projector Displays Using Camera Based Registration](#), IEEE Visualization, 1999.

Overview


- Geometric Registration
- Photometric Registration
- PC Cluster Based Rendering
- Distributed Rendering

The Problem



- Perfect geometric alignment
- Color variation problem not addressed
- Breaks the illusion of a single display



The Problem



Overlapping



Abutting

The Goal






The Goal



Should look like a single display
Cannot tell the number of projectors

Background: Color

- Perceptual Representation
 - Luminance (L)
 - Brightness 
 - Chrominance (x, y)
 - Hue and Saturation 

- Representation Using Primaries
 - Three channels (Red, Green, Blue)
- Color Gamut



Organization


- Properties of Color Variation
- Existing Methods
- Modeling Color Variation
- PRISM

Organization

- **Properties of Color Variation**
- Existing Methods
- Modeling Color Variation
- PRISM

Properties of Color Variation

- Intra-projector
 - Within a single projector
- Inter-projector
 - Across different projectors
- Overlaps



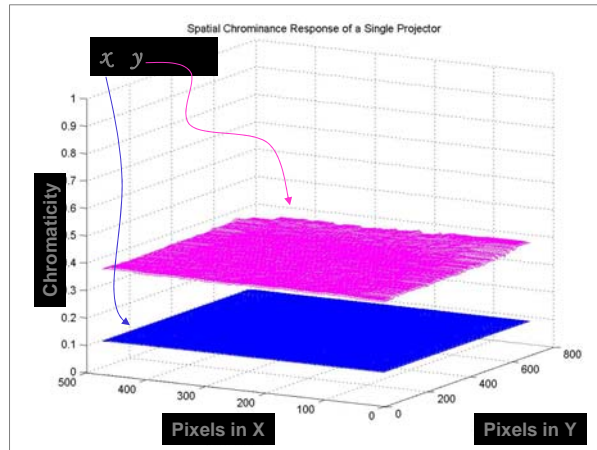
Properties of Color Variation

- Intra-projector
 - Within a single projector
- Inter-projector
 - Across different projectors
- Overlaps



Intra-Projector Variations



- Chrominance is constant



Intra-Projector Variations

- Luminance is not
- Black Offset
 - Always present

Luminance variation is more significant than chrominance variation



Reasons

- Distance attenuation of light
- Non-Lambertian nature of the display
 - Rear projection systems show a higher variation than the front projection systems


Properties of Color Variation

- Intra-projector
 - Within a single projector
- Inter-projector
 - Across different projectors
- Overlaps



Inter-Projector Variations

- Projectors of same model
 - Luminance variation is much more significant than chrominance variation
- Projectors of different models
 - Chrominance variation is relatively very small
 - Luminance variation is significant



Chrominance (x) of a four overlapping projector display

Luminance variation is more significant than chrominance variation

Reasons

- Bulb characteristics
 - Inherent properties of the bulbs
 - Different ages of the bulbs
- Imprecision in generating identical optical elements
 - Lens and filters
- Position and Orientation
- Control settings


Properties of Color Variation

- Intra-projector
 - Within a single projector
- Inter-projector
 - Across different projectors
- Overlaps



Overlaps

- For displays made of same model projectors, at overlap regions
 - Chrominance remains almost constant
 - Luminance almost gets multiplied by the number of overlapping projectors



Luminance variation is more significant than chrominance variation

Properties of Color Variation

References:

- Aditi Majumder, **Properties of Color Variation in Multi Projector Displays**, Proceedings of SID Eurodisplay, 2002.
- Aditi Majumder and Rick Stevens, **Color Non-Uniformity in Multi Projector Displays: Analysis and Solutions**, IEEE Transactions on Visualization and Computer Graphics, Vol. 10, No. 2, Mar-Apr 2003.

Organization

- Properties of Color Variation
- **Existing Methods**
- Modeling Color Variation
- PRISM

Existing Methods

- Manipulation of Projector Controls
- Common Bulb
- Edge Blending
- Gamut Matching

Existing Methods

- Manipulation of Projector Controls
- Common Bulb
- Edge Blending
- Gamut Matching

Manipulation of Projector Controls

- Manually or computer controlled
- Sensor
 - Eye or camera
- Limitations
 - Can be used for abutting configurations only
 - Addresses only inter projector variations
 - Time consuming and labor intensive
 - Not scalable to 40-50 projectors

Existing Methods

- Manipulation of Projector Controls
- Common Bulb
- Edge Blending
- Gamut Matching

Common Bulb

- Using common bulb for all projectors
- Limitations
 - Bulb is not the only cause of color variation
 - Not scalable
 - Expensive (\$100,000 for 3x3 display)
 - Labor intensive
 - Can be used in abutting projector configuration only
 - Addresses only inter projector variation

Common Bulb


References:

- B. Pailthorpe, N. Bordes, W. Bleha, S. Reinsch, and J. Moreland, [High-resolution display with uniform illumination](#), Proceedings Asia Display IDW, 1295-1298, 2001.

Existing Methods

- Manipulation of Projector Controls
- Common Bulb
- Edge Blending
- Gamut Matching

Edge Blending




Edge Blending

- Software Edge Blending
- Aperture Edge Blending
- Optical Edge Blending

Software Edge Blending

Camera image



x has contributions from $P_1(x)$ and $P_2(x)$

Intensity at x :

$$\alpha_1(x)P_1(x) + \alpha_2(x)P_2(x)$$

Find alpha such that:


$$\alpha_1(x) + \alpha_2(x) = 1$$

Algorithm

Assign intensity weights based on x 's distance from projector boundaries


Assigning Intensity Weights

Camera image



$d_1 = x$'s distance to P_1 's boundary
 $d_2 = x$'s distance to P_2 's boundary

Results



Computed Alpha Masks





Without Blending



With Blending


Cannot attenuate the blacks

Aperture Edge Blending



Optical Edge Blending

- Spatial Filter
 - Fabricated in a thermal printer with a photographic dye-set ribbon on a fine grain, high density transparency



Optical Edge Blending

- Engineering feat
 - Movements during insertion of the spatial filter
 - Compensated by intercepting the analog signals from the computer to the projector
- Details are not available

Edge Blending

- Scalable
- Can be used in overlapping configuration only
 - Addresses only overlap variation
- Not enough control (Aperture and Optical)
- No black attenuation (Software)
- Assumes linearity of projector response
- Can get rid of seams entirely if projectors are adjusted to be very similar

Edge Blending

References:


- Lyon Paul, **Edge-blending Multiple Projection Displays On A Dome Surface To Form Continuous Wide Angle Fields-of-View**, Proceedings of 7th I/ITEC, 203-209, 1985. [Software Edge Blending]
- R. Raskar et al, **Seamless Camera-Registered Multi-Projector Displays Over Irregular Surfaces**, Proceedings of IEEE Visualization, 161-168, 1999. [Software Edge Blending]
- K. Li et.al, **Early experiences and challenges in building and using a scalable display wall system**, IEEE Computer Graphics and Applications 20(4), 671-680, 2000. [Aperture Edge Blending]
- C.J. Chen, and M. Johnson , **Fundamentals of scalable high resolution seamlessly tiled projection system**, Proceedings of SPIE Projection Displays VII 4294, 67-74, 2001. [Optical Edge Blending]

Existing Methods

-
- Manipulation of Projector Controls
 - Common Bulb
 - Blending
 - **Gamut Matching**

Gamut Matching

- Use a photometer to capture the color gamut
 - One measurement per projector
- Find the common color gamut that all the projectors can reproduce
- Use linear transformations to achieve the matching



Gamut Matching

- Can be used in abutting configuration only
 - Addresses only inter projector variation
- Not scalable to 40-50 projectors
 - Due to algorithmic complexity

Gamut Matching

References:

- M.C. Stone, [Color balancing experimental projection displays](#), 9th IS&T/SID Color Imaging Conference, 2001.
- M. C. Stone, [Color and brightness appearance issues in tiled displays](#), IEEE Computer Graphics and Applications, 2001.
- G. Wallace, H. Chen, and K. Li, [Color gamut matching for tiled display walls](#), Immersive Projection Technology Workshop, 2003.
- M. Bern and D. Eppstein, [Optimized color gamuts for tiled displays](#), 19th ACM Symposium on Computational Geometry, 2003.
- Aditi Majumder, Zue He, Herman Towles and Greg Welch, [Achieving Color Uniformity in Multi-Projector Displays](#), IEEE Visualization, 2000.
- R. Yang, A. Majumder, M. S. Brown, [Camera-Based Calibration Techniques for Seamless Multi-Projector Displays](#), IEEE Transactions on Visualization and Computer Graphics 11(2), Mar-Apr 2005

Existing Methods

- No comprehensive model
- No single method
 - Flexible
 - Automatic
 - Scalable
- Addresses parts of the problem only
 - Blending : Overlaps
 - Others: Inter Projector Variations
- Intra-projector variation not addressed
- Strict uniformity mindset

Seamless Displays

- Uniformity is not required
- Addressing just the luminance variation can take us a long way
 - Spatial variation in chrominance is negligible
 - Humans are more sensitive to luminance variation than to chrominance variation

Desiderata

- Comprehensive and general framework
 - Addresses intra, inter and overlap variations
 - Design general solutions
 - No special cases
 - Automated
 - Scalable
 - Explain and compare existing methods


Organization

- Properties of Color Variation
- Existing Methods
- Modeling Color Variation
- PRISM


Achieving Color Seamlessness

- To correct, first capture
- Complexity of capture
 - Input color space : 24 bit color
 - Need 2^{24} images
- Reduce complexity by modeling projector color variations


Modeling Color



Modeling Color



Single Pixel, Single Channel Input



- Maximum channel luminance (\mathcal{M}_c)
- Variation in luminance with channel input
 - Channel transfer function ($h_c(i_c)$)

$$C_c(i_c) = h_c(i_c) \mathcal{M}_c$$

Single Pixel, Three Channel Input

At one pixel for one channel:

$$C_c(i_c) = h_c(i_c) \mathcal{M}_c$$

For any input $i = (i_r, i_g, i_b)$:

$$\mathcal{P}(i) = C_r(i_r) + C_g(i_g) + C_b(i_b)$$

With black Offset \mathcal{B} :

$$\mathcal{P}(i) = \begin{aligned} &C_r(i_r) \\ &+ C_g(i_g) \\ &+ C_b(i_b) \\ &+ \mathcal{B} \end{aligned}$$

Single Projector Display

At one pixel for one channel:


$$C_c(i_c, u, v) = \boxed{h_c(i_c)} \quad \mathcal{M}_c(u, v)$$

For any input $\mathbf{i} = (i_r, i_g, i_b)$:

$$\mathcal{P}(\mathbf{i}, u, v) = C_r(i_r, u, v) + C_g(i_g, u, v) + C_b(i_b, u, v)$$

With black Offset \mathcal{B} :

$$\begin{aligned} \mathcal{P}(\mathbf{i}, u, v) &= C_r(i_r, u, v) \\ &+ C_g(i_g, u, v) \\ &+ C_b(i_b, u, v) \\ &+ \boxed{\mathcal{B}(u, v)} \end{aligned}$$



Single Projector Display

At one pixel for one channel:


$$C_c(i_c, u, v) = \boxed{h_c(i_c)} \quad \mathcal{M}_c(u, v)$$

For any input $\mathbf{i} = (i_r, i_g, i_b)$:

$$\mathcal{P}(\mathbf{i}, u, v) = C_r(i_r, u, v) + C_g(i_g, u, v) + C_b(i_b, u, v)$$

With black Offset \mathcal{B} :

$$\begin{aligned} \mathcal{P}(\mathbf{i}, u, v) &= C_r(i_r, u, v) \\ &+ C_g(i_g, u, v) \\ &+ C_b(i_b, u, v) \\ &+ \boxed{\mathcal{B}(u, v)} \end{aligned}$$



Single Projector Display

At one pixel for one channel:


$$C_c(i_c, u, v) = \boxed{h_c(i_c)} \quad \boxed{M_c(u, v)}$$

For any input $i = (i_r, i_g, i_b)$:

$$P(i, u, v) = C_r(i_r, u, v) + C_g(i_g, u, v) + C_b(i_b, u, v)$$

With black Offset B :

$$P(i, u, v) = h_r(i_r)M_r(u, v) + h_g(i_g)M_g(u, v) + h_b(i_b)M_b(u, v) + \boxed{B(u, v)}$$



Single Projector Display


$$P(i, u, v) = \begin{matrix} h_r(i_r) & M_r(u, v) \\ + & h_g(i_g) & M_g(u, v) \\ + & h_b(i_b) & M_b(u, v) \\ + & \boxed{B(u, v)} \end{matrix}$$

Transfer Functions

Luminance Functions

Black Offset

Multi-Projector Display



$$\mathcal{E}(u, v, i) = \sum_{j \in \mathcal{N}(u, v)} \mathcal{P}_j(u, v, i)$$

$$\begin{aligned} \mathcal{E}(i, u, v) = & \sum \hat{h}_r(i_r) \mathcal{M}_r(u, v) \\ & + \sum \hat{h}_g(i_g) \mathcal{M}_g(u, v) \\ & + \sum \hat{h}_b(i_b) \mathcal{M}_b(u, v) \\ & + \sum \mathcal{B}(u, v) \\ & j \in \mathcal{N}(u, v) \end{aligned}$$

Multi-Projector Display

- Intra projector luminance variation
 - $\mathcal{M}_r(u, v)$ and $\mathcal{B}(u, v)$ are not flat
- Inter projector luminance variation
 - $\hat{h}_r(i_r)$ is different
 - $\mathcal{M}_r(u, v)$ and $\mathcal{B}(u, v)$ have different shapes
- Overlap luminance variation
 - \mathcal{N} is different

$$\begin{aligned} \mathcal{E}(i, u, v) = & \sum \hat{h}_r(i_r) \mathcal{M}_r(u, v) \\ & + \sum \hat{h}_g(i_g) \mathcal{M}_g(u, v) \\ & + \sum \hat{h}_b(i_b) \mathcal{M}_b(u, v) \\ & + \sum \mathcal{B}(u, v) \\ & j \in \mathcal{N}(u, v) \end{aligned}$$

Modeling Color Properties

References:

- Aditi Majumder and M. Gopi, **Modeling Color of Multi-Projector Displays**, Computer Graphics Forum, (CGF), 2005.

Organization

- Properties of Color Variation
- Existing Methods
- Modeling Color Variation
- **PRISM: PeRceptual Seamlessness in Multi-Projector Displays**

Luminance Seamlessness

- Luminance uniformity is not required to achieve luminance seamlessness
- Automation
- Unified way of solving all luminance variations: intra, inter and overlap
- Real time correction

Luminance Seamlessness

- Reconstruction
 - Reconstruct E
- Modification
 - Modify E to E'
- Reprojection
 - Reproject E'

Luminance Seamlessness


- Reconstruction
 - Reconstruct \mathcal{E}
- Modification
 - Modify \mathcal{E} to \mathcal{E}'
- Reprojection
 - Reproject \mathcal{E}'

Reconstruction Using Camera


$$\mathcal{E}(i, u, v) = \sum_{r \in \mathcal{N}(u, v)} \hat{h}_r(i_r) \mathcal{M}_r(u, v) + \sum_g \hat{h}_g(i_g) \mathcal{M}_g(u, v) + \sum_b \hat{h}_b(i_b) \mathcal{M}_b(u, v) + \sum_{j \in \mathcal{N}(u, v)} \mathcal{B}(u, v)$$

Transfer Function

Luminance Function



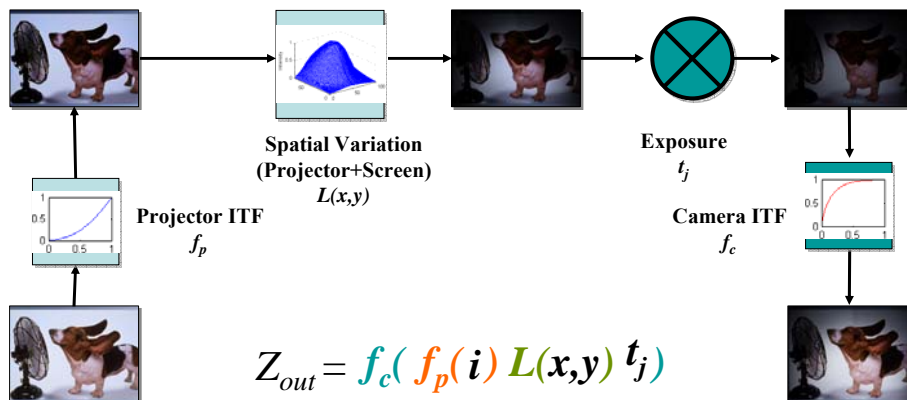
Black Offset



Pre-calibrated Camera

- Known transfer function
 - Use HDR imaging techniques
- Negligible vignetting effect
 - Set to f/32 aperture
- Can we use uncalibrated camera?

Projector-Camera Self Calibration



Algorithm

- Since $f_c(\cdot)$ is monotonic, inverse exists.

$$f_c^{-1}(Z_{out}) = f_p(i) L(x,y) t_j$$

Algorithm

- Take the log of both sides

$$\begin{aligned} f_c^{-1}(Z_{out}) &= f_p(i) L(x,y) t_j \\ \downarrow & \quad \downarrow \quad \downarrow \quad \downarrow \\ \ln f_c^{-1}(Z_{out}) &= \ln f_p(i) + \ln L(x,y) + \ln(t_j) \end{aligned}$$

Algorithm

- Images of multiple projector inputs at multiple camera exposures
- Setup system of equations and solve using least-squares

$$\ln f_c^{-1}(Z_{out}) = \ln f_p(i) + \ln L(x,y) + \ln(t_j)$$

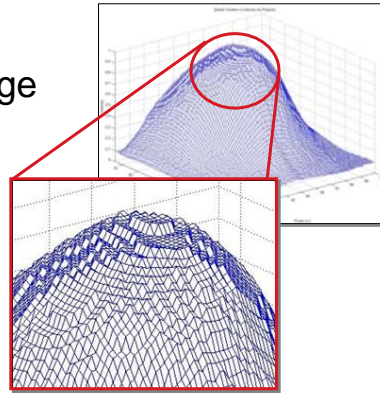
Efficiency Issue

- Imagine a system with:
 - 1024x768 pixel locations
 - 32 colors
 - 6 exposures

 - # of unknowns = $32 + 32 + (1024 \times 768)$
786,000 unknowns
 - # of equations = $1024 \times 768 \times 32 \times 6$
150,000,000 equations

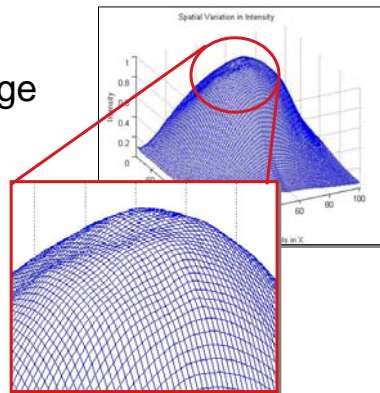
Handling Saturation and Noise

- Saturation Problems
 - Cannot use a single image
- Random Noise
 - Devices
 - Screen



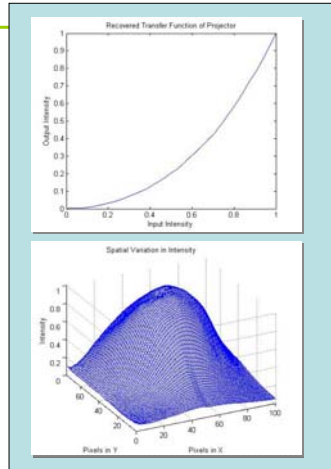
Handling Saturation and Noise

- Saturation Problems
 - Cannot use a single image
- Random Noise
 - Devices
 - Screen




Details in the paper

Estimated Photometric Properties



Projector



Camera

Luminance Seamlessness

- Reconstruction
 - Reconstruct \mathcal{E}
- Modification
 - Modify \mathcal{E} to \mathcal{E}'
- Reprojection
 - Reproject \mathcal{E}'

Modification 1

$$\begin{aligned}
 \mathcal{P}(i, u, v) = & \mathit{h}_r(i_r) \mathcal{M}_r(u, v) \\
 & + \mathit{h}_g(i_g) \mathcal{M}_g(u, v) \\
 & + \mathit{h}_b(i_b) \mathcal{M}_b(u, v) \\
 & + \mathcal{B}(u, v)
 \end{aligned}$$

Single Projector

$$\begin{aligned}
 \mathcal{E}(i, u, v) = & \sum \mathit{h}_r(i_r) \mathcal{M}_r(u, v) \\
 & + \sum \mathit{h}_g(i_g) \mathcal{M}_g(u, v) \\
 & + \sum \mathit{h}_b(i_b) \mathcal{M}_b(u, v) \\
 & + \sum_{j \in \mathcal{N}(u, v)} \mathcal{B}(u, v)
 \end{aligned}$$

Multi Projector

- Match $\mathit{h}_r(i_r)$ of projectors

Modification 1

$$\begin{aligned}
 \mathcal{P}(i, u, v) = & \mathit{h}_r(i_r) \mathcal{M}_r(u, v) \\
 & + \mathit{h}_g(i_g) \mathcal{M}_g(u, v) \\
 & + \mathit{h}_b(i_b) \mathcal{M}_b(u, v) \\
 & + \mathcal{B}(u, v)
 \end{aligned}$$

Single Projector

$$\begin{aligned}
 \mathcal{E}(i, u, v) = & \sum \mathcal{H}_r(i_r) \mathcal{M}_r(u, v) \\
 & + \sum \mathcal{H}_g(i_g) \mathcal{M}_g(u, v) \\
 & + \sum \mathcal{H}_b(i_b) \mathcal{M}_b(u, v) \\
 & + \sum_{j \in \mathcal{N}(u, v)} \mathcal{B}(u, v)
 \end{aligned}$$

Multi Projector

- Match $\mathit{h}_r(i_r)$ of projectors

Modification 1

$$\begin{aligned} \mathcal{P}(i, u, v) = & \color{magenta} h_r(i_r) \mathcal{M}_r(u, v) \\ & + \color{blue} h_g(i_g) \mathcal{M}_g(u, v) \\ & + \color{cyan} h_b(i_b) \mathcal{M}_b(u, v) \\ & + \color{red} \mathcal{B}(u, v) \end{aligned} \quad \text{Single Projector}$$

$$\begin{aligned} \mathcal{E}(i, u, v) = & \color{blue} \mathcal{H}_r(i_r) \sum \mathcal{M}_r(u, v) \\ & + \color{blue} \mathcal{H}_g(i_g) \sum \mathcal{M}_g(u, v) \\ & + \color{blue} \mathcal{H}_b(i_b) \sum \mathcal{M}_b(u, v) \\ & + \sum_{j \in \mathcal{N}(u, v)} \color{red} \mathcal{B}(u, v) \end{aligned} \quad \text{Multi Projector}$$

• Match $h_r(i_r)$ of projectors

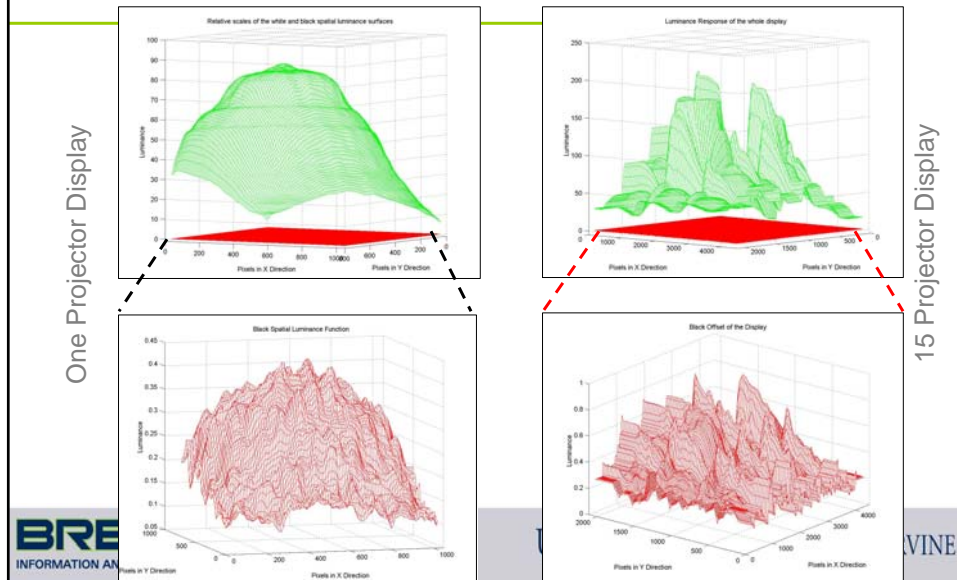
Modification 1

$$\begin{aligned} \mathcal{P}(i, u, v) = & \color{magenta} h_r(i_r) \mathcal{M}_r(u, v) \\ & + \color{blue} h_g(i_g) \mathcal{M}_g(u, v) \\ & + \color{cyan} h_b(i_b) \mathcal{M}_b(u, v) \\ & + \color{red} \mathcal{B}(u, v) \end{aligned} \quad \begin{array}{l} \text{Luminance Functions} \\ \text{of one Projector} \\ \text{Transfer} \\ \text{Function of} \\ \text{one Projector} \end{array} \quad \begin{array}{l} \text{Black Offset of one Projector} \end{array}$$

$$\begin{aligned} \mathcal{E}(i, u, v) = & \color{blue} \mathcal{H}_r(i_r) \sum \mathcal{M}_r(u, v) \\ & + \color{blue} \mathcal{H}_g(i_g) \sum \mathcal{M}_g(u, v) \\ & + \color{blue} \mathcal{H}_b(i_b) \sum \mathcal{M}_b(u, v) \\ & + \sum_{j \in \mathcal{N}(u, v)} \color{red} \mathcal{B}(u, v) \end{aligned} \quad \begin{array}{l} \text{Luminance Functions} \\ \text{of the whole Display} \\ \text{Common} \\ \text{Transfer} \\ \text{Function} \end{array} \quad \begin{array}{l} \text{Black Offset of the whole Display} \end{array}$$

Display is like a single large projector



Display Luminance Functions






Modification 2

- Sharp discontinuities are the cause of photometric seams
- Remove the sharp discontinuities

Strict Luminance Uniformity



Strict Luminance Uniformity




Results




Before

After Strict Photometric Uniformity

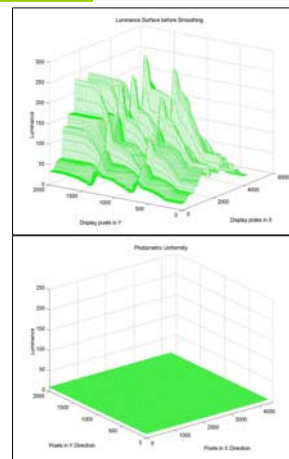


Strict Luminance Uniformity

Suboptimal use of system resources



Significant Contrast/
Dynamic Range
Compression




Results



Before

Which one is better?


After Strict Luminance Uniformity



Optimization Problem

- Smoothing $P(x,y)$ to generate $F(x,y)$
 - Maximize dynamic range $\sum F(x,y)$
 - Smoothing guided by perceptual parameters
 - $F'(x,y) \leq kF(x,y)$, $0.0 \leq k \leq 1.0$
 - Assures no *visible* seams
 - Smoothened profile within the original profile
 - $F(x,y) \leq P(x,y)$
 - Assures with display capability
 - Optimal solution using dynamic programming

Optimization Problem



Strict photometric uniformity is a special case.

Results



Before

After Strict Luminance Uniformity



Results




Before

After Luminance Smoothing





Different Smoothing Parameter (2x2 array of four projectors)


Smooth



Original



Flat



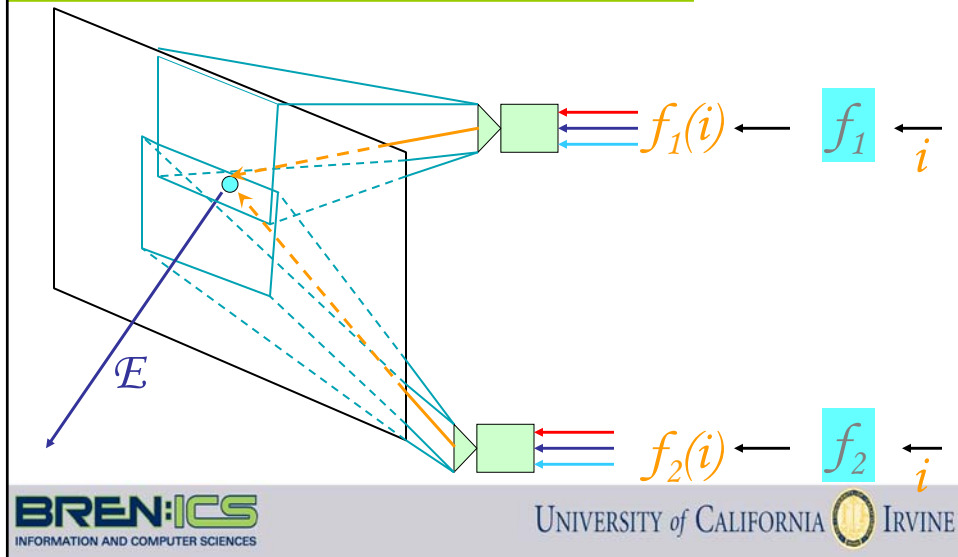
Different Smoothing Parameter (3x5 array of 15 projectors)



Luminance Seamlessness

- Reconstruction
 - Reconstruct \mathcal{E}
- Modification
 - Modify \mathcal{E} to \mathcal{E}'
- Reprojection
 - Reproject \mathcal{E}'


Reprojection



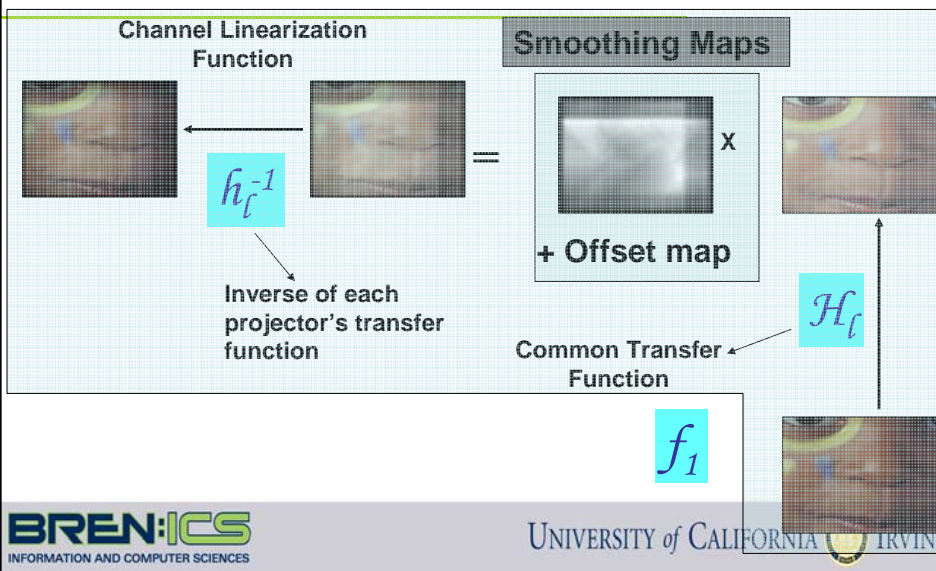
Smoothing Maps


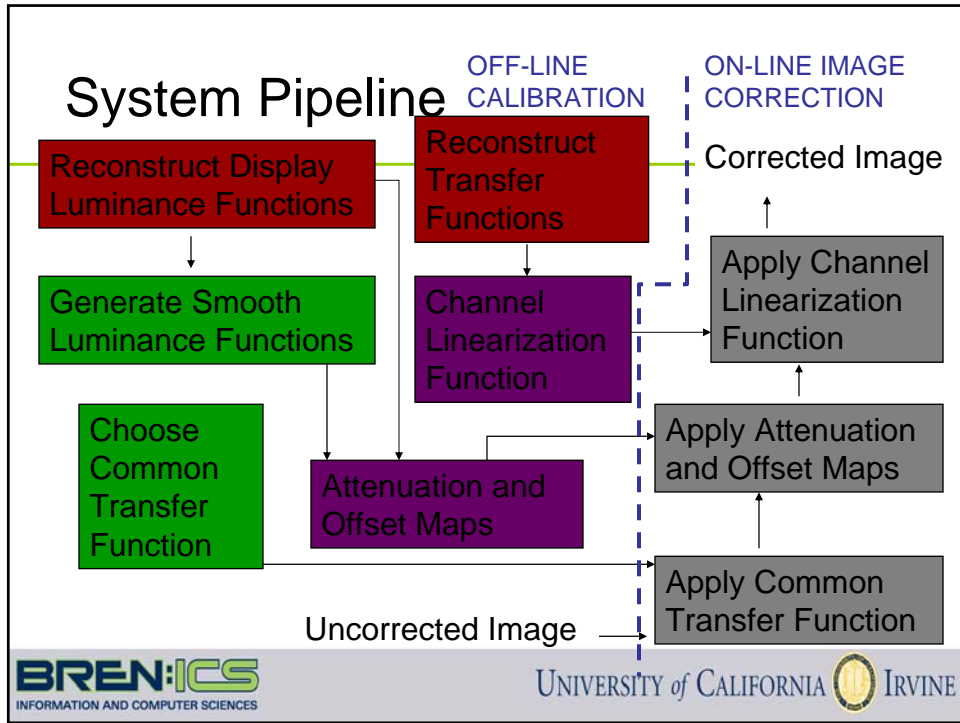
- Attenuation Map
 - Per pixel luminance attenuation to achieve the desired luminance function
- Offset Map
 - Per pixel luminance offset to achieve the desired black offset

Attenuation Map



Per Projector Image Correction






Results (After)



BRENDS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

Results (Before)

15 Projector Display



BRENDS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

Results (After)



BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

PRISM

References:

- Aditi Majumder and Rick Stevens, [LAM: Luminance Attenuation Map for Photometric Uniformity Across Projection Based Displays](#), ACM Virtual Reality Software and Technology, 2002.
- Aditi Majumder and Rick Stevens, [Perceptual Photometric Seamlessness in Tiled Projection Based Displays](#), ACM Transactions on Graphics, Vol. 24, No. 1, 2005.
- Aditi Majumder, [Improving Contrast of Multi-Displays Using Human Contrast Sensitivity](#), International Conference on Computer Vision and Pattern Recognition (CVPR), 2005.
- R. Juang, A. Majumder, [Photometric Self-Calibration of a Projector Camera Pair](#), IEEE CVPR Workshop on Projector Camera Systems, 2007.
- E. Bhasker, R. Juang, A. Majumder, [Registration Techniques for Using Imperfect and Partially Calibrated Devices in Planar Multi-Projector Displays](#), IEEE Visualization, 2007.

BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

Summary

Method	Addresses..	
Controls Manipulation	Inter	Luminance and Chrominance
Common Bulb	Inter	Luminance and Chrominance
Edge Blending	Overlap	Luminance Only
Gamut matching	Inter	Luminance and Chrominance
PRISM	Intra + Inter + Overlap	Luminance Only

Future Work (Handling Chrominance)

Before



Future Work (Handling Chrominance)

After



BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

Overview


- Geometric Registration
- Photometric Registration
- PC Cluster Based Rendering
 - Figures: Courtesy Michael S. Brown
- Distributed Rendering

BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

Traditional Rendering Architecture


- Multi-pipe single machine
 - Each pipe driving one projector
- Advantages
 - Shared Memory
 - Shared Disk Space



PC Cluster Based Rendering

- Graphics card has similar performance
- PC can drive only a single display
 - One PC per projector
- Problem
 - Multiple PCs for the entire display
 - Content needs to be distributed

Distributed Rendering Framework



PC Cluster Rendering Solutions

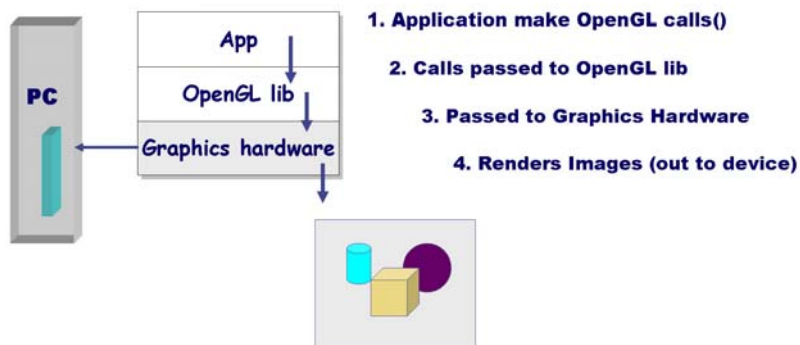
- WireGL
- Chromium
- VR Jugglers
- All use PC cluster + network to render a large “logical” framebuffer
 - Rendering is synchronized via the network

Chromium

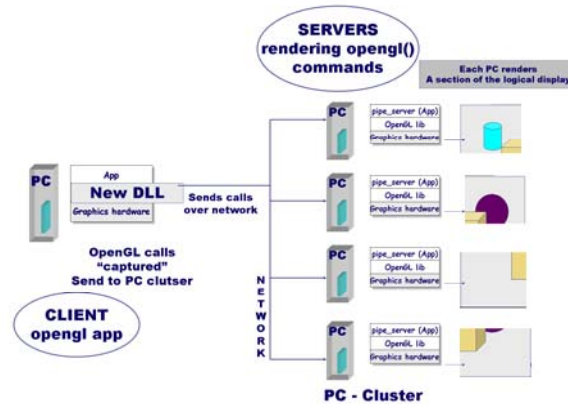
- Designed to support OpenGL API
 - No change to existing OpenGL applications
- Each PC renders a logical *tile*
- Tiles can overlap completely, partially or none
- Well suited for our application
 - Each PC drives a projector
 - Has partial overlap

Basic Rendering on Single PC

Single PC - Running an OpenGL app




Chromium intercepts OpenGL calls



Defining tiles

- Configuration file
- Virtual framebuffer
 - Top left is (0,0)
- Tile: (offset, size)
- Overlapping tiles
 - Rectangular!!




USER DEFINED CONFIG FILE

SERVER & POSITION IN LOGICAL DISPLAY

1	Server1.xyz.org	Position 0 0	size (1024 768)
2	Server2.xyz.org	Position 1000 0	size (1024 768)
3	Server3.xyz.org	Position 0 700	size (1024 768)
4	Server4.xyz.org	Position 1000 700	size (1024 768)

Desired Functionality




We would prefer “casual” alignment
Remove the need for precise alignment



Non-rectangular projectors

- Tile (T) inscribes the non-rectangular projector region (P)
- Generate new configfile

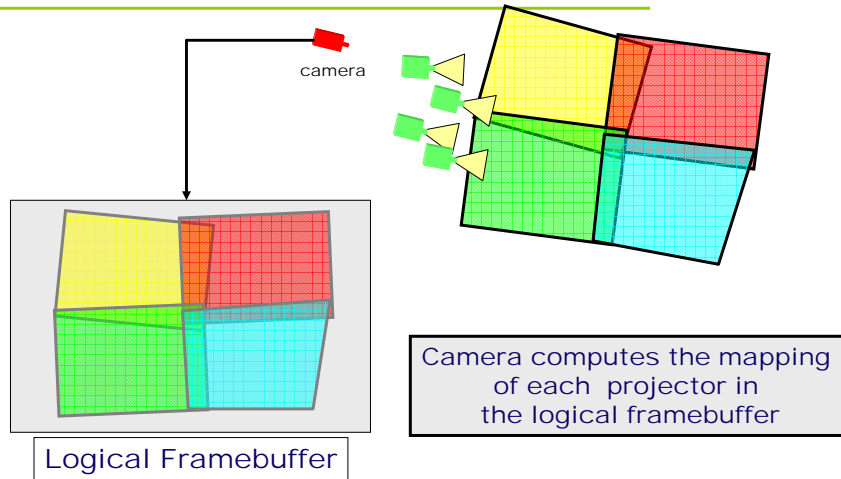


AUTOMATICALLY GENERATED CONFIG FILE

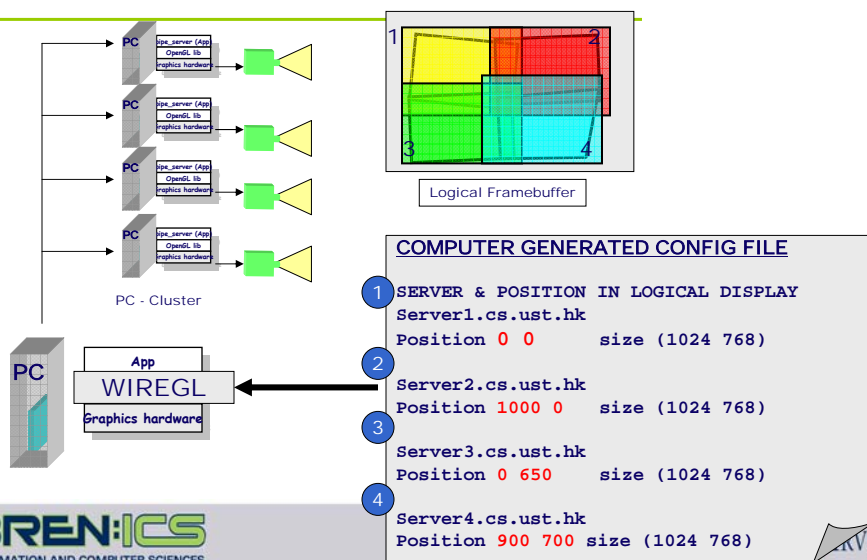
```

SERVER & POSITION IN LOGICAL DISPLAY
Server1.xyz.org
1 Position 0 0 size (1024 768)
Server2.xyz.org
2 Position 1000 0 size (1024 768)
Server3.xyz.org
3 Position 0 650 size (1024 768)
Server4.xyz.org
4 Position 900 700 size (1024 768)
    
```

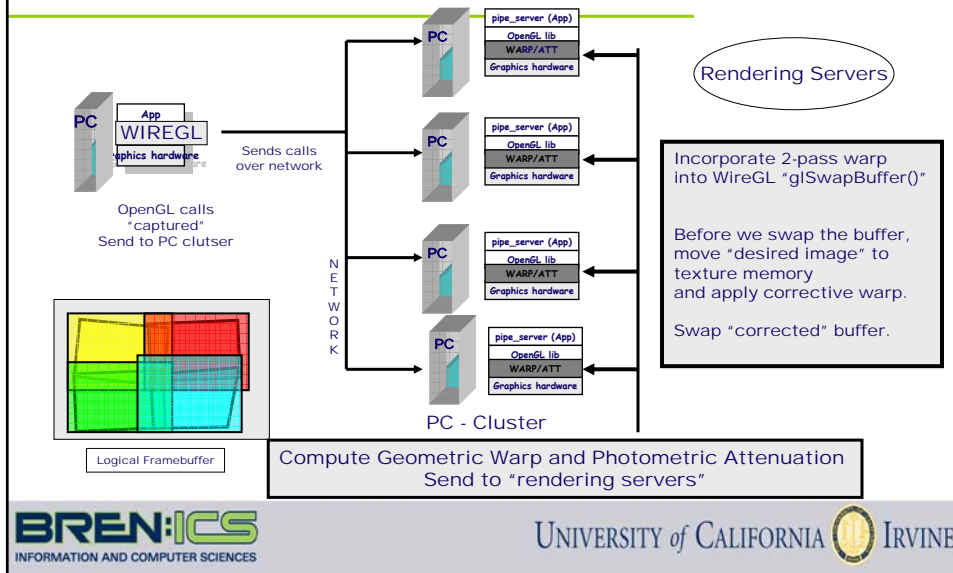
Can be found automatically




Generate the Configuration file



Modification to rendering servers



Correction



Examples



Example Deployment



Display on Curved Surface



Examples




Display Running Quake III



Reconfigured in less than one minute



Also Possible



PC Based Rendering


References:

- A. Majumder, M. S. Brown, [Practical Multi-Projector Display Design](#), A.K. Peters, 2007.
- R. Yang, A. Majumder, M. S. Brown, [Camera-Based Calibration Techniques for Seamless Multi-Projector Displays](#), IEEE Transactions on Visualization and Computer Graphics 11(2), Mar-Apr 2005
- G. Humphreys, P. Hanrahan, [A Distributed Graphics System for Large Tiled Displays](#), IEEE Visualization, 1999.
- G. Humphreys, M. Eldridge, I. Buck, G. Stoll, M. Everett, P. Hanrahan, [WireGL: A Scalable Graphics Systems for Clusters](#), SIGGRAPH 2001.
- G. Humphreys, M. Houston, R. Ng, R. Frank, S. Ahem, P. Kirchner, J. Klosowski, [Chromium: A Stream Processing Framework for Interactive Rendering on Clusters](#), ACM Transactions on Graphics, 2002.

Overview

- Geometric Registration
- Photometric Registration
- PC Cluster Based Rendering
- Distributed Rendering

Centralized Architecture



Centralized Server must use synchronized push

Limitations of Centralized Approach

- Educated User
 - Difficult to deploy
- Not easy to add/remove projectors
 - Not scalable (Limited by camera resolution)
- Not easy to rearrange projectors
 - Not reconfigurable
- Not easy to tolerate faults


Imagine...

- A display that can calibrate itself with no user intervention
- Can detect addition/removal and recalibrate itself
- Can detect faults and function at a limited capability

Distributed Approach

- Plug-and-Play Projector (PPP)
- Distributed Architecture
- Asynchronous Distributed Calibration

Plug-and-Play Projectors (PPP)



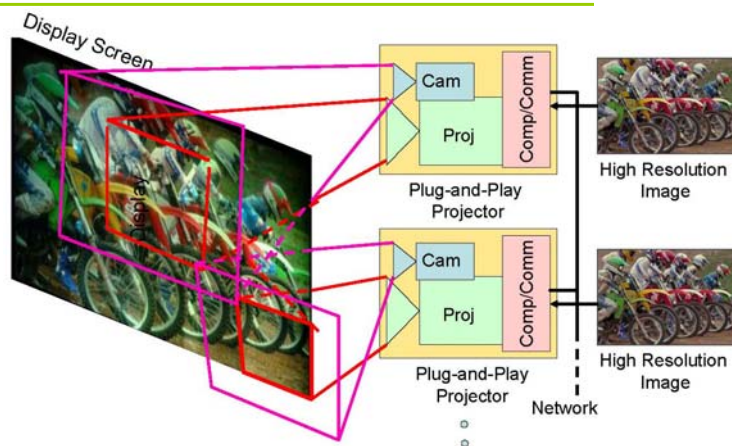
Projector, Camera, Wireless Unit, Embedded Computation Unit
(Inspired by Rasker '03)

Plug-and-Play Projectors (PPP)



Our Prototype

Distributed Architecture



Distributed Architecture

- Data is pulled by each PPP
- Data server does not know that these are displays
 - Acts like any other data client
- Each PPP manages its own pixels

Asynchronous Distributed Calibration

- Each PPP runs asynchronous SPMD algorithm
 - Each PPP discovers its neighbors
 - PPPs discovers the array configuration
 - Using camera-based-communication
 - Self-calibrates accordingly
 - Scalable
 - Reconfigurable
 - Fault-Tolerant

Initially...



BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

SPMD Algorithm

- Neighbor Discovery
- Configuration Identification
- Alignment


BREN:ICS
INFORMATION AND COMPUTER SCIENCES

UNIVERSITY of CALIFORNIA  IRVINE

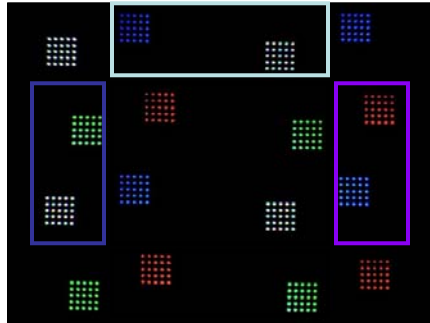
SPMD Algorithm

- Neighbor Discovery
- Configuration Identification
- Alignment

Projected Pattern




Pattern seen by Cameras



From a camera of a PPP with all four neighbors

From a camera of a PPP at the top-left corner of the display



After Neighbor Discovery






- Each PPP knows
 - The number of neighbors it has
 - Their location relative to self (top, bottom, etc.)
- But does not know
 - Total number of projectors
 - Projection array configuration
 - Its own coordinates in the array

SPMD Algorithm

- Neighbor Discovery
- Configuration Identification
- Alignment

Communication Pattern

- Binary-encoded cluster of blobs

Row		Row 1
Column		Column 3
Total Rows		2 Total Rows
Total Columns		3 Total Columns
Status Bits		Not Complete

- Neighbors update beliefs by detecting patterns
- Several rounds of such local updates
 - Parameters diffuse to all PPPs
 - Asynchronously converge to correct global values

Enabling Low Bandwidth Network Communication

- Only camera-based communication till now
- PPPs need to know the IP addresses of its neighbors
- Each PPP broadcasts its IP address and coordinates

After Configuration Identification

- Each PPP knows
 - Total size of display
 - The part of the display it is responsible for
 - IP address of neighbors
- But does not know
 - The relative orientation of its neighbor to warp the image to make a seamless display

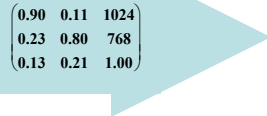
After Configuration Identification



SPMD Algorithm


- Neighbor Discovery
- Configuration Identification
- Alignment

Geometric Alignment




0.90	0.11	1024
0.23	0.80	768
0.13	0.21	1.00

- Local Geometric Calibration
 - Find correspondences
 - Hungarian Method
 - No Pattern Sequences
 - Find local homographies
- Global Geometric Calibration
 - Distributed Homography Tree



Distributed Homography Tree



Photometric Seamlessness

- Local Overlap Blending
 - Addresses only overlap brightness variation
 - No distributed color calibration yet

After Geometry and Color Alignment



Distributed Registration

References:

- E. Bhasker, P. Sinha, A. Majumder, *Asynchronous Distributed Calibration for Scalable and Reconfigurable Multi-Projector Displays*, IEEE Visualization, 2006.

Primary Reference

- Most common issues
 - In lucid terms
- Many Examples
- Sample code for PC cluster rendering





Visually Augmenting the real World with Projectors

Oliver Bimber

Bauhaus-University Weimar



Parts of these slides have been used for a Eurographics'07 State of the Art Report presentation:

"The Visual Computing of Projector-Camera System"

They have been created in cooperation with
Daisuke Iwai, Osaka University
Gordon Wetzstein, University of British Columbia
Anselm Grundhöfer, Bauhaus-University Weimar


Technical details and further information are provided in the accompanying article:

Bimber, O., Iwai, D., Wetzstein, G. and Grundhöfer, A.
The Visual Computing of Projector-Camera Systems
Eurographics State of the Art Report, 2007
To appear in Computer Graphics Forum, 2008

This article is reproduced by kind permission of the Eurographics Association
(c) Eurographics Association 2007

Introduction and Motivation

Projection onto Non-Optimized Surfaces



Bimber et al, IEEE Computer 2005

← normal projection

← corrected projection

...but why?

...sometimes screens are just not possible!

Applications

Theatres



Karl May Festival Eispe, 2006



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Theatres



Karl May Festival Elspe, 2006

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Advertisement



Moritz Immobilien Leinefelde, 2007

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Advertisement



Moritz Immobilien Leinefelde, 2007

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Advertisement



World of Events, 2007

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



Historic Sites



Henrichshütte Hattingen, 2007



Osterburg Weida, 2007



Historic Sites



Dechen Cave Iserlohn, 2007



Museums



Watermuseum Aquarius Mülheim, 2007



Open Air Movies



Castel Klaffenbach, Chemnitz, 2007




Open Air Movies



Castel Klaffenbach, Chemnitz, 2007



Festivals



Abbey Neumünster, Luxembourg, 2007



Festivals



Abbey Neumünster, Luxembourg, 2007



Ad Hoc VR



Architecture Faculty, Bauhaus-University Weimar, 2006

More Possibilities

Paintings





augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Fossils



Senkenberg Museum Frankfurt, 2004



Senkenberg Museum Frankfurt, 2005

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Projector-Based AR Visualization



IEEE/ACM ISMAR, 2007, www.uni-weimar.de/medien/AR/sARc

Bimber, et al.,
IEEE/ACM ISMAR 2005

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



Television Studios




Projector-Based AR Games





Mobile Devices



Challenges of Non-Optimized Surfaces

augmented reality [erweiterte Realität] Bauhaus-Universität Weimar

Smart Projection with Projector-Camera Systems

The diagram illustrates a projector-camera system. A projector (P) and a camera (C) are positioned to project and capture an image of a scene. An eye (O) observes the scene. A light source (E) is also shown. The diagram includes a green dashed line representing a path or boundary, a small inset image of a stone wall, and two side-by-side images: 'original image' and 'observed projection'.

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

augmented reality [erweiterte Realität] Bauhaus-Universität Weimar

Smart Projection with Projector-Camera Systems

The diagram illustrates a projector-camera system. A projector (P) and a camera (C) are positioned to project and capture an image of a scene. A light source (E) is also shown. The diagram includes a green dashed line representing a path or boundary, a small inset image of a stone wall, and two side-by-side images: 'original image' and 'observed projection'.

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Smart Projection with Projector-Camera Systems

The diagram illustrates the setup for smart projection. A projector (E) projects light onto a surface (P). A camera (O) captures the scene. The diagram shows the original image and the observed projection.

original image observed projection

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Local Diffuse Reflections

The image shows a close-up of a stone wall with local diffuse reflections. The reflections are visible on the surface of the stones, showing the texture and color of the wall.

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008



Local Diffuse Reflections



normal projection



Local Diffuse Reflections



radiometric compensation




Local Diffuse Reflections



misaligned projection




Limited Dynamic Range and Brightness






Global Scattering




Global Scattering






Global Scattering



Defocus





Other Complex Modulations



Other Complex Modulations





What this Part will not cover...



Calibration for Tiled Screens



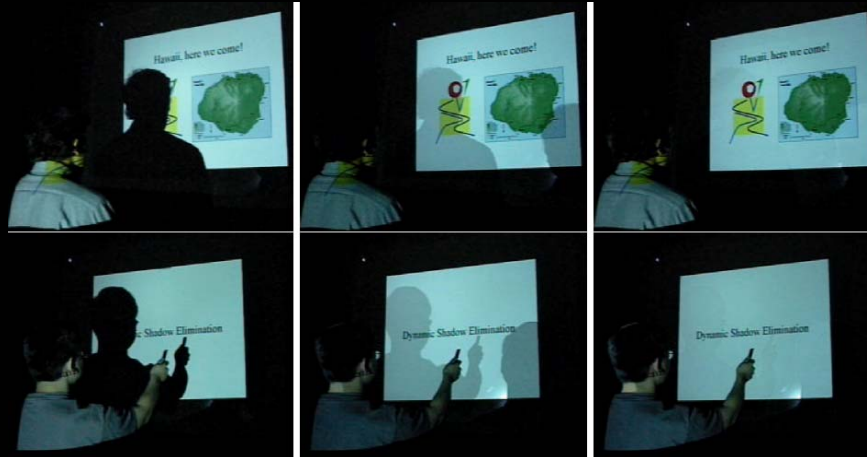
Courtesy: Brown, et al.,
IEEE TVCG, 2005

Book: Majumder and Brown,
Practical Multi-Projector
Display Design, AK Peters 2007






Real-Time Shadow Removal



Shadow Removal
Courtesy: Rahul Sukthankar, Compaq (many similar approaches)



Non-Trivial, Uniformly Colored Surfaces



Similar:

Yang et al.,
WSCG01 (evaluate
deformation of
images to
reconstruct
geometry, projector
and camera are
calibrated)


Courtesy: Johnson, et al., ProCams 2007



Textured Surfaces



Pixel-Displacement Mapping







Radiometric Compensation

Single Projector

determining parameters (textures):

- (1) turn off environment light and project black flood image

$$I=0, E=0 \rightarrow \text{BFM}$$

- (2) turn on environment light and project black flood image

$$I=0, E=1 \rightarrow \text{EM (incl. BFM !)}$$


- (3) turn off environment light and project white flood image

$$I=1, E=0 \rightarrow \text{FM (incl. BFM !)}$$

$$\rightarrow \text{FM} = \text{FM} - \text{BFM}$$

compensation (per pixel):

$$I = (R - EM) / (FM)$$



$$R = IFM + EM$$

I → projected image

B → black-level

F → projector-to-surface form factor

E → environment light

M → surface reflectance (diffuse)

Multiple Projectors

strategy: balance intensity load

- assume: total intensity is equally balanced among multiple low-capacity units

$$I_1 = I_2 = I_3 = \dots = I_N$$

- this is equivalent to the assumption that a single high capacity projector produces the total intensity arriving on the surface virtually


$$R = EM + I_1(F_1M + F_2M + \dots + F_NM)$$

$$\rightarrow EM + I(F_1 + F_2 + \dots + F_N)M$$

compensation (per pixel):

$$I_i = (R - EM) / (F_1M + F_2M + \dots + F_NM)$$

remember: $F_iM = F_iM - B_iF_iM$!
 or $\text{BFM} = B_1F_1M + \dots + B_iF_iM$



$$R = EM + I_1F_1M + I_2F_2M + \dots + I_NF_NM$$

Color Mixing

Nayar et al, ProCams, 2003
 Yoshida et al, VSMM, 2003 (3x4 color mixing matrix)
 Grossberg et al, CVPR, 2004

determining color mixing matrix V :


for un-normalized matrix: (camera and projector response must be known and linearized):

capture 9+ images \rightarrow least squares
 3x4 color mixing matrix: last column \rightarrow EM

for normalized matrix (camera response must be known, projector response can be unknown):

diagonals are 1 (unknown scaling)
 off-diagonals are $V_{ij} = \Delta C_j / \Delta I_i = \Delta C_j / \Delta R_i$
 (since $V_{ii} = 1, \Delta I_i = \Delta C_i$)

capture 6 images C (2 per color channel to determine deltas)



compensation (per pixel):

$$I = V^{-1} * R \text{ (does not consider environment light!)}$$

$$R = V * I$$

$I \rightarrow$ projected image
 $V \rightarrow$ color mixing matrix (projector/camera/reflectance)

Dynamic Adaptation

Fujii et al, CVPR, 2005

determining color mixing matrix V_0 :

similar as before: $V_{ij} = \Delta C_j / \Delta I_i$
 (un-normalized!)

determine reflected environment

light $E_0 * M_0$ at $t=0$:


$E_0 * M_0 = C - V_0 * I$ (project arbitrary I and capture C)


compensation (per pixel at t):

$I_t = V_0^{-1} * (R * M_0 / M_{t-1} - E_{t-1} * M_0)$
 $\rightarrow E_{t-1} * M_0$ approx. $E_0 * M_0$
 $\rightarrow M_0 / M_{t-1} = C_0 / C_{t-1}$

$$R_t = M_t / M_0 * (E_t * M_0 + V_0 * I_t)$$

$t \rightarrow$ time index
 $I_t \rightarrow$ projected image at t
 $V_0 \rightarrow$ un-normalized color mixing matrix at $t=0$ (const.)
 $M_t \rightarrow$ material at t
 $M_0 \rightarrow$ material at $t=0$
 $E_t \rightarrow$ environment light ($E_t = E_0$)





Adaptive Radiometric Compensation



Content Dependent Compensation

Wang, ProCams 2005
(only gray scale)


- solution for gray scale image content and surfaces only
- calculation of maximum non-perceivable luminance threshold (threshold map)
- error minimization by reducing global image contrast until perceivable clipping errors disappear



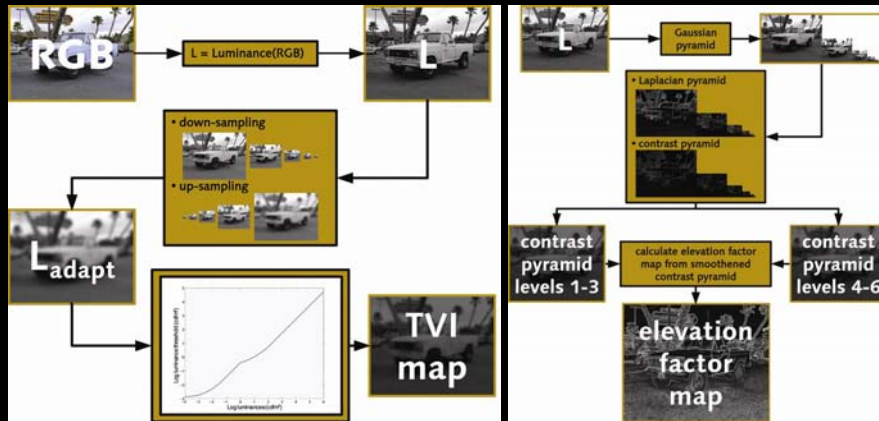
Content Dependent Compensation

Ashdown et al, ProCams, 2006 (only static)

- analyze and manipulate the image content to guarantee an optimized display quality
- conversion into perceptually-uniform CIE Luv color space
- chrominance fitting into the gamut of the projector
- automatic spatially varying luminance matching into displayable range with a relaxation method used to solve elliptic partial differential equations

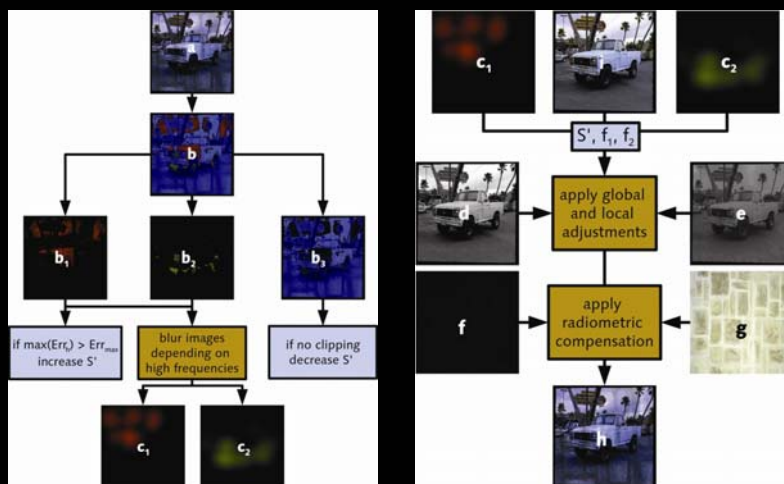


Real-Time Adaptive Radiometric Compensation



Grundhöfer and Bimber, Siggraph Poster 2006, IEEE TVCG (to appear 2008)


Real-Time Adaptive Radiometric Compensation



Grundhöfer and Bimber, Siggraph Poster 2006, IEEE TVCG (to appear 2008)



Real-Time Adaptive Radiometric Compensation



Grundhöfer and Bimber, Siggraph Poster 2006, IEEE TVCG (to appear 2008)



Karl May Festival Elspe, 2006





normal projection



radiometric compensation




misaligned projection



Embedded Structured Light



Structured Light Projection





Temporal Coded Projection

Raskar et al, Siggraph'98
(temporal coding of binary information)




**Bauhaus-University
Weimar**




Color and Intensities


Bimber et al, J. VRB 2006



recorded image



compensation image



perceptible projection




contrast enhanced recording



Perceptible Temporal Coded Projection

Waschbüsch et al. :
Scalable 3D Video of
Dynamic Scenes. The
Visual Computer 21
2005



Viera et al.:
A Camera-Projector System for
Real-Time 3D
Video. ProCams 2005



Artefacts during Eye Motions





Artefacts during Eye Motions



← eye motion




Artefacts during Code Transitions



← time
1 0 1 0 1




Artefacts during Code Transitions




Adaptive Temporal Coded Projection

Grundhöfer, Seeger,
Häntsch, Bimber,
ISMAR 07

deriving local coding
contrast from
contrast sensitivity
parameters (local
spatial F and L) in
code and display
images



0
Delta
(% of
intensity)



10
Blending
Steps
(frames)

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Zollmann and Bimber, Eurographics 2007

Imperceptible Calibration for Radiometric Compensation

binary points (quick geometry) flat image (reflectance) binary Gray code (surface discontinuities) phase shifted cosine (precise geometry)

(b) (c) (d) (e)

(f) (g) (h) (i)

rough calibration (1.7 seconds) pixel-precise calibration (6.4 seconds)

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Zollmann and Bimber, Eurographics 2007

Imperceptible Calibration for Radiometric Compensation

camera image → ← projection

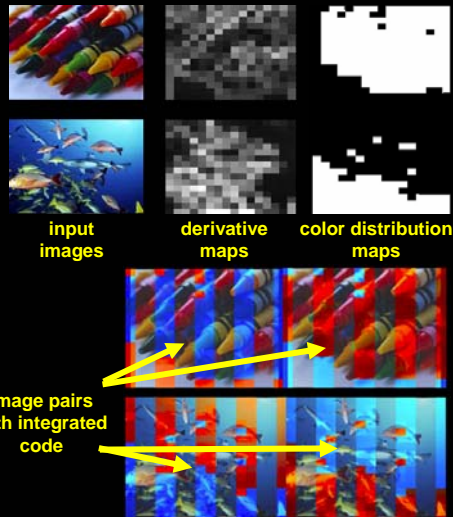
O. Bimber Visually Augmenting the real World with Projectors 17/05/2008



Adaptive Temporal Coded Projection


Park et al., HCI International 2007

- image dependent adaptation of embedded code intensity
- calculation of local spatial variations in the input image
- chrominance analysis in the YIQ color space
- integration of code pattern into RED or BLUE color channel depending on spatial variation and maximum values in the IQ chrominance channels



Binary Image Exposure Period

Cotting et al., ISMAR 2004 (spatial coding)






Complex Light Modulations

Specular Reflection Elimination

- depending on viewer-position, select projector pixel that produces less specular reflections at common surface spot
- black-level cannot be corrected (remaining hotspots)



Global Scattering

Seitz et al., ICCV 2005
 Bimber et al., IEEE VR + Tech Rep. BUW, 2006
 Mukaigawa et al., VRST 2006
 Habe et al, ProCams 2007


$$B_i = E_i + M_i \sum_{j,j=1}^P B_j F_{ji}$$

In matrix notation:


Radiosity: $B = E(I - F)^{-1}$
 E=original, B=original+scattering

Reverse Radiosity: $B(I - F) = E$
B=original, E=compensated
 $\rightarrow (I - F) = \text{const.}$

Reverse Radiosity




Global Scattering



Bimber et al, IEEE VR 2006




Global Scattering



Bimber et al, IEEE VR 2006




Global Scattering



Bimber et al, IEEE VR 2006



Global Scattering



Bimber et al, IEEE VR 2006




Global Scattering



Bimber et al, IEEE VR 2006




Global Scattering



Bimber et al, IEEE VR 2006



Measuring Impulse Scatter Functions (ISF)




measurement of ISF with laser pointer, cancellation of interreflections in images

Seitz et al., ICCV 2005



measurement of ISF with projector, cancellation of interreflections in projection



Habe et al, ProCams 2007



Nayar, et al, Siggraph'06

separation of direct and indirect illumination in images via high-frequency (projected) lighting

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Defocus

Bimber et al, IEEE TVCG 2006

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

augmented reality [erweiterte realität] Bauhaus-Universität Weimar

Defocus

Bimber et al, IEEE TVCG 2006

O. Bimber Visually Augmenting the real World with Projectors 17/05/2008

Defocus

Zhang et al, Siggraph 2006
(similar: Brown, CVPR, 2006) and Oyamada, ProCams 2007

- measure spatially varying defocus kernel using projected structured light




- radiance I is given by

$$I = \alpha f(x; z) * P(x) + \beta$$

surface reflectance
defocus kernel
illumination pattern
environment light

- compute compensation image by solving with respect to P

$$P^+ = (\alpha f)^{-1} * (I - \beta)$$



Coded Aperture Projection

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR



Coded Aperture Projection

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR



Coded Aperture Projection

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR



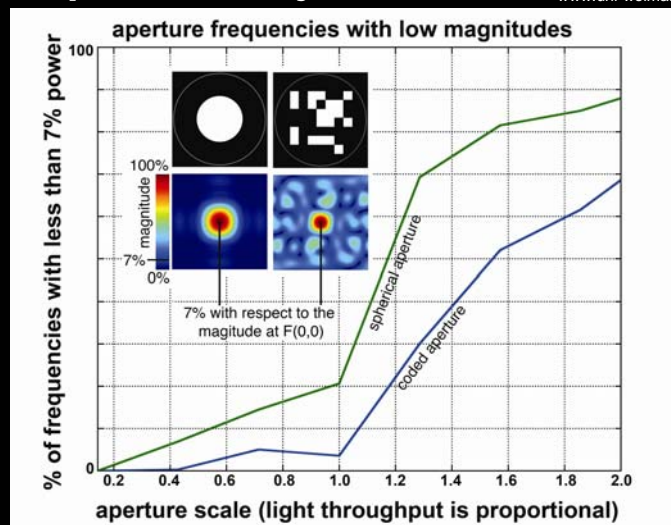
Coded Aperture Projection

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR




Coded Aperture Projection

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR



Increasing Depth of Field




Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR



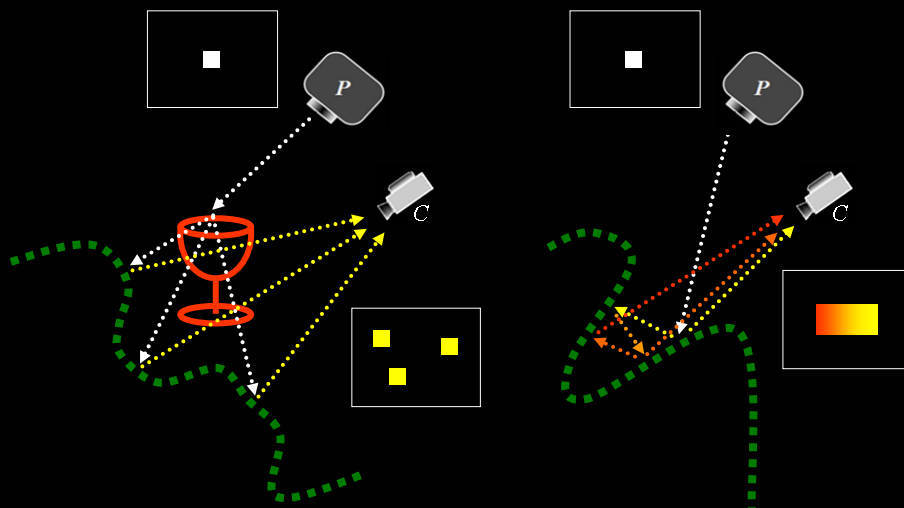
Depixelation

Grosse and Bimber
ACM Siggraph 2008 (poster)
Tech. Report available at
www.uni-weimar.de/medien/AR





One Technique for all...?




no direct mapping: refractions

no direct mapping: inter-reflections

Complex Modulations




Measuring Light Transport






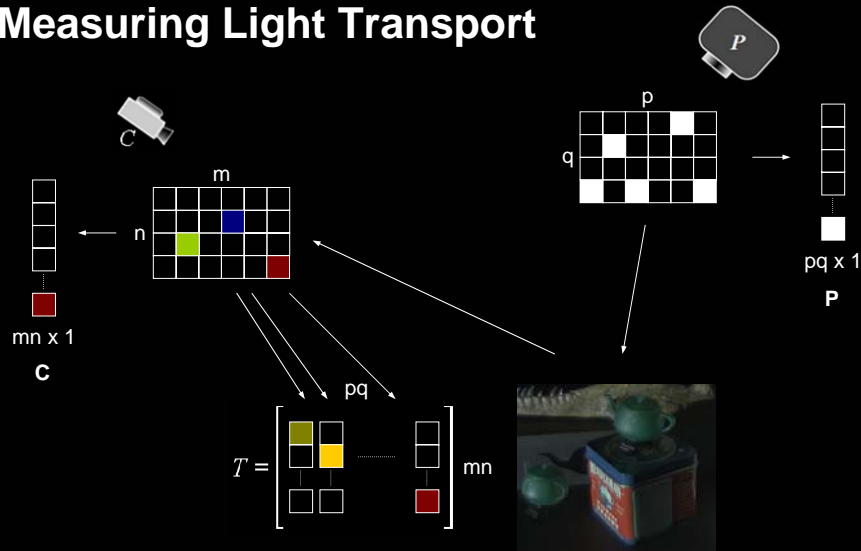
Measuring Light Transport




Measuring Light Transport



Measuring Light Transport




Measuring Light Transport



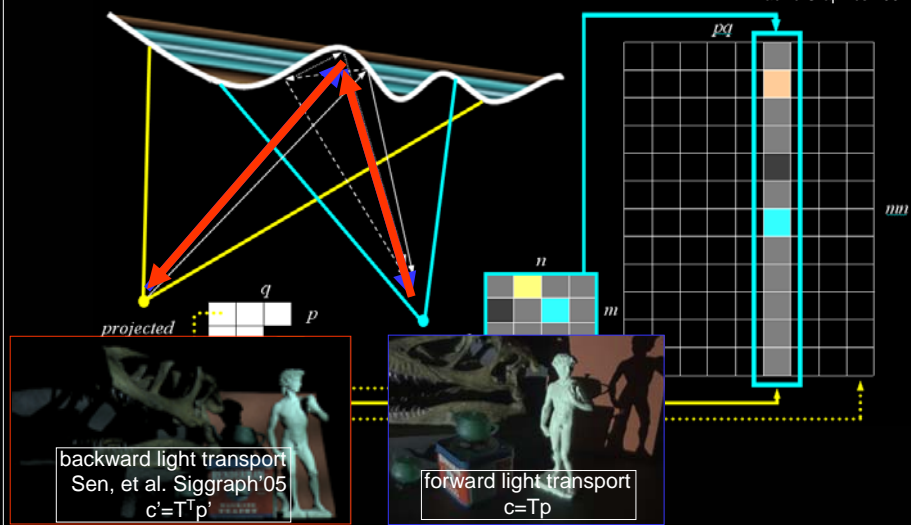
Measuring Light Transport

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007



Inverse Light Transport

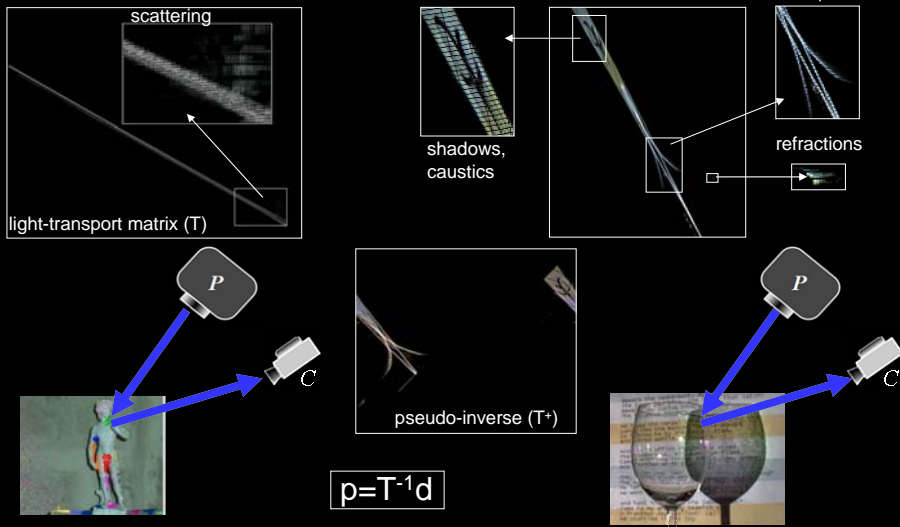
Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007






Inverting Light Transport

Wetzstein and Bimber, Siggraph 2006, Pacific Graphics 2007



Clustering and HW Acceleration

Wetzstein and Bimber, Siggraph 2006, Pacific Graphics 2007




Refractions

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007




Local Diffuse Reflections

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007



Scattering and Inter-Reflections

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007



Scattering and Inter-Reflections

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007

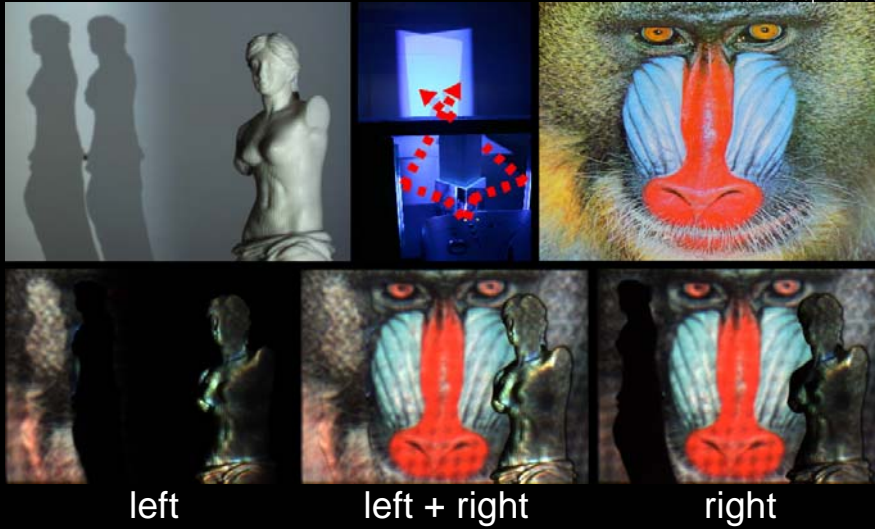


shadows cannot be compensated with single projector



Shadows

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007



Defocus

Wetzstein and Bimber,
Siggraph 2006,
Pacific Graphics 2007







Overcoming Technical Limitations




low resolution projector and camera



low dynamic range projector and camera



low speed projector and camera



Zollmann and Bimber, Eurographics, 2007



super-resolution capturing technique



Courtesy: Tokyo Institute of Technology

high dynamic range Imaging technique




high speed capturing system



Wilburn et al., CVPR, 2004




super-resolution projection




Damera-Venkata et al., ProCams, 2007

high dynamic range projection





Kusakabe et al., IDW, 2006

high speed projection




Jones et al., CVMP, 2006




Super-Resolution Projection


- super-resolution projector-camera systems improve the accuracy of geometric warping and consequently have the potential to enhance radiometric compensation
- 4096 x 2160 pixel projector is commercially available for cinema use
- methods to enhance resolution of usual off-the-shelf projector
 1. **single projector** with special optical actuator (Wobulation)
 2. **multi-projector** approach




Wobulation



Wobulation



Wobulation

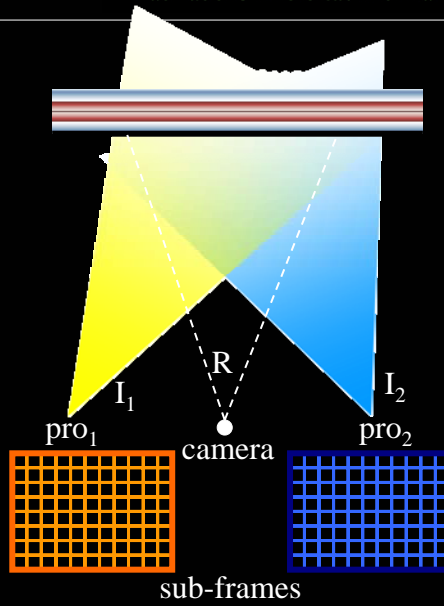


Allen, et al., SID, 2005

Multi-Projector



Jaynes, et al., ProCams, 2003
Damera-Venkata, et al., ProCams, 2007





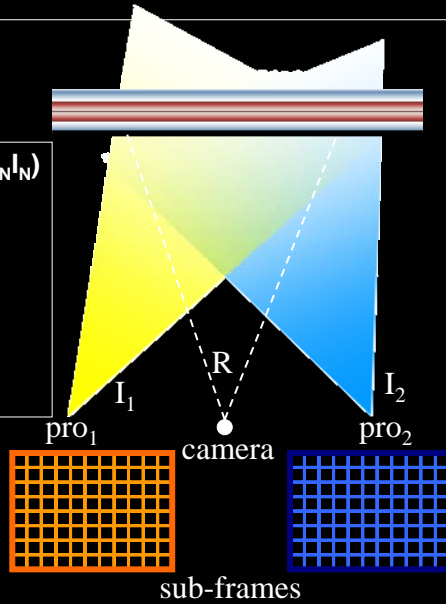
Multi-Projector

$$R = EM + (A_1 V_1 I_1 + A_2 V_2 I_2 + \dots + A_N V_N I_N)$$

- R → captured image
- I_i → projected image
- A_i → geometric warping matrix
- V_i → color mixing matrix
- E → environment light
- M → surface reflectance (diffuse)

Note:

- R, EM, I_i are images
- A_i, V_i transform the whole image



Jaynes, et al., ProCams, 2003
 Damera-Venkata, et al., ProCams, 2007



Multi-Projector

global optimization

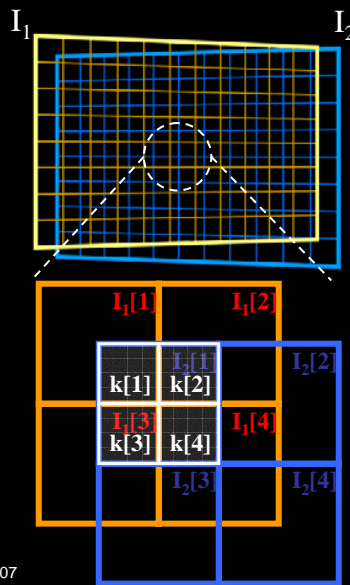
- goal: to find the intensities and colors of corresponding projector pixels in I₁ and I₂ that approximate k as close as possible by assuming that the perceived result is I₁+I₂

$$I_i = \underset{I_i}{\operatorname{argmin}} \|O - R\|^2$$

- this can be solved numerically using an iterative gradient descent algorithm

real-time sub-frame rendering

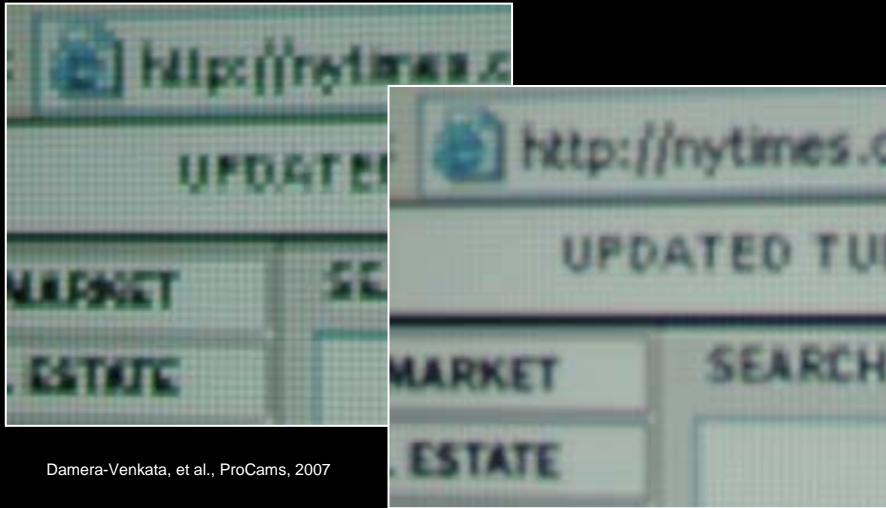
- near-optimal results can be produced with a non-iterative approximation
- implemented with fragment shader



Damera-Venkata, et al., ProCams, 2007



Multi-Projector



Damera-Venkata, et al., ProCams, 2007

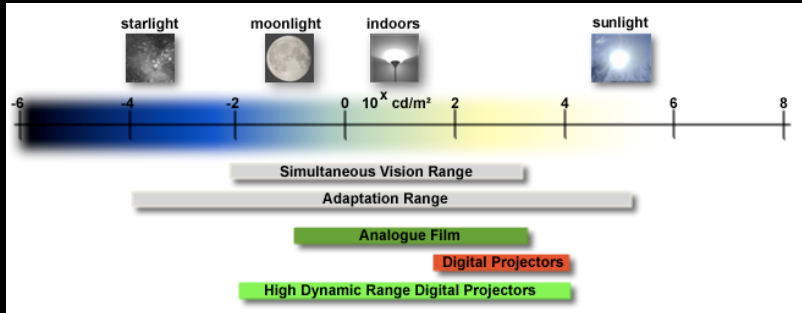


brightness and contrast are modified



High-Dynamic Range Projection

- contrast ratio w/o auto-iris
 - LCD → 600 : 1
 - DMD → 1,000 ~ 2,000 : 1
 - LCoS → 5,000 ~ 20,000 : 1




Damberg et al., SID, 2007




Combination of Light Modulating Device

- modulating in image path




- modulating in illumination path





Combination of Light Modulating Device

1. low resolution LCoS panel (chrominance modulator)
2. high resolution LCoS panel (luminance modulator)



Kusakabe et al., IDW, 2006



Combination of Light Modulating Device


- final contrast ratio is the product of two modulation blocks
 - $c_1 * c_2 : 1$
(chrominance modulator = $c_1 : 1$, luminance modulator = $c_2 : 1$)
- much lower resolution of chrominance modulator can be used
 - human vision features high spatial frequency response with respect to luminance more than chrominance.
- experimental result
 - contrast ratio = 1,100,000 : 1 (Kusakabe et al., IDW, 2006)



Image modulated in chrominance
(low resolution)



Image modulated in luminance
(high resolution)



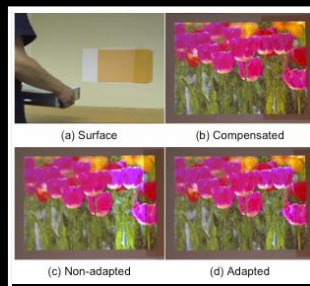
Output image

High-Dynamic Range Laser Projector


- LaserCave (Biehling et al, 2004)
- achieves the contrast ratio of 100,000 : 1
- wide color gamut can be covered




Courtesy, The Corporation for Laser Optics Research (COLOR)



Fujii et al, CVPR, 2005



Cotting et al., ISMAR 2004



Zollmann and Bimber, Eurographics, 2007

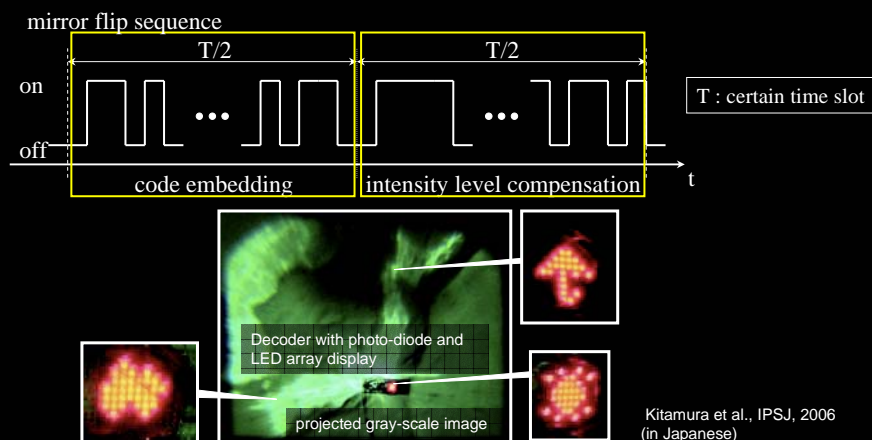
High Speed Projection

- commercially available DLP stereo projector
 - 120 Hz (DepthQ, InFocus)
- DMD Discovery
 - mirror flip sequence can be controlled in 16,300 Hz (binary)




High Speed Imperceptible Coding

- temporal coding technique
 - each mirror flip state can be used as a binary data



High Speed Projector-Camera System

- simulating spatially varying lighting on a live performance




high speed camera

high speed projector




the performance relit by a stained-glass window



Jones et al., CVMP, 2006



High Speed Projector-Camera System

- fast range scanning
 - McDowall et al., EDT'05, Takei et al., IROS'07
 - 3,000 Hz shape measurement with tracking of the object



high speed camera

high speed projector




Takei et al., IROS, 2007

these approaches do not project pictorial image content,
but rather represent encouraging examples of High Speed ProCams



LED Projector

- Raskar et al. proposed 10,000 Hz binary frame-rate graycode projector
→ LED can be better light source than conventional UHP lamps




Raskar et al., Siggraph, 2007



LED Projector

- Raskar et al. proposed 10,000 Hz binary frame-rate graycode projector
→ LED can be better light source than conventional UHP lamps



Coming up...

Augmented Studio




Deutsche Forschungsgemeinschaft (DFG)
www.uni-weimar.de/medien/AR/AugmentedStudio



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Augmented Studio



Deutsche Forschungsgemeinschaft (DFG)
www.uni-weimar.de/medien/AR/AugmentedStudio

O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



augmented reality
[erweiterte realität]

Bauhaus-Universität Weimar

Dynamic Bluescreens



Grundhoefer and Bimber, Siggraph 2008 Poster (Tech. Report available at www.uni-weimar.de/medien/AR)


O. Bimber

Visually Augmenting the real World with Projectors

17/05/2008



Superimposing Dynamic Range



Bimber and Iwai, Siggraph 2008 TechDemo (Tech. Report available at www.uni-weimar.de/medien/AR)



I would like to thank all current and former members of the ARGroup@BUW...

- Anselm Grundhöfer
- Erich Bruns
- Christian Tonn
- Daisuke Iwai
- Ferry Häntsch
- Martin Triebel
- Alexander Schiewe
- Max Große
- Bruno Fernandes
- Miki Matsumura
- Daniel Kurz
- Tobias Langlotz
- Saskia Groenewegen
- Stefanie Zollmann
- Christian Nitschke
- Daniel Kolster
- Alexander Kleppe
- Daniel Danich
- Thomas Zeidler
- Franz Corland
- Uwe Hähne
- Mathias Mörhing
- Sebastian Knedel
- Gordon Weitzstein
- Tobias Langlotz
- Tim Golub
- Andreas Einmieding
- Arnd Oberländer
- Michael Rang
- Sebastian Thiele
- Benjamin Brombach
- Mari Man Fu
- Thomas Klemmer
- Sebastian Derkau
- Christian Lessig
- Daniel Danich
- Petro Kappalos
- Paul Fackler
- Manja Seeger
- Steffen Hippeli
- Franz Corland

Thank you! Questions?

...and our sponsors

- Deutsche Forschungsgemeinschaft (DFG)
- Stiftung für Technologie und Innovation Thüringen (STIFT)
- Orange (France Telecom) Foundation
- Volkswagen AG
- VIOSO GmbH
- Nokia
- Bennert GmbH

The following article is reproduced by kind permission of the Eurographics Association

Reference:

Bimber, O., Iwai, D., Wetzstein, G. and Grundhoefer, A.
The Visual Computing of Projector-Camera Systems
Eurographics State of the Art Report, 2007
To appear in Computer Graphics Forum, 2008

(c) Eurographics Association 2007

The Visual Computing of Projector-Camera Systems

Oliver Bimber¹, Daisuke Iwai^{1,2}, Gordon Wetzstein³ and Anselm Grundhöfer¹

¹Bauhaus-University Weimar, Germany, {bimber, iwai, grundhoefer}@uni-weimar.de

²Osaka University, Japan, iwai@sens.sys.es.osaka-u.ac.jp

³University of British Columbia, Canada, wetzste1@cs.ubc.ca

Abstract

This article focuses on real-time image correction techniques that enable projector-camera systems to display images onto screens that are not optimized for projections, such as geometrically complex, colored and textured surfaces. It reviews hardware accelerated methods like pixel-precise geometric warping, radiometric compensation, multi-focal projection, and the correction of general light modulation effects. Online and offline calibration as well as invisible coding methods are explained. Novel attempts in super-resolution, high dynamic range and high-speed projection are discussed. These techniques open a variety of new applications for projection displays. Some of them will also be presented in this report.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation I.4.8 [Image Processing and Computer Vision]: Scene Analysis I.4.9 [Image Processing and Computer Vision]: Applications

Keywords: Projector-Camera Systems, Image-Correction, GPU Rendering, Virtual and Augmented Reality

1. Introduction

Their increasing capabilities and declining cost make video projectors widespread and established presentation tools. Being able to generate images that are larger than the actual display device virtually anywhere is an interesting feature for many applications that cannot be provided by desktop screens. Several research groups discover this potential by applying projectors in unconventional ways to develop new and innovative information displays that go beyond simple screen presentations.

Today's projectors are able to modulate the displayed images spatially and temporally. Synchronized camera feedback is analyzed to support a real-time image correction that enables projections on complex everyday surfaces that are not bound to projector-optimized canvases or dedicated screen configurations.

This article reviews current projector-camera-based image correction techniques. It starts in section 2 with a discussion on the problems and challenges that arise when projecting images onto non-optimized screen surfaces. Geometric warping techniques for surfaces with different topology and

reflectance are described in section 3. Section 4 outlines radiometric compensation techniques that allow the projection onto colored and textured surfaces of static and dynamic scenes and configurations. It also explains state-of-the-art techniques that consider parameters of human visual perception to overcome technical limitations of projector-camera systems. In both sections (3 and 4), conventional structured light range scanning as well as imperceptible coding schemes are outlined that support projector-camera calibration (geometry and radiometry). While the previously mentioned sections focus on rather simple light modulation effects, such as diffuse reflectance, the compensation of complex light modulations, such as specular reflection, interreflection, refraction, etc. are explained in section 5. It also shows how the inverse light transport can be used for compensating all measurable light modulation effects. section 6 is dedicated to a discussion on how novel (at present mainly experimental) approaches in high speed, high dynamic range, large depth of field and super-resolution projection can overcome the technical limitations of today's projector-camera systems in the future.

Such image correction techniques have proved to be use-

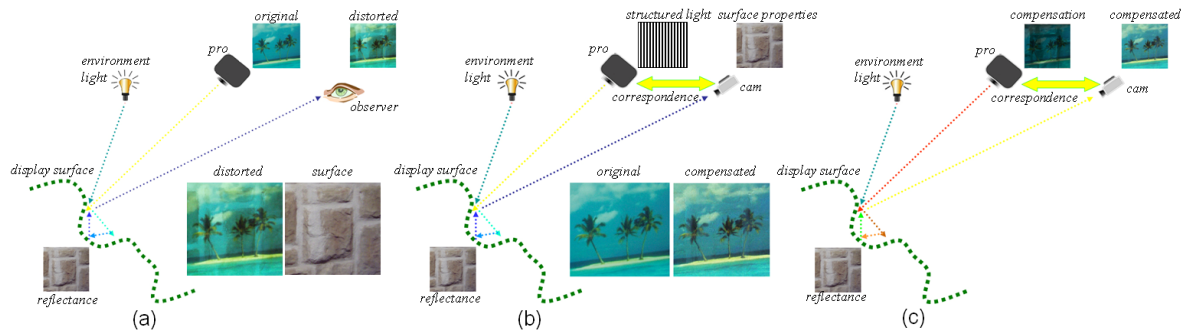


Figure 1: Projecting onto non-optimized surfaces can lead to visual artifacts in the reflected image (a). Projector-camera systems can automatically scan surface and environment properties (b) to compute compensation images during run-time that neutralize the measured light modulations on the surface (c).

ful tools for scientific experiments, but also for real-world applications. Some examples are illustrated in figures 25-29. (on the last page of this report). They include on-site architectural visualization, augmentations of museum artifacts, video installations in cultural heritage sites, outdoor advertisement displays, projections onto stage settings during live performances, and ad-hoc stereoscopic VR/AR visualizations within everyday environments. Besides these rather individual application areas, real-time image correction techniques hold the potential of addressing future mass markets, such as flexible business presentations with quickly approaching pocket projector technology, upcoming projection technology integrated in mobile devices - like cellphones, or game-console driven projections in the home-entertainment sector.

2. Challenges of Non-Optimized Surfaces

For conventional applications, screen surfaces are optimized for a projection. Their reflectance is usually uniform and mainly diffuse (although with possible gain and anisotropic properties) across the surface, and their geometrical topologies range from planar and multi-planar to simple parametric (e.g., cylindrical or spherical) surfaces. In many situations, however, such screens cannot be applied. Some examples are mentioned in section 1. The modulation of the projected light on these surfaces, however, can easily exceed a simple diffuse reflection modulation. In addition, blending with different surface pigments and complex geometric distortions can degrade the image quality significantly. This is outlined in figure 1.

The light of the projected images is modulated on the surface together with possible environment light. This leads to a color, intensity and geometry distorted appearance (cf. figure 1a). The intricacy of the modulation depends on the complexity of the surface. It can contain interreflections, diffuse and specular reflections, regional defocus effects, refractions, and more. To neutralize these modulations in real-time, and

consequently to reduce the perceived image distortions is the aim of many projector-camera approaches.

In general, two challenges have to be mastered to reach this goal: First, the modulation effects on the surface have to be measured and evaluated with computer vision techniques and second, they have to be compensated in real-time with computer graphics approaches. Structured light projection and synchronized camera feedback enables the required parameters to be determined and allows a geometric relation between camera(s), projector(s) and surface to be established (cf. figure 1b). After such a system is calibrated, the scanned surface and environment parameters can be used to compute compensation images for each frame that needs to be projected during run-time. If the compensation images are projected, they are modulated by the surface together with the environment light in such a way that the final reflected images approximate the original images from the perspective of the calibration camera/observer (cf. figure 1c).

The sections below will review techniques that compensate individual modulation effects.

3. Geometric Registration

The amount of geometric distortion of projected images depends on how much the projection surface deviates from a plane, and on the projection angle. Different *geometric projector-camera registration techniques* are applied for individual surface topologies. While simple homographies are suited for registering projectors with planar surfaces, projective texture mapping can be used for non-planar surfaces of known geometry. This is explained in subsection 3.1. For geometrically complex and textured surfaces of unknown geometry, image warping based on look-up operations has frequently been used to achieve a pixel-precise mapping, as discussed in subsection 3.2. Most of these techniques require structured light projection to enable a fully automatic calibration. Some modern approaches integrate the structured code information directly into the projected image content in such

a way that an imperceptible calibration can be performed during run-time. They are presented in subsection 3.3. Note, that image warping techniques for parametric surfaces, such as explained in [RvBWR04] are out of the scope of this article.

3.1. Uniformly Colored Surfaces

For surfaces whose reflectance is optimized for projection (e.g., surfaces with a homogenous white reflectance), a geometric correction of the projected images is sufficient to provide an undistorted presentation to an observer with known perspective. Slight misregistrations of the images on the surface in the order of several pixels lead to geometric artifacts that -in most cases- can be tolerated. This section gives a brief overview over general geometry correction techniques that support a single or multiple projectors for such surfaces.




Figure 2: Camera-based projector registration for untextured planar (a) and non-planar (b) surfaces.

If multiple projectors (*pro*) have to be registered with a planar surface via camera (*cam*) feedback (cf. figure 2a), collineations with the plane surface can be expressed as 3×3 camera-to-projector homography matrix H :

$$H_{3 \times 3} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix}$$

A homography matrix can be automatically determined numerically by correlating a projection pattern to its corresponding camera image. Knowing the homography matrix H_i for projector pro_i and the calibration camera *cam*, allows the mapping from camera pixel coordinates $cam(x, y)$ to the corresponding projector pixel coordinates $pro_i(x, y)$ with $pro_i(x, y, 1) = H_i \cdot cam(x, y)$. The homographies are usually extended to homogenous 4×4 matrices to make them compatible with conventional transformation pipelines and to consequently benefit from single pass rendering [Ras99]:

$$A_{4 \times 4} = \begin{bmatrix} h_{11} & h_{12} & 0 & h_{13} \\ h_{21} & h_{22} & 0 & h_{23} \\ 0 & 0 & 1 & 0 \\ h_{31} & h_{32} & 0 & h_{33} \end{bmatrix}$$

Multiplied after the projection transformation, they map

normalized camera coordinates into normalized projector coordinates. An observer located at the position of the (possibly off-axis aligned) calibration camera perceives a correct image in this case. Such a camera-based approach is frequently used for calibrating tiled screen projection displays. A sparse set of point correspondences is determined automatically using structured light projection and camera feedback [SPB04]. The correspondences are then used to solve for the matrix parameters of H_i for each projector i . In addition to a geometric projector registration, a camera-based calibration can be used for photometric (luminance and chrominance) matching among multiple projectors. A detailed discussion on the calibration of tiled projection screens is out of the scope of this report. It does not cover multi-projector techniques that are suitable for conventional screen surfaces. The interested reader is referred to [BMY05] for a state-of-the-art overview over such techniques. Some other approaches apply mobile projector-camera systems and homographies for displaying geometrically corrected images on planar surfaces (e.g., [RBvB*04]).

Once the geometry of the projection surface is non-planar but known (cf. figure 2b), a two-pass rendering technique can be applied for projecting the images in an undistorted way [RWC*98, RBY*99]: In the first pass, the image that has to be displayed is off-screen rendered from a target perspective (e.g. the perspective of the camera or an observer). In the second step, the geometry model of the display surface is texture-mapped with the previously rendered image while being rendered from the perspective of each projector *pro*. For computing the correct texture coordinates that ensure an undistorted view from the target perspective *projective texture mapping* is applied. This hardware accelerated technique dynamically computes a texture matrix that maps the 3D vertices of the surface model from the perspectives of the projectors into the texture space of the target perspective.

A camera-based registration is possible in this case as well. For example, instead of a visible (or an invisible - as discussed in section 3.3) structured light projection, features of the captured distorted image that is projected onto the surface can be analyzed directly. A first example was presented in [YW01] that evaluates the deformation of the image content when projected onto the surface to reconstruct the surface geometry, and refine it iteratively. This approach assumes a calibrated camera-projector system and an initial rough estimate of the projection surface. If the surface geometry has been approximated, the two-pass method outlined above can be applied for warping the image geometry in such a way that it appears undistorted. In [JF07] a similar method is described that supports a movable projector and requires a stationary and calibrated camera, as well as the known surface geometry. The projector's intrinsic parameters and all camera parameters have to be known in both cases. While the method in [YW01] results in the estimated surface geometry, the approach of [JF07] leads to the projector's extrinsic parameters. The possibility of establishing the correspondence

between projector and camera pixels in these cases, however, depends always on the quality of detected image features and consequently on the image content itself. To improve their robustness, such techniques apply a predictive feature matching rather than a direct matching for features in projector and camera space.

However, projective texture mapping in general assumes a simple pinhole camera/projector model and normally does not take the lens distortion of projectors into account (yet, a technique that considers the distortion of the projector for planar untextured screens has been described in [BJM07]). This -together with flaws in feature matching or numerical minimization errors- can cause misregistrations of the projected images in the range of several pixels – even if other intrinsic and extrinsic parameters have been determined precisely. These slight geometric errors are normally tolerable on uniformly colored surfaces. Projecting corrected images onto textured surfaces with misregistrations in this order causes -even with applying a radiometric compensation (see section 4)- immediate visual intensity and color artifacts that are well visible. Consequently, more precise registration techniques are required for textured surfaces.

3.2. Textured Surfaces

Mapping projected pixels precisely onto different colored pigments of textured surfaces is essential for an effective radiometric compensation (described in section 4). To achieve a precision on a pixel basis is not practical with the registration techniques outlined in section 3.1. Instead of registering projectors by structured light sampling followed by numerical optimizations that allow the computation of projector-camera correspondences via homographies or other projective transforms, they can be measured pixel-by-pixel and queried through look-up operations during runtime. Well known structured light techniques [BMS98, SPB04] (e.g., gray code scanning) can be used as well for scanning the 1-to-n mapping of camera pixels to projector pixels. This mapping is stored in a 2D look-up-texture having a resolution of the camera, which in the following is referred to as $C2P$ map (cf. figure 3). A corresponding texture that maps every projector pixel to one or many camera pixels can be computed by reversing the $C2P$ map. This texture is called $P2C$ map. It has the resolution of the projector.

The 1-to-n relations (note that n can also become 0 during the reversion process) are finally removed from both maps through averaging and interpolation (e.g., via a Delaunay triangulation of the transformed samples in the $P2C$ map, and a linear interpolation of the pixel colors that store the displacement values within the computed triangles). Figure 3b illustrates the perspective of a camera onto a scene and the scanned and color-coded (red=x, green=y) $C2P$ texture that maps camera pixels to their corresponding projector pixel coordinates. Note, that all textures contain floating point numbers.




Figure 3: Camera-based projector registration for textured surfaces (a). The camera perspective onto a scene (b-top) and the scanned look-up table that maps camera pixels to projector pixels. Holes are not yet removed in this example (b-bottom).

These look-up textures contain only the 2D displacement values of corresponding projector and camera pixels that map onto the same surface point. Thus, neither the 3D surface geometry, nor the intrinsic or extrinsic parameters of projectors and camera are known.

During runtime, a fragment shader maps all pixels from the projector perspective into the camera perspective (via texture look-ups in the $P2C$ map) to ensure a geometric consistency for the camera view. We want to refer to this as *pixel displacement mapping*. If multiple projectors are involved, a $P2C$ map has to be determined for each projector. Projector-individual fragment shaders will then perform a customized pixel-displacement mapping during multiple rendering steps, as described in [BEK05].

In [BWEN05] and in [ZLB06], pixel-displacement mapping has been extended to support moving target perspectives (e.g., of the camera and/or the observer). In [BWEN05] an image-based warping between multiple $P2C$ maps that have been pre-scanned for known camera perspectives is applied. The result is an estimated $P2C$ map for a new target perspective during runtime. Examples are illustrated in figures 27 and 28. While in this case, the target perspective must be measured (e.g., using a tracking device), [ZLB06] analyzes image features of the projected content to approximate a new $P2C$ as soon as the position of the calibration camera has changed. If this is not possible because the detected features are too unreliable, a structured light projection is triggered to scan a correct $P2C$ map for the new perspective.

3.3. Embedded Structured Light

Section 3.1 has already discussed registration techniques (i.e., [YW01, JF07]) that do not require the projection of structured calibration patterns, like gray codes. Instead, they analyze the distorted image content, and thus depend on matchable image features in the projected content. Structured light techniques, however, are more robust because they generate such features

synthetically. Consequently, they do not depend on the image content. Overviews over different general coding schemes are given in [BMS98, SPB04].

Besides a spatial modulation, a temporal modulation of projected images allows integrating coded patterns that are not perceivable due to limitations of the human visual system. Synchronized cameras, however, are able to detect and extract these codes. This principle has been described by Raskar et al. [RWC*98], and has been enhanced by Cotting et al. [CNGF04]. It is referred to as *embedded imperceptible pattern projection*. Extracted code patterns, for instance, allow the simultaneous acquisition of the scenes' depth and texture for 3D video applications [WWC*05], [VVSC05]. These techniques, however, can be applied to integrate the calibration code directly into the projected content to enable an invisible online calibration. Thus, the result could be, for instance, a *P2C* map scanned by a binary gray code or an intensity phase pattern that is integrated directly into the projected content.

The first applicable imperceptible pattern projection technique was presented in [CNGF04], where a specific time slot (called *BIEP*=binary image exposure period) of a DLP projection sequence is occupied exclusively for displaying a binary pattern within a single color channel (multiple color channels are used in [CZGF05] to differentiate between multiple projection units). Figure 4 illustrates an example.




Figure 4: Mirror flip (on/off) sequences for all intensity values of the red color channel and the chosen binary image exposure period. ©2004 IEEE [CNGF04]

The *BIEP* is used for displaying a binary pattern. A camera that is synchronized to exactly this projection sequence will capture the code. As it can be seen in the selected *BIEP* in figure 4, the mirror flip sequences are not evenly distributed over all possible intensities. Thus, the intensity of each projected original pixel might have to be modified to ensure that the mirror state is active which encodes the desired binary value

at this pixel. This, however, can result in a non-uniform intensity fragmentation and a substantial reduction of the tonal values. Artifacts are diffused using a dithering technique. A coding technique that benefits from re-configurable mirror flip sequences using the DMD discovery board is described in section 6.4.

Another possibility of integrating imperceptible code patterns is to modulate the intensity of the projected image I with a spatial code. The result is the code image I_{cod} . In addition, a compensation image I_{com} is computed in such a way that $(I_{cod} + I_{com})/2 = I$. If both images are projected alternately with a high speed, human observers will perceive I due to the slower temporal integration of the human visual system. This is referred to as *temporal coding* and was shown in [RWC*98]. The problem with this simple technique is that the code remains visible during eye movements or code transitions. Both cannot be avoided for the calibration of projector-camera systems using structured light techniques. In [GSHB07] properties of human perception, like adaptation limitations to local contrast changes, are taken into account for adapting the coding parameters depending on local characteristics, such as spatial frequencies and local luminance values of image and code. This makes a truly imperceptible temporal coding of binary information possible. For binary codes, I is regionally decreased ($I_{cod} = I - \Delta$ to encode a binary 0) or increased ($I_{cod} = I + \Delta$ to encode a binary 1) in intensity by the amount of Δ , while the compensation image is computed with $I_{com} = 2I - I_{cod}$. The code can then be reconstructed from the two corresponding images (C_{cod} and C_{com}) captured by the camera with $C_{cod} - C_{com} \leq 0$. Thereby, Δ is one coding parameter that is locally adapted.

In [PLJP07] another technique for adaptively embedding complementary patterns into projected images is presented. In this work the embedded code intensity is regionally adapted depending on the spatial variation of neighbouring pixels and their color distribution in the YIQ color space. The final code contrast of Δ is then calculated depending on the estimated local spatial variations and color distributions.

In [ZB07], the binary temporal coding technique was extended to encoding intensity values as well. For this, the code image is computed with $I_{cod} = I\Delta$ and the compensation image with $I_{com} = I(2 - \Delta)$. The code can be extracted from the camera images with $\Delta = 2C_{cod}/(C_{cod} + C_{com})$. Using binary and intensity coding, an imperceptible multi-step calibration technique is presented in [ZB07] which is visualized in figure 5, and is outlined below.

A re-calibration is triggered automatically if misregistrations between projector and camera are detected (i.e., due to motion of camera, projector or surface). This is achieved by continuously comparing the correspondences of embedded point samples. If necessary, a first rough registration is carried out by sampling binary point patterns (cf. figure 5b) that leads to a mainly interpolated *P2C* map (cf. figure 5f). This step is followed by an embedded measurement of the surface




Figure 5: Imperceptible multi-step calibration for radiometric compensation. A series of invisible patterns (b-e) integrated into an image (a) and projected onto a complex surface (g) results in surface measurements (f-i) used for radiometric compensation (j). ©2007 Eurographics [ZB07].

reflectance (cf. figures 5c,g), which is explained in section 4.2. Both steps lead to quick but imprecise results. Then a more advanced 3-step phase shifting technique (cf. figure 5e) is triggered that results in a pixel-precise *P2C* registration (cf. figure 5i). For this, intensity coding is required (cf. figure 5h). An optional gray code might be necessary for surfaces with discontinuities (cf. figure 5d). All steps are invisible to the human observer and are executed while dynamic content can be projected with a speed of 20Hz.

In general, temporal coding is not limited to the projection of two images only. Multiple code and compensation images can be projected if the display frame-rate is high enough. This requires fast projectors and cameras, and will be discussed in section 6.4.

An alternative to embedding imperceptible codes in the visible light range would be to apply infrared light as shown in [SMO03] for augmenting real environments with invisible information. Although it has not been used for projector-camera calibration, this would certainly be possible.

4. Radiometric Compensation

For projection screens with spatially varying reflectance, color and intensity compensation techniques are required in addition to a pixel-precise geometric correction. This is known as *radiometric compensation*, and is generally used to minimize the artifacts caused by the local light modulation between projection and surface. Besides the geometric mapping between projector and camera, the surface's reflectance parameters need to be measured on a per-pixel basis before using them for real-time image corrections during run-time. In most cases, a one-time calibration process applies visible structured light projections and camera feedback to establish the correspondence between camera and projector pixels (see section 3.2) and to measure the surface pigment's radiometric behavior.

A pixel precise mapping is essential for radiometric compensation since slight misregistrations (in the order of only a few pixels) can lead to significant blending artifacts - even if the geometric artifacts are marginal. Humans are extremely sensitive to even small (less than 2%) intensity variations.

This section reviews different types of radiometric compensation techniques. Starting with methods that are suited for static scenes and projector-camera configurations in subsection 4.1, it will then discuss more flexible techniques that support dynamic situations (i.e., moving projector-camera systems and surfaces) in subsection 4.2. Finally, most recent approaches are outlined that dynamically adapt the image content before applying a compensation based on pure radiometric measurements to overcome technical and physical limitations of projector-camera systems. Such techniques take properties of human visual perception into account and are explained in subsection 4.3.

4.1. Static Techniques

In its most basic configuration (cf. figure 6a), an image is displayed by a single projector (*pro*) in such a way that it appears correct (color and geometry) for a single camera view (*cam*). Thereby, the display surfaces must be diffuse, but can have an arbitrary color, texture and shape. The first step is to determine the geometric relations of camera pixels and projector pixels over the display surface. As explained in section 3, the resulting *C2P* and *P2C* look-up textures support a pixel-precise mapping from camera space to projector space and vice versa.

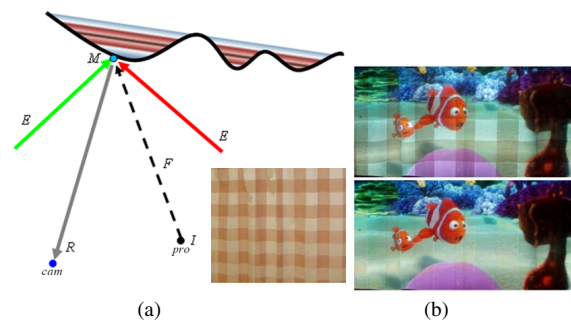


Figure 6: Radiometric compensation with a single projector (a) and sample images projected without and with compensation onto window curtains (b). ©2007 IEEE [BEK05]

Once the geometric relations are known, the radiometric parameters are measured. One of the simplest radiometric compensation approaches is described in [BEK05]: With respect to figure 6a, it can be assumed that a light ray with intensity I is projected onto a surface pigment with reflectance M . The fraction of light that arrives at the pigment depends on the geometric relation between the light source (i.e., the projector) and the surface. A simple representation of the form factor can be used for approximating this fraction: $F = f * \cos(\alpha)/r^2$, where α is the angular correlation between the light ray and the surface normal and r is the distance (considering square distance attenuation) between the light source and the surface. The factor f allows scaling the intensity to avoid clipping (i.e., intensity values that exceed the luminance capabilities of the projector) and to consider the simultaneous contributions of multiple projectors. Together with the environment light E , the projected fraction of I is blended with the pigment's reflectance M : $R = EM + IFM$. Thereby, R is the diffuse radiance that can be captured by the camera. If R , F , M , and E are known, a compensation image I can be computed with:

$$I = (R - EM)/FM \quad (1)$$

In a single-projector configuration, E , F , and M cannot be determined independently. Instead, FM is measured by projecting a white flood image ($I = 1$) and turning off the entire environment light ($E = 0$), and EM is measured by projecting a black flood image ($I = 0$) under environment light. Note, that EM also contains the black level of the projector. Since this holds for every discrete camera pixel, R , E , FM and EM are entire textures and equation 1 can be computed together with pixel displacement mapping (see section 3.2) in real-time by a fragment shader. Thus, every rasterized projector pixel that passes through the fragment shader is displaced and color compensated through texture look-ups. The projection of the resulting image I onto the surface leads to a geometry and color corrected image that approximates the desired original image $R = O$ for the target perspective of the camera.

One disadvantage of this simple technique is that the optical limitations of color filters used in cameras and projectors are not considered. These filters can transmit a quite large spectral band of white light rather than only a small monochromatic one. In fact, projecting a pure red color, for instance, usually leads to non-zero responses in the blue and green color channels of the captured images. This is known as the *color mixing* between projector and camera, which is not taken into account by equation 1.

Color mixing can be considered for radiometric compensation: Nayar et al. [NPG03], for instance, express the color transform between each camera and projector pixel as pixel-individual 3x3 *color mixing matrices*:

$$V = \begin{bmatrix} v_{RR} & v_{RG} & v_{RB} \\ v_{GR} & v_{GG} & v_{GB} \\ v_{BR} & v_{BG} & v_{BB} \end{bmatrix}$$

Thereby, v_{RG} represents the green color component in the red color channel, for example. This matrix can be estimated from measured camera responses of multiple projected sample images. It can be continuously refined over a closed feedback loop (e.g., [FGN05]) and is used to correct each pixel during runtime. In the case the camera response is known while the projector response can remain unknown, it can be assumed that $v_{ii} = 1$. This corresponds to an unknown scaling factor, and V is said to be normalized. The off-diagonal values can then be computed with $v_{ij} = \Delta C_j / \Delta P_i$, where ΔP_i is the difference between two projected intensities ($P_{1i} - P_{2i}$) of primary color i , and ΔC_j is the difference of the corresponding captured images ($C_{1j} - C_{2j}$) in color channel j . Thus, 6 images have to be captured (2 per projected color channel) to determine all v_{ij} . The captured image R under projection of I can now be expressed with: $R = VI$. Consequently, the compensation image can be computed with the inverse color mixing matrix:

$$I = V^{-1}R \quad (2)$$

Note, that V is different for each camera pixel and contains the surface reflectance, but not the environment light. Another way of determining V is to numerically solve equation 2 for V^{-1} if enough correspondences between I and R are known. In this case, V is un-normalized and v_{ii} is proportional to $[FM_R, FM_G, FM_B]$. Consequently, the off-diagonal values of V are 0 if no color mixing is considered. Yoshida et al. [YHS03] use an un-normalized 3x4 color mixing matrix. In this case, the fourth column represents the constant environment light contribution. A refined version of Nayar's technique was used for controlling the appearance of two- and three-dimensional objects, such as posters, boxes and spheres [GPNB04]. Sections 4.2 and 4.3 also discuss variations of this method for dynamic situations and image adaptations. Note, that a color mixing matrix was also introduced in the context of shape measurement based on a color coded pattern projection [CKS98].

All of these techniques support image compensation in real-time, but suffer from the same problem: if the compensation image I contains values above the maximal brightness or below the black level of the projector, clipping artifacts will occur. These artifacts allow the underlying surface structure to become visible. The intensity range for which radiometric compensation without clipping is possible depends on the surface reflectance, on the brightness and black level of the projector, on the required reflected intensity (i.e., the desired original image), and on the environment light contribution.

Figure 7 illustrates an example that visualizes the reflection

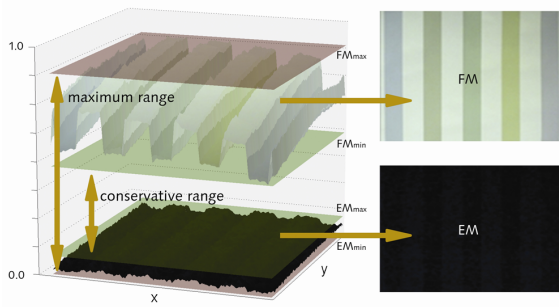


Figure 7: Intensity range reflected by a striped wall paper. ©2007 IEEE [GB07]

properties for a sample surface. By analyzing the responses in both datasets (FM and EM), the range of intensities for a conservative compensation can be computed. Thus, only input pixels of the desired original image $R = O$ within this global range (bound by the two green planes - from the maximum value EM_{max} to the minimum value FM_{min}) can be compensated correctly for each point on the surface without causing clipping artifacts. All other intensities can potentially lead to clipping and incorrect results. This conservative intensity range for radiometric compensation is smaller than the maximum intensity range achieved when projecting onto optimized (i.e., diffuse and white) surfaces.

Different possibilities exist to reduce these clipping problems. While applying an amplifying transparent film material is one option that is mainly limited to geometrically simple surfaces, such as paintings [BCK*05], the utilization of multiple projectors is another option.




Figure 8: Radiometric compensation with multiple projectors. Multiple individual low-capacity projection units (a) are assumed to equal one single high-capacity unit (b).

The simultaneous contribution of multiple projectors increases the total light intensity that reaches the surface. This can overcome the limitations of equation 1 for extreme situations (e.g., small FM values or large EM values) and can consequently avoid an early clipping of I . Therefore, [BEK05]

presents a multi-projector approach for radiometric compensation: If N projectors are applied (cf. figure 8a), the measured radiance captured by the camera can be approximated with: $R = EM + \sum_i^N (I_i FM_i)$. One strategy is to balance the projected intensities equally among all projectors i , which leads to:

$$I_i = (R - EM) / \sum_j^N (I_j FM_j) \quad (3)$$

Conceptually, this is equivalent to the assumption that a single high capacity projector (pro_v) produces the total intensity arriving on the surface virtually (cf. figure 8b). This equation can also be solved in real-time by projector-individual fragment shaders (based on individual parameter textures FM_i , $C2P_i$ and $P2C_i$ - but striving for the same final result R). Note, that EM also contains the accumulated black level of all projectors. If all projectors provide linear transfer functions (e.g., after a linearization) and identical brightness, a scaling of $f_i = 1/N$ used in the form factor balances the load among them equally. However, f_i might be decreased further to avoid clipping and to adapt for differently aged bulbs. Note however, that the total black level increases together with the total brightness of a multiple projector configuration. Thus, an increase in contrast cannot be achieved. Possibilities for dynamic range improvements are discussed in section 6.3.

Since the required operations are simple, a pixel-precise radiometric compensation (including geometric warping through pixel-displacement mapping) can be achieved in real-time with fragment shaders of modern graphics cards. The actual speed depends mainly on the number of pixels that have to be processed in the fragment shader. For example, frame-rates of $>100\text{Hz}$ can be measured for radiometric compensations using equation 1 for PAL-resolution videos projected in XGA resolution.

4.2. Dynamic Surfaces and Configurations

The techniques explained in section 4.1 are suitable for purely static scenes and fixed projector-camera configurations. They require a one-time calibration before runtime. For many applications, however, a frequent re-calibration is necessary because the alignment of camera and projectors with the surfaces changes over time (e.g., due to mechanical expansion through heating, accidental offset, intended readjustment, mobile projector-camera systems, or dynamic scenes). In these cases, it is not desired to disrupt a presentation with visible calibration patterns. While section 3 discusses several online calibration methods for geometric correction, this section reviews online radiometric compensation techniques.

Fujii et al. have described a dynamically adapted radiometric compensation technique that supports changing projection surfaces and moving projector-camera configurations [FGN05]. Their system requires a fixed co-axial alignment

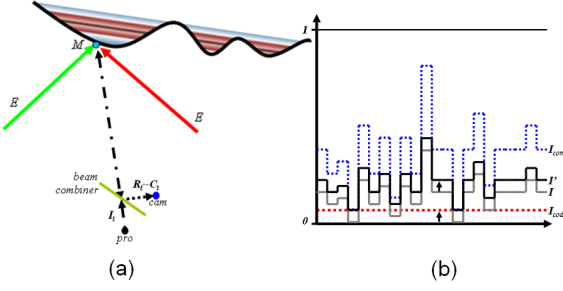


Figure 9: Co-axial projector-camera alignment (a) and reflectance measurements through temporal coding (b).

of projector and camera (cf. figure 9a). An optical registration of both devices makes a frequent geometric calibration unnecessary. Thus, the fixed mapping between projector and camera pixels does not have to be re-calibrated if either surface or configuration changes. At an initial point in time 0 the surface reflectance is determined under environment light (E_0M_0). To consider color mixing as explained in section 4.1, this can be done by projecting and capturing corresponding images I_0 and C_0 . The reflected environment light E_0 at a pigment with reflectance M_0 can then be approximated by $E_0M_0 = C_0 - V_0I_0$, where V_0 is the un-normalized color mixing matrix at time 0, which is constant. After initialization, the radiance R_t at time t captured by the camera under projection of I_t can be approximated with: $R_t = M_t/M_0(E_tM_0 + V_0I_t)$. Solving for I_t results in:

$$I_t = V_0^{-1}(R_tM_0/M_{t-1} - E_{t-1}M_0) \quad (4)$$

Thereby, $R_t = O_t$ is the desired original image and I_t the corresponding compensation image at time t . The environment light contribution cannot be measured during runtime. It is approximated to be constant. Thus, $E_{t-1}M_0 = E_0M_0$. The ratio M_0/M_{t-1} is then equivalent to the ratio C_0/C_{t-1} . In this closed feedback loop, the compensation image I_t at time t depends on the captured parameters (C_{t-1}) at time $t-1$. This one-frame delay can lead to visible artifacts. Furthermore, the surface reflectance M_{t-1} is continuously estimated based on the projected image I_{t-1} . Thus, the quality of the measured surface reflectance depends on the content of the desired image R_{t-1} . If R_{t-1} has extremely low or high values in one or multiple color channels, M_{t-1} might not be valid in all samples. Other limitations of such an approach might be the strict optical alignment of projector and camera that might be too inflexible for many large scale applications, and that it does not support multi-projector configurations.

Another possibility of supporting dynamic surfaces and projector-camera configurations that do not require a strict optical alignment of both devices was described in [ZB07]. As outlined in section 3.3, imperceptible codes can be em-

bedded into a projected image through a temporal coding to support an online geometric projector-camera registration. The same approach can be used for embedding a uniform gray image I_{cod} into a projected image I . Thereby, I_{cod} is used to illuminate the surface with a uniform flood-light image to measure the combination of surface reflectance and projector form factor FM , as explained in section 4.1. To ensure that I_{cod} can be embedded correctly, the smallest value in I must be greater than or equal I_{cod} . If this is not the case, I is transformed to I' to ensure this condition (cf. figure 9b). A (temporal) compensation image can then be computed with $I_{com} = 2I' - I_{cod}$. Projecting I_{cod} and I_{com} with a high speed, one perceives $(I_{cod} + I_{com})/2 = I'$. Synchronizing a camera with the projection allows I_{cod} and therefore also FM to be captured. In practice, I_{cod} is approximately 3-5% of the total intensity range - depending on the projector brightness and the camera sensitivity of the utilized devices. One other advantage of this method is, that in contrast to [FGN05] the measurements of the surface reflectance do not depend on the projected image content. Furthermore, equations 1 or 3 can be used to support radiometric compensation with single or multiple projectors. However, projected (radiometric) compensation images I have to be slightly increased in intensity which leads to a smaller (equal only if $FM = 1$ and $EM = 0$) global intensity increase of $R = O$. However, since I_{cod} is small, this is tolerable. One main limitation of this method in contrast to the techniques explained in [FGN05], is that it does not react to changes quickly. Usually a few seconds (approx. 5-8s) are required for an imperceptible geometric and radiometric re-calibration. In [FGN05] a geometric re-calibration is not necessary. As explained in [GSHB07], a temporal coding requires a sequential blending of multiple code images over time, since an abrupt transition between two code images can lead to visible flickering. This is another reason for longer calibration times.

In summary we can say that fixed co-axial projector-camera alignments as in [FGN05] support real-time corrections of dynamic surfaces for a single mobile projector-camera system. The reflectance measurements' quality depends on the content in O . A temporal coding as in [ZB07] allows unconstrained projector-camera alignments and supports flexible single- or multi-projector configurations - but no real-time calibration. The quality of reflectance measurements is independent on O in the latter case. Both approaches ensure a fully invisible calibration during runtime, and enable the presentation of dynamic content (such as movies) at interactive rates (≥ 20 Hz).

4.3. Dynamic Image Adaptation

The main technical limitations for radiometric compensation are the resolution, frame-rate, brightness and dynamic range of projectors and cameras. Some of these issues will be addressed in section 6. This section presents alternative techniques that adapt the original images O based on the hu-

man perception and the projection surface properties before carrying out a radiometric compensation to reduce the effects caused by brightness limitations, such as clipping.

All compensation methods described so far take only the reflectance properties of the projection surface into account. Particular information about the input image, however, does not influence the compensation directly. Calibration is carried out once or continuously, and a static color transformation is applied as long as neither surface nor projector-camera configuration changes - regardless of the individual desired image O . Yet, not all projected colors and intensities can be reproduced as explained in section 4.1 and shown in figure 7.

Content dependent radiometric and photometric compensation methods extend the traditional algorithms by applying additional image manipulations depending on the current image content to minimize clipping artifacts while preserving a maximum of brightness and contrast to generate an optimized compensation image.

Such a content dependent radiometric compensation method was presented by Wang et al. [WSOS05]. In this method, the overall intensity of the input image is scaled until clipping errors that result from radiometric compensation are below a perceivable threshold. The threshold is derived by using a perceptually-based physical error metric that was proposed in [RPG99], which considers the image luminance, spatial frequencies and visual masking. This early technique, however, can only be applied to static monochrome images and surfaces. The numerical minimization that is carried out in [WSOS05] requires a series of iterations that make real-time rates impossible.

Park et al. [PLKP06] describe a technique for increasing the contrast in a compensation image by applying a histogram equalization to the colored input image. While the visual quality can be enhanced in terms of contrast, this method does not preserve the contrast ratio of the original image. Consequently, the image content is modified significantly, and occurring clipping errors are not considered.

A complex framework for computing an optimized photometric compensation for colored images is presented by Ashdown et al. [AOSS06]. In this method the device-independent CIE L^*u^*v color space is used, which has the advantage that color distances are based on the human visual perception. Therefore, an applied high dynamic range (HDR) camera has to be color calibrated in advance. The input images are adapted depending on a series of global and local parameters to generate an optimized compensated projection: The captured surface reflectance as well as the content of the input image are transformed into the CIE L^*u^*v color space. The chrominance values of all input image's pixels are fitted into the gamut of the corresponding projector pixels. In the next step, a luminance fitting is applied by using a relaxation method based on differential equations. Finally, the compensated adapted input image is transformed back into the RGB color space for projection.




Figure 10: Results of a content-dependent photometric compensation. The uncompensated image leads to visible artifacts (b) when being projected onto a colored surface (a). The projection of an adapted compensation image (c) minimizes the visibility of these artifacts (d). ©2006 IEEE [AOSS06]

This method achieves optimal compensation results for surfaces with varying reflectance properties. Furthermore, a compensation can be achieved for highly saturated surfaces due to the fact that besides a luminance adjustment, a chrominance adaptation is applied as well. Its numerical complexity, however, allows the compensation of still images only. Figure 10 shows a sample result: An uncompensated projection of the input image projected onto a colored surface (a) results in color artifacts (b). Projecting the adapted compensation image (c) onto the surface leads to significant improvements (d).

Ashdown et al. proposed another fitting method in [ASOS07] that uses the chrominance threshold model of human vision together with the luminance threshold to avoid visible artifacts.

Content-dependent adaptations enhance the visual quality of a radiometric compensated projection compared to static methods that do not adapt to the input images. Animated content like movies or TV-broadcasts, however, cannot be compensated in real-time with the methods reviewed above. While movies could be pre-corrected frame-by-frame in advance, real-time content like interactive applications cannot be presented.

In [GB07], a real-time solution for adaptive radiometric compensation was introduced that is implemented entirely on the GPU. The method adapts each input image in two steps: First it is analyzed for its average luminance that leads to an approximate global scaling factor which depends on the surface reflectance. This factor is used to scale the input image's intensity between the conservative and the maximum

intensity range (cf. figure 7 in section 4.1). Afterwards, a compensation image is calculated according to equation 1. Instead of projecting this compensation image directly, it is further analyzed for potential clipping errors. Errors are extracted and blurred in addition. In a final step, the input image is scaled globally again depending on its average luminance and on the calculated maximum clipping error. In addition, it is scaled locally based on the regional error values. The *threshold map* explained in [RPG99] is used to constrain the local image manipulation based on the contrast and the luminance sensitivity of human observers. Radiometric compensation (equation 1) is applied again to the adapted image, and the result is finally projected. Global, but also local scaling parameters are adapted over time to reduce abrupt intensity changes in the projection which would lead to a perceived and irritating flickering.




Figure 11: Two frames of a movie (b,e) projected onto a natural stone wall (a) with static (c,f) and real-time adaptive radiometric compensation (d,g) for bright and dark input images. ©2007 IEEE [GB07]

This approach does not apply numerical optimizations and consequently enables a practical solution to display adapted dynamic content in real-time and in increased quality (compared to traditional radiometric compensation). Yet, small clipping errors might still occur. However, especially for content with varying contrast and brightness, this adaptive technique enhances the perceived quality significantly. An example is shown in figure 11: Two frames of a movie (b,e) are projected with a static compensation technique [BEK05] (c,f) and with the adaptive real-time solution [GB07] (d,g) onto a natural stone wall (a). While clipping occurs in case (c), case (f) appears too dark. The adaptive method reduces the clipping errors for bright images (d) while maintaining details in the darker image (g).

5. Correcting Complex Light Modulations

All image correction techniques that have been discussed so far assume a simple geometric relation between camera and projector pixels that can be automatically derived using homography matrices, structured light projections, or co-axial projector-camera alignments.

When projecting onto complex everyday surfaces, however, the emitted radiance of illuminated display elements is often subject to complex lighting phenomena. Due to diffuse or specular interreflections, refractions and other global illumination effects, multiple camera pixels at spatially distant regions on the camera image plane may be affected by a single projector pixel.

A variety of projector-camera based compensation methods for specific global illumination effects has been proposed. These techniques, as well as a generalized approach to compensating light modulations using the inverse light transport will be discussed in the following subsections. We start with discussions on how diffuse interreflections (subsection 5.1) and specular highlights (subsection 5.2) can be compensated. The inverse light transport approach is introduced as the most general image correction scheme in subsection 5.3.

5.1. Interreflections

Eliminating diffuse interreflections or scattering for projection displays has recently gained a lot of interest in the computer graphics and vision community. Cancellation of interreflections has been proven to be useful for improving the image quality of immersive virtual and augmented reality displays [BGZ*06]. Furthermore, such techniques can be employed to remove indirect illumination from photographs [SMK05]. For compensating global illumination effects, these need to be acquired, stored and processed, which will be discussed for each application.

Seitz et al. [SMK05], for instance, measured an *impulse scatter function* (ISF) matrix B with a camera and a laser pointer on a movable gantry. The camera captured diffuse objects illuminated at discrete locations. Each of the samples' centroid represents one row/column in the matrix as depicted in figure 12.




Figure 12: A symmetric ISF matrix is acquired by illuminating a diffuse surface at various points, sampling their locations in the camera image and inserting captured color values into the matrix.

The ISF matrix can be employed to remove interreflections from photographs. Therefore, an *interreflection cancellation*

operator $C^1 = B^1 B^{-1}$ is defined that, when multiplied to a captured camera image R , extracts its direct illumination. B^{-1} is the ISF matrix's inverse and B^1 contains only direct illumination. For a diffuse scene, this can easily be extracted from B by setting its off-diagonal elements to zero. A related technique that quickly separates direct and indirect illumination for diffuse and non-diffuse surfaces was introduced by Nayar et al. [NKGR06].

Experimental results in [SMK05] were obtained by sampling the scene at approx. 35 locations in the camera image under laser illumination. Since B is in this case a very small and square matrix it is trivial to be inverted for computing B^{-1} . However, inverting a general light transport matrix in a larger scale is a challenging problem and will be discussed in section 5.3.

Compensating indirect diffuse scattering for immersive projection screens was proposed in [BGZ*06]. Assuming a known screen geometry, the scattering was simulated and corrected with a customized reverse radiosity scheme. Bimber et al. [Bim06] and Mukaigawa et al. [MKO06] showed that a compensation of diffuse light interaction can be performed in real-time by reformulating the radiosity equation as $I = (1 - \rho F)O$. Here O is the desired original image, I the projected compensation image, 1 the identity matrix and ρF the precomputed form-factor matrix. This is equivalent to applying the interreflection cancellation operator, introduced in [SMK05], to an image O that does not contain interreflections. The quality of projected images for a two-sided projection screen can be greatly enhanced as depicted in figure 13. All computations are performed with a relatively coarse patch resolution of about 128×128 as seen in figure 13 (c).




Figure 13: Compensating diffuse scattering: An uncompensated (a) and a compensated (b) stereoscopic projection onto a two-sided screen. Scattering and color bleeding can be eliminated (d) if the form factors (c) of the projection surface are known. ©2006 IEEE [BGZ*06]

While the form factor matrix in [Bim06, MKO06] was

precomputed, Habe et al. [HSM07] presented an algorithm that automatically acquires all photometric relations within the scene using a projector-camera system. They state also that this theoretically allows specular interreflections to be compensated for a fixed viewpoint. However, such a compensation has not been validated in the presented experiments. For the correction, a form-factor matrix inverse is required, which again is trivial to be calculated for a low patch resolution.

5.2. Specular Reflections

When projecting onto non-diffuse screens, not only diffuse and specular interreflections affect the quality of projected imagery, but a viewer may also be distracted by specular highlights. Park et al. [PLKP05] presented a compensation approach that attempts to minimize specular reflections using multiple overlapping projectors. The highlights are not due to global illumination effects, but to the incident illumination that is reflected directly toward the viewer on a shiny surface. Usually, only one of the projectors creates a specular highlight at a point on the surface. Thus, its contribution can be blocked while display elements from other projectors that illuminate the same surface area from a different angle are boosted.

For a view-dependent compensation of specular reflections, the screen's geometry needs to be known and registered with all projectors. Displayed images are pre-distorted to create a geometrically seamless projection as described in section 3. The amount of specularly for a projector i at a surface point s with a given normal n is proportional to the angle θ_i between n and the sum of the vector from s to the projector's position p_i and the vector from s to the viewer u :

$$\theta_i = \cos^{-1} \left(\frac{-n \cdot (p_i + u)}{|p_i + u|} \right) \quad (5)$$

Assuming that k projectors illuminate the same surface, a weight w_i is multiplied to each of the incident light rays for a photometric compensation:

$$w_i = \frac{\sin(\theta_i)}{\sum_{j=1}^k \sin(\theta_j)} \quad (6)$$

Park et al. [PLS*06] extended this model by an additional radiometric compensation to account for the color modulation of the underlying projection surface (cf. figure 14). Therefore, Nayar's model [NPGB03] was implemented. The required one-to-one correspondences between projector and camera pixels were acquired with projected binary gray codes [SPB04].




Figure 14: Radiometric compensation in combination with specular reflection elimination. Projection onto a specular surfaces (a) – before (b) and after (c) specular highlight compensation. ©2006 IEEE [PLS*06]

5.3. Radiometric Compensation through Inverse Light Transport

Although the previously discussed methods are successful in compensating particular aspects of the light transport between projectors and cameras, they lead to a fragmented understanding of the subject. A unified approach that accounts for many of the problems that were individually addressed in previous works was described in [WB07]. The full light transport between a projector and a camera was employed to compensate direct and indirect illumination effects, such as interreflections, refractions and defocus, with a single technique in real-time. Furthermore, this also implies a pixel-precise geometric correction. In the following subsection we refer to the approach as performing radiometric compensation. However, geometric warping is always implicitly included.

In order to compensate direct and global illumination as well as geometrical distortions in a generalized manner, the full light transport has to be taken into account. Within a projector-camera system, this is a matrix T_λ that can be acquired in a pre-processing step, for instance as described by Sen et al. [SCG*05]. Therefore, a set of illumination patterns is projected onto the scene and recorded using HDR imaging techniques (e.g. [DM97]). Individual matrix entries can then be reconstructed from the captured camera images. As depicted in figure 15, a camera image with a single lit projector pixel represents one column in the light transport matrix. Usually, the matrix is acquired in a hierarchical manner by simultaneously projecting multiple pixels.

For a single-projector-camera configuration the forward light transport is described by a simple linear equation as

$$\begin{bmatrix} r_R - e_R \\ r_G - e_G \\ r_B - e_B \end{bmatrix} = \begin{bmatrix} T_R^R & T_R^G & T_R^B \\ T_G^R & T_G^G & T_G^B \\ T_B^R & T_B^G & T_B^B \end{bmatrix} \begin{bmatrix} i_R \\ i_G \\ i_B \end{bmatrix}, \quad (7)$$

where each r_λ is a single color channel λ of a camera




Figure 15: The light transport matrix between a projector and a camera.

image with resolution $m \times n$, i_λ is the projection pattern with a resolution of $p \times q$, and e_λ are direct and global illumination effects caused by the environment light and the projector's black level captured from the camera. 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 . The model can easily be extended for k projectors and l cameras:

$$\begin{bmatrix} r_R - e_R \\ r_G - e_G \\ \vdots \\ r_B - e_B \end{bmatrix} = \begin{bmatrix} T_R^R & T_R^G & \dots & T_R^B \\ T_G^R & T_G^G & \dots & T_G^B \\ \vdots & \vdots & \ddots & \vdots \\ T_l^R & T_l^G & \dots & T_l^B \end{bmatrix} \begin{bmatrix} i_R \\ i_G \\ \vdots \\ i_B \end{bmatrix} \quad (8)$$

For a generalized radiometric compensation the camera image r_λ is replaced by a desired image o_λ of camera resolution and the system can be solved for the projection pattern i_λ that needs to be projected. This accounts for color modulations and geometric distortions of projected imagery. Due to the matrix's enormous size, sparse matrix representations and operations can help to save storage and increase performance.

A customized clustering scheme that allows the light transport matrix's pseudo-inverse to be approximated is described in [WB07]. Inverse impulse scatter functions or form-factor matrices had already been used in previous algorithms [SMK05, Bim06, MKO06, HSM07], but in a much smaller scale, which makes an inversion trivial. Using the light transport matrix's approximated pseudo-inverse, radiometric compensation reduces to a matrix-vector multiplication:

$$i_\lambda = T_\lambda^+ (o_\lambda - e_\lambda), \quad (9)$$

In [WB07], this was implemented on the GPU and yielded real-time frame-rates.

Figure 16 shows a compensated projection onto highly




Figure 16: Real-time radiometric compensation (f) of global illumination effects (a) with the light transport matrix's (b) approximated pseudo-inverse (c).

refractive material (f), which is impossible with conventional approaches (e), because a direct correspondence between projector and camera pixels is not given. The light transport matrix (cf. figure 16b) and its approximated pseudo-inverse (visualized in c) contain local and global illumination effects within the scene (global illumination effects in the matrix are partially magnified in b).

It was shown in [WB07] that all measurable light modulations, such as diffuse and specular reflections, complex interreflections, diffuse scattering, refraction, caustics, defocus, etc. can be compensated with the multiplication of the inverse light transport matrix and the desired original image. Furthermore, a pixel-precise geometric image correction is implicitly included and becomes feasible - even for surfaces that are unsuited for a conventional structured light scanning. However, due to the extremely long acquisition time of the light transport matrix (up to several hours), this approach will not be practical before accelerated scanning techniques have been developed.

6. Overcoming Technical Limitations

Most of the image correction techniques that are described in this report are constrained by technical limitations of projector and camera hardware. An insufficient resolution or dynamic range of both devices leads to a significant loss of image quality. A too short focal depth results in regionally defocused image areas when projected onto surfaces with an essential depth variance. Slow projection frame-rates will cause the perception of temporally embedded codes. This section is dedicated to giving an overview of novel (at present mainly experimental) approaches that might lead to future improvements of projector-camera systems in terms of focal depth (subsection 6.1), high resolution (subsection 6.2), dynamic range (subsection 6.3), and high speed (subsection 6.4).

6.1. Increasing Focal Depth

Projections onto geometrically complex surfaces with a high depth variance generally do not allow the displayed content to be in focus everywhere. Common DLP or LCD projectors usually maximize their brightness with large apertures. Thus, they suffer from narrow depths of field and can only generate focused imagery on a single fronto-parallel screen. Laser projectors, which are commonly used in planetaria, are an exception. These emit almost parallel light beams, which make very large depths of field possible. However, the cost of a single professional laser projector can exceed the cost of several hundred conventional projectors. In order to increase the depth of field of conventional projectors, several approaches for deblurring unfocused projections with a single or with multiple projectors have been proposed.

Zhang and Nayar [ZN06] presented an iterative, spatially-varying filtering algorithm that compensates for projector defocus. They employed a coaxial projector-camera system to measure the projection's spatially-varying defocus. Therefore, dot patterns as depicted in figure 17a are projected onto the screen and captured by the camera (b). The defocus kernels for each projector pixel can be recovered from the captured images and encoded in the rows of a matrix B . Given the environment light EM including the projector's black level and a desired input image O , the compensation image I can be computed by minimizing the sum-of-squared pixel difference between O and the expected projection $BI + EM$ as

$$\arg \min_{I, 0 \leq I \leq 255} \|BI + EM - O\|^2, \quad (10)$$

which can be solved with a constrained, iterative steepest gradient solver as described in [ZN06].

An alternative approach to defocus compensation for a single projector setup was presented by Brown et al. [BSC06]. Projector defocus is modeled as a convolution of a projected original image O and Gaussian *point spread functions* (PSFs) as $R(x, y) = O(x, y) \otimes H(x, y)$, where the blurred image that can be captured by a camera is R . The PSFs are estimated by projecting features on the canvas and capturing them with a camera. Assuming a spatially-invariant PSF, a compensation image I can be synthesized by applying a Wiener deconvolution filter to the original image:

$$I(x, y) = \mathcal{F}^{-1} \left\{ \frac{\tilde{H}^*(u, v) \tilde{O}(u, v)}{|\tilde{H}(u, v)|^2 + 1/SNR} \right\}. \quad (11)$$

The signal-to-noise ration (SNR) is estimated a priori, \tilde{O} and \tilde{H} are the Fourier transforms of O and H , respectively, and \tilde{H}^* is \tilde{H} 's complex conjugate. \mathcal{F}^{-1} denotes the inverse Fourier transform. Since the defocus kernel H is generally not spatially-invariant (this would only be the case for a fronto-parallel plane) Wiener filtering cannot be applied directly.




Figure 17: Defocus compensation with a single projector: An input image (c) and its defocused projection onto a planar canvas (d). Solving equation 10 results in a compensation image (e) that leads to a sharper projection (f). For this compensation, the spatially-varying defocus kernels are acquired by projecting dot patterns (a) and capturing them with a camera (b). ©2006 ACM [ZN06]

Therefore, basis compensation images are calculated for each of the uniformly sampled feature points using equation 11. The final compensation image is then generated by interpolating the four closest basis responses for each projector pixel.

Oyamada and Saito [OS07] presented a similar approach to single projector defocus compensation. Here, circular PSFs are used for the convolution and estimated by comparing the original image to various captured compensation images that were generated with different PSFs.

The main drawback of these single projector defocus compensation approaches is that the quality is highly dependent on the projected content. All of the discussed methods result in a pre-sharpened compensation image that is visually closer to the original image after being optically blurred by the defocused projection. While soft contours can be compensated, this is generally not the case for sharp features.

Inverse filtering for defocus compensation can also be seen as the division of the original image by the projector's aperture image in frequency domain. Low magnitudes in the Fourier transform of the aperture image, however, lead to intensity values in spatial domain that exceed the displayable range. Therefore, the corresponding frequencies are not considered, which then results in visible ringing artifacts in the final projection. This is the main limitation of the approaches discussed above, since in frequency domain the Gaussian PSF of spherical apertures does contain a large fraction of low Fourier magnitudes. As shown above, applying only small kernel scales will reduce the number of low Fourier magnitudes (and consequently the ringing artifacts) – but will also lead only to minor focus improvements. To overcome this problem, a coded aperture whose Fourier transform has

initially less low magnitudes was applied in [GB08]. Consequently, more frequencies are retained and more image details are reconstructed (cf. figure 18).




Figure 18: The power spectra of the Gaussian PSF of a spherical aperture and of the PSF of a coded aperture: Fourier magnitudes that are too low are clipped (black), which causes ringing artifacts. Image projected in focus, and with the same optical defocus (approx. 2m distance to focal plane) in three different ways: with spherical aperture – untreated and deconvolved with Gaussian PSF, with coded aperture and deconvolved with PSF of aperture code.

An alternative approach that is less dependent on the actual frequencies in the input image was introduced in [BE06]. Multiple overlapping projectors with varying focal depths illuminate arbitrary surfaces with complex geometry and reflectance properties. Pixel-precise focus values $\Phi_{i,x,y}$ are automatically estimated at each camera pixel (x,y) for every projector. Therefore, a uniform grid of circular patterns is displayed by each projector and recorded by a camera. In order to capture the same picture (geometrically and color-wise) for each projection, these are pre-distorted and radiometrically compensated as described in sections 3 and 4.

Once the relative focus values are known, an image from multiple projector contributions with minimal defocus can be composed in real-time. A weighted image composition represents a tradeoff between intensity enhancement and focus refinement as:

$$I_i = \frac{w_i (R - EM)}{\sum_j^N w_j FM_j}, \quad w_{i,x,y} = \frac{\Phi_{i,x,y}}{\sum_j^N \Phi_{j,x,y}}, \quad (12)$$

where I_i is the compensation image for projector i if N projectors are applied simultaneously. Display contributions with high focus values are up-weighted while contributions of projectors with low focus values are down-weighted proportionally. A major advantage of this method, compared to single projector approaches, is that the focal depth of the

entire projection scales with the number of projectors. An example for two projectors can be seen in figure 19.

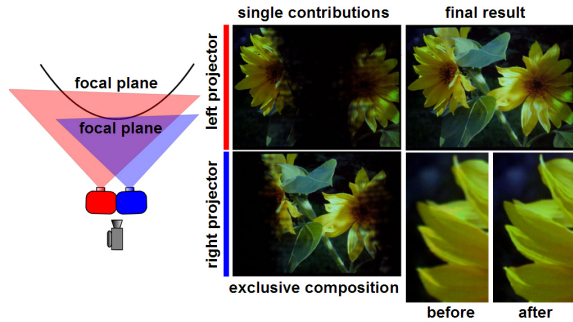


Figure 19: Defocus compensation with two overlapping projectors that have differently adjusted focal planes. ©2006 IEEE [BE06]

6.2. Super-Resolution

Super-resolution techniques can improve the accuracy of geometric warping (see section 3) and consequently have the potential to enhance radiometric compensation (see section 4) due to a more precise mapping of projector pixels onto surface pigments. Over the past years, several researches have proposed super-resolution camera techniques to overcome the inherent limitation of low-resolution imaging systems by using signal processing to obtain super-resolution images (or image sequences) with multiple low-resolution devices [PPK03]. Using a single camera to obtain multiple frames of the same scene is most popular. Multi-camera approaches have also been proposed [WJV*05].

On the other hand, super-resolution projection systems are just beginning to be researched. This section introduces recent work on such techniques that can generally be categorized into two different groups. The first group proposes super-resolution rendering with a single projector [AU05]. Other approaches achieve this with multiple overlapping projectors [JR03,DVC07].

In single projector approaches, so-called *wobulation* techniques are applied: Multiple sub-frames are generated from an original image. An optical image shift displaces the projected image of each sub-frame by a fraction of a pixel [AU05]. Each sub-frame is projected onto the screen with slightly different positions using an opto-mechanical image shifter. This light modulator must be switched fast enough so that all sub-frames are projected in one frame. Consequently, observers perceive this rapid sequence as a continuous and flicker-free image while the resolution is spatially enhanced. Such techniques have been already realized with a DLP system (SmoothPicture®, Texas Instruments Incorporated).

The goal of multi-projector super-resolution methods is

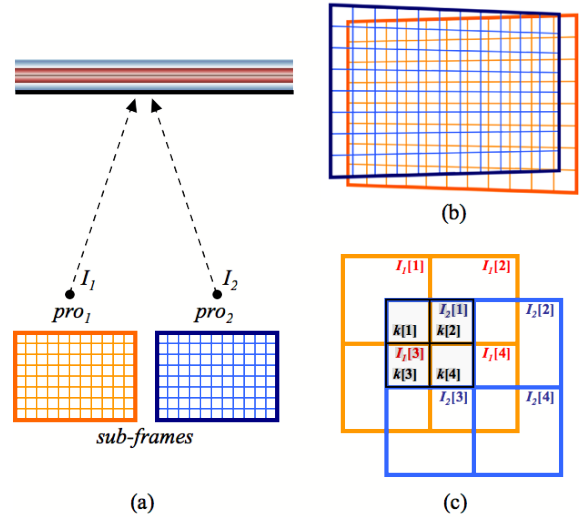


Figure 20: Super-resolution projection with a multi-projector setup (a), overlapping images on the projection screen (b) and close-up of overlapped pixels (c).

to generate a high resolution image with the superposition of multiple low resolution sub-frames produced by different projection units. Thereby, the resolutions of each sub-frame differ and the display surfaces are assumed to be diffuse. Super-resolution pixels are defined by the overlapping sub-frames that are shifted on a sub-pixel basis as shown in figure 20. Generally, the final image is estimated as the sum of the sub-frames. If N sub-frames $I_{i=1..N}$ are displayed, this is modeled as:

$$R = \sum_i^N A_i V_i I_i + EM \quad (13)$$

Note, that in this case the parameters R , I_i , and EM are images, and that A_i and V_i are the geometric warping matrix and the color mixing matrix that transform the whole image (in contrast to sections 3 and 4, where these parameters represent transformations of individual pixels).

Figure 20c shows a close-up of overlapping pixels to illustrate the problem that has to be solved: While $I_1[1..4]$ and $I_2[1..4]$ are the physical pixels of two projectors, $k[1..4]$ represent the desired “super-resolution” pixel structure. The goal is to find the intensities and colors of corresponding projector pixels in I_1 and I_2 that approximate k as close as possible by assuming that the perceived result is $I_1 + I_2$. This is obviously a global optimization problem, since k and I have different resolutions. Thus, if O is the desired original image and R is the captured result, the estimation of sub-frame I_i for projector i is in general achieved by minimizing $\|O - R\|^2$:

$$I_i = \arg \min_{I_i} \|O - R\|^2 \quad (14)$$

Jaynes et al. first demonstrated resolution enhancement with multiple superimposed projections [JR03]. Homographies are used for initial geometric registration of multiple sub-frames onto a planar surface. However, homographic transforms lead to uniform two-dimensional shifts and sampling rates with respect to the camera image rather than to non-uniform ones of general projective transforms.

To reduce this effect, a warped sub-frame is divided into smaller regions that are shifted to achieve sub-pixel accuracy. Initially, each such frame is estimated in the frequency domain by phase shifting the frequencies of the original image. Then, a greedy heuristic process is used to recursively update pixels with the largest global error with respect to equation 14. The proposed model does not consider V_i and EM in equation 13 and a camera is used only for geometric correction. The iterations of the optimization process are terminated manually in [JR03].

Damera-Venkata et al. proposed a real-time rendering algorithm for computing sub-frames that are projected by superimposed lower-resolution projectors [DVC07]. In contrast to the previous method, they use a camera to estimate the geometric and photometric properties of each projector during a calibration step. Image registration is achieved on a sub-pixel basis using gray code projection and coarse-to-fine multi-scale corner analysis and interpolation. In the proposed model, A_i encapsulates the effects of geometric distortion, pixel reconstruction point spread function and resample filtering operations.

Furthermore, V_i and EM are obtained during calibration by analyzing the camera response for projected black, red, green, and blue flood images of each projector. In principle, this model could be applied to a projection surface with arbitrary color, texture and shape. However, this has not been shown in [DVC07]. Once the parameters are estimated, equation 14 can be solved numerically using an iterative gradient descent algorithm. This generates optimal results but does not achieve real-time rendering rates.

For real-time sub-frame rendering, it was shown in [DVC07] that near-optimal results can be produced with a non-iterative approximation. This is accomplished by introducing a linear filter bank that consists of impulse responses of the linearly approximated results which are pre-computed with the non-linear iterative algorithm mentioned above. The filter bank is applied to the original image for estimating the sub-frames.

In an experimental setting, this filtering process is implemented with fragment shaders and real-time rendering is achieved. Figure 21 illustrates a close-up of a single projected sub-frame (a) and four overlapping projections with super-resolution rendering enabled (b). In this experiment, the original image has a higher resolution than any of the sub-frames.




Figure 21: Experimental result for four superimposed projections: Single sub-frame image (a) and image produced by four superimposed projections with super-resolution enabled. ©2007 IEEE [DVC07]

6.3. High Dynamic Range

To overcome the contrast limitations that are related to radiometric compensation (see figure 7), high dynamic range (HDR) projector-camera systems are imaginable. Although there has been much research and development on HDR camera and capturing systems, little work has been done so far on HDR projectors.

In this section, we will focus on state-of-the-art HDR projector technologies rather than on HDR cameras and capturing techniques. A detailed discussion on HDR capturing/imaging technology and techniques, such as recovering camera response functions and tone mapping/reproduction is out of the scope of this report. The interested reader is referred to [RWPD06].

Note, that for the following we want to use the notation of *dynamic range* (unit decibel, dB) for cameras, and the notation of *contrast ratio* (unit-less) for projectors.

The dynamic range of common CCD or CMOS chips is around 60 dB while recent logarithmic CMOS image sensors for HDR cameras cover a dynamic range of 170 dB (HDRC®, Omron Automotive Electronics GmbH). Besides special HDR sensors, low dynamic range (LDR) cameras can be applied for capturing HDR images.

The most popular approach to HDR image acquisition involves taking multiple images of the same scene with the same camera using different exposures, and then merging them into a single HDR image.

There are many ways for making multiple exposure measurements with a single camera [DM97] or with multiple coaxially aligned cameras [AA01]. The interested reader is referred to [NB03] for more information. As an alternative to merging multiple LDR images, the exposure of individual sensor pixels in one image can be controlled with additional light modulators, like an LCD panel [NB03] or a DMD chip [NBB04] in front of the sensor or elsewhere within the optical path. In these cases, HDR images are acquired directly.

The contrast ratio of DMD chips and LCoS panels (without additional optics) is about 2,000:1 [DDS03] and 5,000:1 (SXRD®, Sony Corporation) respectively. Currently, a contrast ratio of around 15,000:1 is achieved for high-end projectors with auto-iris techniques that dynamically adjust the amount of the emitting light according to the image content. Auto-iris techniques, however, cannot expand the dynamic range within a single frame. On the other hand, a laser projection system achieved the contrast ratio of 100,000:1 in [BDD*04] because of the absence of light in dark regions.

Multi-projector systems can enhance spatial resolution (see section 6.2) and increase the intensity range of projections (see section 4.1). However, merging multiple LDR projections does not result in an HDR image. Majumder et al., for example, have rendered HDR images with three overlapped projectors to demonstrate that a larger intensity range and resolution will result in higher quality images [MW01]. Although the maximum intensity level is increased with each additional projector unit, the minimum intensity level (i.e., the black level) is also increased. The contrast of overlapping regions is never greater than the largest one of each individual projector.

Theoretically, if the maximum and the minimum intensities of the i th projector are I_i^{max} and I_i^{min} , its contrast ratio is I_i^{max}/I_i^{min} . If N projectors are overlapped, the contrast ratio of the final image is $\sum_i^N I_i^{max}/\sum_i^N I_i^{min}$. For example, if two projectors are used whose intensities are $I_1^{min} = 10$, $I_1^{max} = 100$ and $I_2^{min} = 100$, $I_2^{max} = 1000$ (thus both contrast ratios are 10 : 1), the contrast ratio of the image overlap is still 10 : 1 ($10 = (I_1^{max} + I_2^{max})/(I_1^{min} + I_2^{min})$).

Recently, HDR display systems have been proposed that combine projectors and external light modulators. Seetzen et al. proposed an HDR display that applies a projector as a backlight of an LCD panel instead of a fluorescent tube assembly [SHS*04]. As in figure 22a, the projector is directed to the rear of a transmissive LCD panel. The light that corresponds to each pixel on the HDR display is effectively modulated twice: first by the projector and then by the LCD panel. Theoretically, the final contrast ratio is the product of the individual contrast ratio of the two modulators. If a projector with a contrast ratio of c_1 : 1 and an LCD panel with a contrast ratio of c_2 : 1 are used in this example, the contrast of the combined images is $(c_1 \cdot c_2)$: 1. In an experimental setup, this approach achieved a contrast ratio of 54,000 : 1 using an LCD panel and a DMD projector with a contrast ratio of 300 : 1 and 800 : 1 respectively. The reduction of contrast is due to noise and imperfections in the optical path.

The example described above does not really present a projection system since the image is generated behind an LCD panel, rather than on a projection surface. True HDR projection approaches are discussed in [DRW*06, DSW*07]. The basic idea of realizing an HDR projector is to combine a normal projector and an additional low resolution light modulating device. Double modulation decreases the black

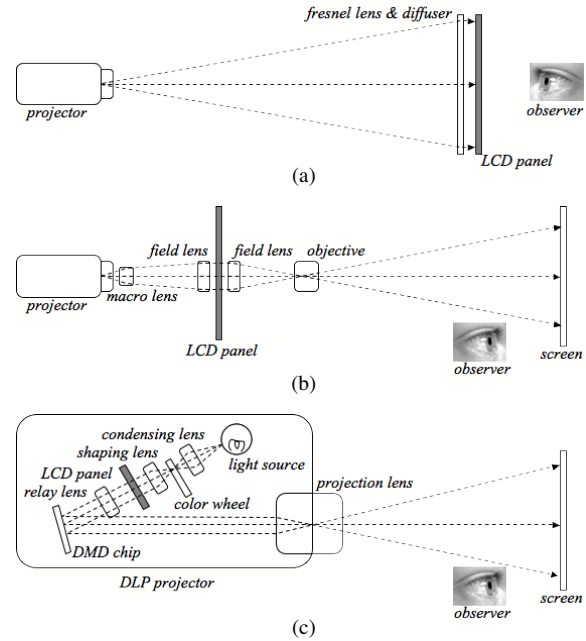


Figure 22: Different HDR projection setups: using a projector as backlight of an LCD (a), modulating the image path (b), and modulating the illumination path (c).

level of the projected image, and increases the dynamic range as well as the number of addressable intensity levels. Thereby, LCD panels, LCoS panels, DMD chips can serve as light modulators.

HDR projectors can be categorized into systems that modulate the image path (cf. figure 22b), and into systems that modulate the illumination path (22c). In the first case, an image is generated with a high resolution light modulator first, and then modulated again with an additional low resolution light modulator. In the latter case, the projection light is modulated in advance with a low resolution light modulator before the image is generated with a high resolution modulator.

In each approach, a compensation for the optical blur caused by the low resolution modulator is required. The degree of blur can be measured and can be described with a point spread function (PSF) for each low resolution pixel in relation to corresponding pixels on the higher resolution modulator. A division of the desired output image by the estimated blurred image that is simulated by the PSF will result in the necessary compensation mask which will be displayed on the high resolution modulator.

Pavlovych et al. proposed a system that falls into the first category [PS05]. This system uses an external attachment (an LCD panel) in combination with a regular DLP projector (cf. figure 22b). The projected image is resized and focused

first on the LCD panel through a set of lenses. Then it is modulated by the LCD panel and projected through another lens system onto a larger screen.




Figure 23: Photographs of a part of an HDR projected image: image modulated with low resolution chrominance modulators (a), image modulated with a high resolution luminance modulator (b), output image (c). ©2006 ITE/SID [KKN*06]

Kusakabe et al. proposed an HDR projector that applies LCoS panels that falls into the second category [KKN*06]. In this system, three low resolution (RGB) modulators are used first for chrominance modulation of the projection light. Finally, the light is modulated again with a high resolution luminance modulator which forms the image.

The resolution of the panel that is applied for chrominance modulation can be much lower than the one for luminance modulation because the human visual system is sensitive only to a relatively low chrominance contrast. An experimental result is shown in figure 23. The proposed projector has a contrast ratio of 1, 100,000 : 1.

6.4. High Speed

High speed projector-camera systems hold the enormous potential to significantly improve high frequent temporal coded projections (see sections 3.3 and 4.2). They enable, for instance, projecting and capturing imperceptible spatial patterns that can be efficiently used for real-time geometric registration, fast shape measurement and real-time adaptive radiometric compensation while a flicker-free content is perceived by the observer at the same time. The faster the projection and the capturing process can be carried out, the more information per unit of time can be encoded. Since high speed capturing systems are well established, this section focuses mainly on the state-of-the-art of high speed projection systems. Both together, however, could be merged into future high speed projector-camera systems. For this reason, we first want to give only a brief overview of high speed capturing systems.

Commercially available single-chip high speed cameras exist that can record 512x512 color pixels at up to 16,000 fps (FASTCAM SA1, Photron Ltd.). However, these systems are typically limited to storing just a few seconds of data directly on the camera because of the huge bandwidth that is necessary to transfer the images. Other CMOS devices are on the market that enable a 500 fps (A504k, Basler AG) capturing and transfer rates.

Besides such single-camera systems, a high capturing speed can also be achieved with multi-camera arrays. Wilburn et al., for example, proposed a high speed video system for capturing 1,560 fps videos using a dense array of 30 fps CMOS image sensors [WJV*04]. Their system captures and compresses images from 52 cameras in parallel. Even at extremely high frame-rates, such a camera array architecture supports continuous streaming to disk from all of the cameras for minutes.

In contrast to this, however, the frame-rate of commercially available DLP projectors is normally less than or equal to 120 fps (DepthQ®, InFocus Corporation). Although faster projectors that can be used in the context of our projector-camera system are currently not available, we want to outline several projection approaches that achieve higher frame-rates - but do not necessarily allow the projection of high quality images.

Raskar et al., for instance, developed a high speed optical motion capture system with an LED-based code projector [RNdD*07]. The system consists of a set of 1-bit gray code infrared LED beamers. Such a beamer array is effectively emitting 10,000 binary gray coded patterns per second, and is applied for object tracking. Each object to be tracked is tagged with a photosensor that detects and decodes the temporally projected codes. The 3D location of the tags can be computed at a speed of 500 Hz when at least three such beamer arrays are applied.

In contrast to this approach which does not intent to project pictorial content in addition to the code patterns, Nii et al. proposed a visible light communication (VLC) technique that does display simple images [NSI05]. They developed an LED-based high speed projection system (with a resolution of 4x5 points produced with an equally large LED matrix) that is able to project alphabetic characters while applying an additional pulse modulation for coding information that is detected by photosensors. This system is able to transmit two data streams with 1 kHz and 2 kHz respectively at different locations while simultaneously projecting simple pictorial content. Although LEDs can be switched with a high speed (e.g., the LEDs in [NSI05] are temporally modulated at 10.7 MHz), such simple LED-based projection systems offer a too low spatial resolution at the moment.

In principle, binary frame-rates of up to 16,300 fps can currently be achieved with DMDs for a resolution of 1024x768. The DMD discovery board enables developers to implement their own mirror timings for special purpose application

[DDS03]. Consequently, due to this high binary frame-rate some researchers utilized the Discovery boards for realizing high speed projection techniques. McDowall et al., for example, demonstrated the possibility of projecting 24 binary code and compensation images at a speed of 60 Hz [MBHF04]. Viewers used time-encoded shutter glasses to make individual images visible.

Kitamura et al. also developed a high speed projector based on the DMD discovery board [KN06]. In their approach, photosensors can be used to detect temporal code patterns that are embedded into the mirror flip sequence. In contrast to the approach by Cotting et al. [CNGF04] that was described in section 3.3, the mirror flip sequence can be freely re-configured.

The results of an initial basic experiment with this system are shown in figure 24a: The projected image is divided into 10 regions. Different on/off mirror flip frequencies are used in each region (from 100 Hz to 1,000 Hz at 100 Hz intervals), while a uniformly bright image with a 50 % intensity appears in all regions - regardless of the locally applied frequencies. The intensity fall-off in the projection is mainly due to imperfections in applied optics. The signal waves are received by photosensors that are placed within the regions. They can detect the individual frequency.




Figure 24: Regionally different mirror flip frequencies and corresponding signal waves received by photosensors at different image areas. The overall image appears mostly uniform in intensity (a). Binary codes can be embedded into the first half of the exposure sequence while the second half can compensate the desired intensity (b). ©2006 IPSJ [KN06]

Instead of using a constant on-off flip frequency for each region, binary codes can be embedded into a projected frame. This is illustrated in figure 24b: For a certain time slot of T ,

the first half of the exposure sequence contains a temporal code pattern (modulated with different mirror flip states) that is compensated with the second half of the exposure sequence to modulate a desired intensity. Yet, contrast is lost in this case due to the modulated intensity level created by the code pattern. Here, the number of on-states always equals the number of off-states in the code period. This leads to a constant minimum intensity level of 25 %. Since also 25 % of the off states are used during this period, intensity values between 25 % and 75 % can only be displayed.

All systems that have been outlined above, apply photosensors rather than cameras. Thus, they cannot be considered as suitable projector-camera systems in our application context. Yet, McDowall et al. combined their high speed projector with a high speed camera to realize fast range scanning [MB05]. Takei et al. proposed a 3,000 fps shape measurement system (shape reconstruction is performed off-line in this case) [TKH07].

In an image-based rendering context, Jones et al. proposed to simulate spatially varying lighting on a live performance based on a fast shape measurement using a high-speed projector-camera system [JGB*06]. However, all of these approaches do not project pictorial image content, but rather represent encouraging examples of fast projector-camera techniques.

The mirrors on a conventional DMD chip can be switched much faster than alternative technologies, such as ordinary LCD or LCoS panels whose refresh rate can be up to 2.5 ms (= 400 Hz) at the moment.

LEDs are generally better suited for high-speed projectors than a conventional UHP lamp (we do not want to consider brightness issues for the moment), because three or more different LEDs that correspond to each color component can be switched at a high speed (even faster than a DMD) for modulating colors and intensities. Therefore, a combination of DMD and LED technologies seems to be optimal for future projection units.

Let's assume that the mirrors of a regular DLP projector can be switched at $15\mu\text{s}$ (= 67,000 binary frames per second). For projecting 256 different intensity levels (i.e., an 8 bit encoded gray scale image), the gray scale frame rate is around 260 Hz (= 67,000 binary frames per second / 256 intensity levels). Consequently, the frame rate for full color images is around 85 Hz (= 260 gray scale frames per second / 3 color channels) if the color wheel consists of three filter segments.

Now, let's consider DLP projectors that apply LEDs instead of a UHP lamps and a color wheel. If, for example, the intensities of three (RGB) color LEDs can be switched between eight different levels (1,2,4,8,16,32,64,128,256) at a high speed, a full color image can theoretically be projected at around 2,800 Hz (= 67,000 binary frames per second / 8 (8-bit encoded) intensity levels / 3 color channels).

To overcome the bandwidth limitation for transferring

the huge amount of image data in high-speed, the MULE projector adopts a custom programmed FPGA-based circuitry [JMY*07]. The FPGA decodes a standard DVI signal from the graphics card. Instead of rendering a color image, the FPGA takes each 24 bit color frame of video and displays each bit sequentially as separate frames. Thus, if the incoming digital video signal is 60 Hz, the projector displays $60 \times 24 = 1,440$ frames per second. To achieve even faster rates, the refresh rate of a video card is set at 180-240 Hz. At 200 Hz, for instance, the projector can display 4,800 binary frames per second.

7. Conclusion

This article reviewed the state-of-the-art of projector-camera systems with a focus on real-time image correction techniques that enable projections onto non-optimized surfaces. It did not discuss projector-camera related areas, such as camera supported photometric calibration of conventional projection displays (e.g., [BMY05], [JM07], [BM07]), real-time shadow removal techniques (e.g., [STJS01], [JWS*01], [JWS04]), or projector-camera based interaction approaches (e.g., [Pin01], [EHH04], [FR06]).

While most of the presented techniques are still on a research level, others found already practical applications in theatres, museums, historic sites, open-air festivals, trade shows, and advertisement. Some examples are shown in figures 25-27.

Future projectors will become more compact in size and will require little power and cooling. Reflective technology (such as DLP or LCOS) will more and more replace transmissive technology (e.g., LCD). This leads to an increased brightness and extremely high update rates. They will integrate GPUs for real-time graphics and vision processing. While resolution and contrast will keep increasing, production costs and market prizes will continue to fall. Conventional UHP lamps will be replaced by powerful LEDs or multi-channel lasers. This will make them suitable for mobile applications.

Imagining projector-camera technology to be integrated into, or coupled with mobile devices, such as cellphones or laptops, will support a truly flexible way for presentations. There is no doubt that this technology is on its way. Yet, one question needs to be addressed when thinking about mobile projectors: What can we project onto, without carrying around screen canvases? It is clear that the answer to this question can only be: Onto available everyday surfaces. With this in mind, the future importance of projector-camera systems in combination with appropriate image correction techniques becomes clear.

Acknowledgments

We wish to thank the entire ARGroup at the Bauhaus-University Weimar who were involved in developing

projector-camera techniques over the last years, as well as the authors who gave permission to use their images in this article. Special thanks go to Stefanie Zollmann and Mel for proof-reading. Projector-camera activities at BUW were partially supported by the Deutsche Forschungsgemeinschaft (DFG) under contract numbers BI 835/1-1 and PE 1183/1-1.

References

- [AA01] AGGARWAL M., AHUJA N.: Split Aperture Imaging for High Dynamic Range. In *Proc. of IEEE International Conference on Computer Vision (ICCV)* (2001), vol. 2, pp. 10–17.
- [AOSS06] ASHDOWN M., OKABE T., SATO I., SATO Y.: Robust Content-Dependent Photometric Projector Compensation. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2006).
- [ASOS07] ASHDOWN M., SATO I., OKABE T., SATO Y.: Perceptual Photometric Compensation for Projected Images. *IEICE Transaction on Information and Systems J90-D*, 8 (2007), 2115–2125. in Japanese.
- [AU05] ALLEN W., ULICHNEY R.: Wobulation: Doubling the Addressed Resolution of Projection Displays. In *Proc. of SID Symposium Digest of Technical Papers* (2005), vol. 36, pp. 1514–1517.
- [BCK*05] BIMBER O., CORIAND F., KLEPPE A., BRUNS E., ZOLLMANN S., LANGLOTZ T.: Superimposing Pictorial Artwork with Projected Imagery. *IEEE MultiMedia* 12, 1 (2005), 16–26.
- [BDD*04] BIEHLING W., DETER C., DUBE S., HILL B., HELLING S., ISAKOVIC K., KLOSE S., SCHIEWE M.: LaserCave - Some Building Blocks for Immersive Screens -. In *Proc. of International Status Conference Virtual and Augmented Reality* (2004).
- [BE06] BIMBER O., EMMERLING A.: Multifocal Projection: A Multiprojector Technique for Increasing Focal Depth. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 12, 4 (2006), 658–667.
- [BEK05] BIMBER O., EMMERLING A., KLEMMER T.: Embedded Entertainment with Smart Projectors. *IEEE Computer* 38, 1 (2005), 56–63.
- [BGZ*06] BIMBER O., GRUNDHÖFER A., ZEIDLER T., DANCH D., KAPAKOS P.: Compensating Indirect Scattering for Immersive and Semi-Immersive Projection Displays. In *Proc. of IEEE Virtual Reality (IEEE VR)* (2006), pp. 151–158.
- [Bim06] BIMBER O.: Projector-Based Augmentation. In *Emerging Technologies of Augmented Reality: Interfaces and Design*, Haller M., Billinghurst M., Thomas B., (Eds.). Idea Group, 2006, pp. 64–89.
- [BJM07] BHASKER E. S., JUANG R., MAJUMDER A.: Registration techniques for using imperfect and partially

- calibrated devices in planar multi-projector displays. *IEEE Trans. Vis. Comput. Graph.* 13, 6 (2007), 1368–1375.
- [BM07] BHASKER E., MAJUMDER A.: Geometric Modeling and Calibration of Planar Multi-Projector Displays using Rational Bezier Patches. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2007).
- [BMS98] BATLLE J., MOUADDIB E. M., SALVI J.: Recent progress in coded structured light as a technique to solve the correspondence problem: a survey. *Pattern Recognition* 31, 7 (1998), 963–982.
- [BMY05] BROWN M., MAJUMDER A., YANG R.: Camera Based Calibration Techniques for Seamless Multi-Projector Displays. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 11, 2 (2005), 193–206.
- [BSC06] BROWN M. S., SONG P., CHAM T.-J.: Image Pre-Conditioning for Out-of-Focus Projector Blur. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2006), vol. II, pp. 1956–1963.
- [BWEN05] BIMBER O., WETZSTEIN G., EMMERLING A., NITSCHKE C.: Enabling View-Dependent Stereoscopic Projection in Real Environments. In *Proc. of IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR)* (2005), pp. 14–23.
- [CKS98] CASPI D., KIRYATI N., SHAMIR J.: Range Imaging With Adaptive Color Structured Light. *IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE* 20, 5 (1998), 470–480.
- [CNGF04] COTTING D., NÄF M., GROSS M. H., FUCHS H.: Embedding Imperceptible Patterns into Projected Images for Simultaneous Acquisition and Display. In *Proc. of IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR)* (2004), pp. 100–109.
- [CZGF05] COTTING D., ZIEGLER R., GROSS M. H., FUCHS H.: Adaptive Instant Displays: Continuously Calibrated Projections using Per-Pixel Light Control. In *Proc. of Eurographics* (2005), pp. 705–714.
- [DDS03] DUDLEY D., DUNCAN W. M., SLAUGHTER J.: Emerging Digital Micromirror Device (DMD) Applications. In *Proc. of SPIE* (2003), vol. 4985, pp. 14–25.
- [DM97] DEBEVEC P. E., MALIK J.: Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. of ACM SIGGRAPH* (1997), pp. 369–378.
- [DRW*06] DEBEVEC P., REINHARD E., WARD G., MYSZKOWSKI K., SEETZEN H., ZARGARPOUR H., MCTAGGART G., HESS D.: High Dynamic Range Imaging: Theory and Applications. In *Proc. of ACM SIGGRAPH (Courses)* (2006).
- [DSW*07] DAMBERG G., SEETZEN H., WARD G., HEIDRICH W., WHITEHEAD L.: High-Dynamic-Range Projection Systems. In *Proc. of SID Symposium Digest of Technical Papers* (2007), vol. 38, pp. 4–7.
- [DVC07] DAMERA-VENKATA N., CHANG N. L.: Realizing Super-Resolution with Superimposed Projection. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2007).
- [EHH04] EHNE J., HIROTA K., HIROSE M.: Projected Augmentation - Augmented Reality using Rotatable Video Projectors. In *Proc. of IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR)* (2004), pp. 26–35.
- [FGN05] FUJII K., GROSSBERG M., NAYAR S.: A Projector-Camera System with Real-Time Photometric Adaptation for Dynamic Environments. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2005), vol. I, pp. 814–821.
- [FR06] FLAGG M., REHG J. M.: Projector-Guided Painting. In *Proc. of ACM Symposium on User Interface Software and Technology (UIST)* (2006), pp. 235–244.
- [GB07] GRUNDHÖFER A., BIMBER O.: Real-Time Adaptive Radiometric Compensation. To appear in *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2007).
- [GB08] GROSSE M., BIMBER O.: Coded aperture projection, 2008.
- [GPNB04] GROSSBERG M., PERI H., NAYAR S., BELHUMEUR P.: Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Jun 2004), vol. I, pp. 452–459.
- [GSHB07] GRUNDHÖFER A., SEEGER M., HÄNTSCH F., BIMBER O.: Dynamic Adaptation of Projected Imperceptible Codes. *Proc. of IEEE International Symposium on Mixed and Augmented Reality* (2007).
- [HSM07] HABE H., SAEKI N., MATSUYAMA T.: Inter-Reflection Compensation for Immersive Projection Display. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams) (poster)* (2007).
- [JF07] JOHNSON T., FUCHS H.: Real-Time Projector Tracking on Complex Geometry using Ordinary Imagery. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2007).
- [JGB*06] JONES A., GARDNER A., BOLAS M., MCDOWALL I., DEBEVEC P.: Simulating Spatially Varying Lighting on a Live Performance. In *Proc. of European Conference on Visual Media Production (CVMP)* (2006), pp. 127–133.
- [JM07] JUANG R., MAJUMDER A.: Photometric Self-Calibration of a Projector-Camera System. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2007).
- [JMY*07] JONES A., MCDOWALL I., YAMADA H., BOLAS M., DEBEVEC P.: Rendering for an Interactive 360° Light Field Display. In *Proc. of ACM SIGGRAPH* (2007).

- [JR03] JAYNES C., RAMAKRISHNAN D.: Super-Resolution Composition in Multi-Projector Displays. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2003).
- [JWS*01] JAYNES C., WEBB S., STEELE R., BROWN M., SEALES W.: Dynamic Shadow Removal from Front Projection Displays. In *Proc. of IEEE Visualization* (2001), pp. 175–555.
- [JWS04] JAYNES C., WEBB S., STEELE R. M.: Camera-Based Detection and Removal of Shadows from Interactive Multiprojector Displays. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 10, 3 (2004), 290–301.
- [KKN*06] KUSAKABE Y., KANAZAWA M., NOJIRI Y., FURUYA M., YOSHIMURA M.: YC-separation Type Projector with Double Modulation. In *Proc. of International Display Workshop (IDW)* (2006), pp. 1959–1962.
- [KN06] KITAMURA M., NAEMURA T.: A Study on Position-Dependent Visible Light Communication using DMD for ProCam. In *IPSJ SIG Notes. CVIM-156* (2006), pp. 17–24. in Japanese.
- [MB05] MCDOWALL I. E., BOLAS M.: Fast Light for Display, Sensing and Control Applications. In *Proc. of IEEE VR 2005 Workshop on Emerging Display Technologies (EDT)* (2005), pp. 35–36.
- [MBHF04] MCDOWALL I. E., BOLAS M. T., HOBERMAN P., FISHER S. S.: Snared Illumination. In *Proc. of ACM SIGGRAPH (Emerging Technologies)* (2004), p. 24.
- [MKO06] MUKAIGAWA Y., KAKINUMA T., OHTA Y.: Analytical Compensation of Inter-reflection for Pattern Projection. In *Proc. of ACM Symposium on Virtual Reality Software and Technology (VRST) (short paper)* (2006), pp. 265–268.
- [MW01] MAJUMDER A., WELCH G.: COMPUTER GRAPHICS OPTIQUE: Optical Superposition of Projected Computer Graphics. In *Proc. of Immersive Projection Technology - Eurographics Workshop on Virtual Environment (IPT-EGVE)* (2001).
- [NB03] NAYAR S. K., BRANZOI V.: Adaptive Dynamic Range Imaging: Optical Control of Pixel Exposures over Space and Time. In *Proc. of IEEE International Conference on Computer Vision (ICCV)* (2003), vol. 2, pp. 1168–1175.
- [NBB04] NAYAR S. K., BRANZOI V., BOULT T. E.: Programmable Imaging using a Digital Micromirror Array. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2004), vol. 1, pp. 436–443.
- [NKGR06] NAYAR S., KRISHNAN G., GROSSBERG M. D., RASKAR R.: Fast Separation of Direct and Global Components of a Scene using High Frequency Illumination. In *Proc. of ACM SIGGRAPH* (2006), pp. 935–944.
- [NPGB03] NAYAR S. K., PERI H., GROSSBERG M. D., BELHUMEUR P. N.: A Projection System with Radiometric Compensation for Screen Imperfections. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2003).
- [NSI05] NII H., SUGIMOTO M., INAMI M.: Smart Light-Ultra High Speed Projector for Spatial Multiplexing Optical Transmission. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2005).
- [OS07] OYAMADA Y., SAITO H.: Focal Pre-Correction of Projected Image for Deblurring Screen Image. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2007).
- [Pin01] PINHANEZ C.: Using a Steerable Projector and a Camera to Transform Surfaces into Interactive Displays. In *Proc. of CHI (extended abstracts)* (2001), pp. 369–370.
- [PLJP07] PARK H., LEE M.-H., JIN B.-K. S. Y., PARK J.-I.: Content adaptive embedding of complementary patterns for nonintrusive direct-projected augmented reality. In *HCI International 2007* (2007), vol. 14.
- [PLKP05] PARK H., LEE M.-H., KIM S.-J., PARK J.-I.: Specularity-Free Projection on Nonplanar Surface. In *Proc. of Pacific-Rim Conference on Multimedia (PCM)* (2005), pp. 606–616.
- [PLKP06] PARK H., LEE M.-H., KIM S.-J., PARK J.-I.: Contrast Enhancement in Direct-Projected Augmented Reality. In *Proc. of IEEE International Conference on Multimedia and Expo (ICME)* (2006).
- [PLS*06] PARK H., LEE M.-H., SEO B.-K., SHIN H.-C., PARK J.-I.: Radiometrically-Compensated Projection onto Non-Lambertian Surface using Multiple Overlapping Projectors. In *Proc. of Pacific-Rim Symposium on Image and Video Technology (PSIVT)* (2006), pp. 534–544.
- [PPK03] PARK S. C., PARK M. K., KANG M. G.: Super-Resolution Image Reconstruction: A Technical Overview. *IEEE Signal Processing Magazine* 20, 3 (2003), 21–36.
- [PS05] PAVLOVYCH A., STUERZLINGER W.: A High-Dynamic Range Projection System. In *Proc. of SPIE* (2005), vol. 5969.
- [Ras99] RASKAR R.: Oblique Projector Rendering on Planar Surfaces for a Tracked User. In *Proc. of ACM SIGGRAPH (Sketches and Applications)* (1999).
- [RBvB*04] RASKAR R., BEARDSLEY P., VAN BAAR J., WANG Y., DIETZ P., LEE J., LEIGH D., WILLWACHER T.: RFIG Lamps: Interacting with a Self-Describing World via Photosensing Wireless Tags and Projectors. In *Proc. of ACM SIGGRAPH* (2004), pp. 406–415.
- [RBY*99] RASKAR R., BROWN M., YANG R., CHEN W., WELCH G., TOWLES H., SEALES B., FUCHS H.: Multi-Projector Displays using Camera-Based Registration. In *Proc. of IEEE Visualization* (1999), pp. 161–168.
- [RNdD*07] RASKAR R., NII H., DE DECKER B.,

- HASHIMOTO Y., SUMMET J., MOORE D., ZHAO Y., WESTHUES J., DIETZ P., INAMI M., NAYAR S. K., BARNWELL J., NOLAND M., BEKAERT P., BRANZOI V., BRUNS E.: *Prakash: Lighting Aware Motion Capture using Photosensing Markers and Multiplexed Illuminators*. In *Proc. of ACM SIGGRAPH* (2007).
- [RPG99] RAMASUBRAMANIAN M., PATTANAİK S. N., GREENBERG D. P.: *A Perceptually Based Physical Error Metric for Realistic Image Synthesis*. In *Proc. of ACM SIGGRAPH* (1999), pp. 73–82.
- [RvBWR04] RASKAR R., VAN BAAR J., WILLWACHER T., RAO S.: *Quadric transfer for immersive curved screen displays*. In *Proc. of Eurographics* (Aug 2004).
- [RWC*98] RASKAR R., WELCH G., CUTTS M., LAKE A., STESIN L., FUCHS H.: *The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays*. In *Proc. of ACM SIGGRAPH* (1998), pp. 179–188.
- [RWPD06] REINHARD E., WARD G., PATTANAİK S., DEBEVEC P.: *High Dynamic Range Imaging - Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, 2006.
- [SCG*05] SEN P., CHEN B., GARG G., MARSCHNER S. R., HOROWITZ M., LEVOY M., LENSCH H. P. A.: *Dual Photography*. In *Proc. of ACM SIGGRAPH* (2005), pp. 745–755.
- [SHS*04] SEETZEN H., HEIDRICH W., STUERZLINGER W., WARD G., WHITEHEAD L., TRENTACOSTE M., GHOSH A., VOROZCOVS A.: *High Dynamic Range Display Systems*. In *Proc. of ACM SIGGRAPH* (2004), pp. 760–768.
- [SMK05] SEITZ S. M., MATSUSHITA Y., KUTULAKOS K. N.: *A Theory of Inverse Light Transport*. In *Proc. of IEEE International Conference on Computer Vision (ICCV)* (2005), vol. 2, pp. 1440–1447.
- [SMO03] SHIRAI Y., MATSUSHITA M., OHGURO T.: *HIEI Projector: Augmenting a Real Environment with Invisible Information*. In *Proc. of Workshop on Interactive Systems and Software (WISS)* (2003), pp. 115–122. in Japanese.
- [SPB04] SALVI J., PAGÈS J., BATLLE J.: *Pattern Codification Strategies in Structured Light Systems*. *Pattern Recognition* 37, 4 (2004), 827–849.
- [STJS01] SUKTHANKAR R., TAT-JEN C., SUKTHANKAR G.: *Dynamic Shadow Elimination for Multi-Projector Displays*. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2001), vol. II, pp. 151–157.
- [TKH07] TAKEI J., KAGAMI S., HASHIMOTO K.: *3,000-fps 3-D Shape Measurement Using a High-Speed Camera-Projector System*. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2007), pp. 3211–3216.
- [VVSC05] VIEIRA M. B., VELHO L., SA A., CARVALHO P. C.: *A Camera-Projector System for Real-Time 3D Video*. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2005).
- [WB07] WETZSTEIN G., BIMBER O.: *Radiometric Compensation through Inverse Light Transport*. *Proc. of Pacific Graphics* (2007).
- [WJV*04] WILBURN B., JOSHI N., VAISH V., LEVOY M., HOROWITZ M.: *High-Speed Videography using a Dense Camera Array*. In *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2004), vol. II, pp. 294 – 301.
- [WJV*05] WILBURN B., JOSHI N., VAISH V., TALVALA E.-V., ANTUNEZ E., BARTH A., ADAMS A., HOROWITZ M., LEVOY M.: *High Performance Imaging using Large Camera Arrays*. In *Proc. of ACM SIGGRAPH* (2005), pp. 765–776.
- [WSOS05] WANG D., SATO I., OKABE T., SATO Y.: *Radiometric Compensation in a Projector-Camera System Based on the Properties of Human Vision System*. In *Proc. of IEEE International Workshop on Projector-Camera Systems (ProCams)* (2005).
- [WWC*05] WASCHBÜSCH M., WÜRMLIN S., COTTING D., SADLO F., GROSS M. H.: *Scalable 3D Video of Dynamic Scenes*. *The Visual Computer* 21, 8-10 (2005), 629–638.
- [YHS03] YOSHIDA T., HORII C., SATO K.: *A Virtual Color Reconstruction System for Real Heritage with Light Projection*. In *Proc. of International Conference on Virtual Systems and Multimedia (VSMM)* (2003), pp. 161–168.
- [YW01] YANG R., WELCH G.: *Automatic and Continuous Projector Display Surface Calibration using Every-Day Imagery*. In *Proc. of International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG)* (2001).
- [ZB07] ZOLLMANN S., BIMBER O.: *Imperceptible Calibration for Radiometric Compensation*. In *Proc. of Eurographics (short paper)* (2007), pp. 61–64.
- [ZLB06] ZOLLMANN S., LANGLOTZ T., BIMBER O.: *Passive-Active Geometric Calibration for View-Dependent Projections onto Arbitrary Surfaces*. *Proc. of Workshop on Virtual and Augmented Reality of the GI-Fachgruppe AR/VR 2006 (re-print to appear in Journal of Virtual Reality and Broadcasting 2007)* (2006).
- [ZN06] ZHANG L., NAYAR S. K.: *Projection Defocus Analysis for Scene Capture and Image Display*. In *Proc. of ACM SIGGRAPH* (2006), pp. 907–915.




Figure 25: 360 degree surround projection onto the natural stone walls of castle Osterburg in Weida, Germany.



Figure 26: Outdoor projection onto a building front in Leinefelde, Germany.



Figure 27: Projecting onto the stage settings during live performance of the Karl May festival in Elspe, Germany.




Figure 28: View-dependent stereoscopic projection onto concrete walls (top) and onto wall paper (bottom).
© 2005 IEEE [BWEN05]






Figure 29: Projection onto an iron plate in the Henrichshütte in Hattingen, Germany (top). View-dependent stereoscopic projection onto a natural stone wall (bottom).
© 2005 IEEE [BWEN05]

Mobile Projectors and Optical Communication



Ramesh Raskar

Media Lab, MIT
Cambridge, MA

Mobile Projectors
and
Optical Communication

Mobility and Communication

- Portable Projectors
 - Technology and issues
 - Projector Interaction
 - Single-handed interaction, Image stabilization and resizing
 - Research Prototypes
 - iLamps: Geometrically aware pocket projectors
-
- Optical communication for space labeling in robotics, games
 - Optical and Radio Frequency Tags
 - RFID for Augmented Reality: Location sensing RFID and automatic authoring
 - Imperceptible projection
 - High speed motion capture (Prakash system)

Micro and Pico portable projectors are still in development phase. Excellent review at http://www.economist.com/science/tq/displaystory.cfm?story_id=10789401. So called pocket projectors, e.g. from Mitsubishi Electric, are already available.

We will look at research in mobile projectors and the interaction with real world. Then we will explore the opportunities in optical communication with projectors.


Race for Space

Smallest display device with Largest image

- Rollable Screens
- Projective Devices
- Eye-worn display
- Retinal Display



Courtesy:
Microvision, Inc.



There are only a limited solution to create a large display from a small device.


Projectors are showing the potential to create new ways of interacting with information in everyday life. Desktop screens, laptops and TVs have a basic constraint on their size – they can never be smaller than the display area. Hand-helds such as PDAs are compact but the display size is too limited for many uses. In contrast, projectors of the near future will be compact, portable, and with the built-in awareness which will enable them to automatically create satisfactory displays on many of the surfaces in the everyday environment.

Mitsubishi Pocket Projector






TI DLP + 3 Lasers



On the left is a standard DLP chip that you might find in a large-screen TV. On the right is a DLP for a projector phone. Although smaller, the chip can still project fairly large images. By putting these chips into phones, TI hopes to make watching video on phones more pleasing and, of course, sell more chips. A complete pico projector. It fits into a phone. The projector contains three lasers, a DLP chip and a power supply and measures about 1.5 inches in length. A mock-up of TI's functioning projector in a phone. Granted, the phone's a little larger than a lot of the phones on the market today, but those don't have projectors in them.


http://www.news.com/TI-demos-its-movie-projector-in-a-phone/2100-1041_3-6170619.html?tag=ne.gall.related

Microvision: Rotating mirror + Lasers



Microvision's ultra-miniature full-color digital projection display approximately the size of a Thin Mint candy, designed to be embedded into handheld electronic devices such as cellphones, PDAs, or multimedia handhelds. (January 2007 Consumer Electronic Show). The projector developed at Microvision is composed of two main parts: a set of red, blue, and green lasers made of semiconductor material, such as gallium indium arsenide, and a mirror--one millimeter across--that tilts on two axes. The lasers shine on the mirror, and the mirror reflects the pixel of light onto a wall or other surface. The intensities of the lasers change to produce different colors: when all three are pumping out light full blast, the pixel is white; when all three are off, the pixel is black. Other colors are produced from various combinations in between.


As the lasers flash on the mirror, the mirror gimbals on its two axes, flickering to produce 30 million pixels a second, each illuminating a surface for 20 nanoseconds. Using this laser and single-mirror setup, the projector paints a scene onto a surface one pixel at a time, says Sprague. It does this so quickly that our eyes perceive a static image or a continuous movie. One of the challenges is to design a rapidly gyrating mirror that can coordinate with the lasers that turn on and off 100 million times a second. Integrated into the Microvision mirror are silicon mechanical structures that measure strain on the mirror, detecting what position it's in. This information is fed back into the laser modulator--the device that determines when a laser is emitting light or not--and the feedback loop allows the system to constantly adjust, depending on the demands of the projected image. The mirror, its mount, and the other mechanical components are all made of silicon, putting the projector in a class of device called MEMS (microelectromechanical systems). <http://www.technologyreview.com/Biztech/17860/>



LBO's approach to miniature projection has a range of differentiating features and benefits. The term "holographic" refers not to the projected image, but to the method of projection. A diffraction pattern of the desired 2D image, calculated using LBO's patented holographic algorithms, is displayed on a custom-designed phase-modulating reflective Liquid Crystal on Silicon (LCoS) microdisplay. When illuminated by coherent laser light, the desired 2D image is projected. The projector uses a reflective LCoS array to create a constantly varying diffraction pattern that is carefully calculated to produce the desired two-dimensional image when illuminated by red, green and blue lasers.

The resulting image is in focus at all distances from the projector, and the projector has no moving parts. The drawback of this approach is that the LCoS array takes up more space than a single mirror. But the nature of the technology means that the size of the array can be kept to a minimum. Because the array displays a diffraction pattern, not the actual image, the resolution of the projected image can be higher than that of the array. LBO's current prototype projects a 1,600-by-1,024-pixel image using an 864-by-480-pixel array.

Interaction via Handheld Projector



Pocket Projector 2004-05





iLamps 2002



RFIG Lamps 2003-04

There has been significant research in modifying and interfacing projectors for novel scenarios and applications. Two of the projects we will discuss : iLamps and RFIG lamps.

iLamps : Intelligent Locale-aware Mobile Projectors






Ramesh Raskar, Jeroen van Baar,
Paul Beardsley, Thomas Willwacher,
Srinivas Rao, Clifton Forlines

Siggraph 2003

iLamps stands for Intelligent Locale-aware Mobile Projectors. The idea is to augment a mobile phone with a projector, camera and tilt sensor.

Geometrically Aware Projector



Ramesh Raskar, Jeroen van Baar, Paul Beardsley, Thomas Willwacher, Srinivas Rao, Clifton Forlines
Siggraph 2003

Projectors are currently undergoing a transformation as they evolve from static output devices to portable, environment-aware, communicating systems. An enhanced projector can determine and respond to the geometry of the display surface, and can be used in an ad-hoc cluster to create a self-configuring display. Information display is such a prevailing part of everyday life that new and more flexible ways to present data are likely to have significant impact. This project examines geometrical issues for enhanced projectors, relating to customized projection for different shapes of display surface, object augmentation, and co-operation between multiple units.


Single-handed Desktop-like Interaction



Selecting
tags


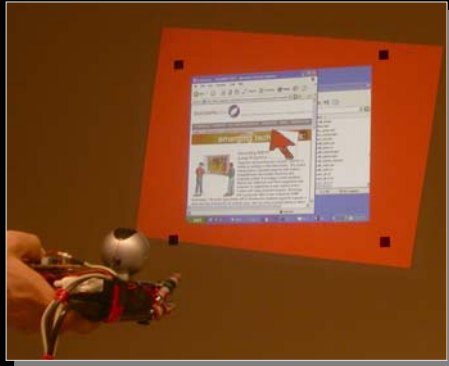
Consider interactive projection to allow a user to interact with projected information e.g. to navigate or update the projected information. A single-handed desktop-like interaction with projected illumination is possible.

Support for handheld projection





Mouse Simulation

- Cursor follows handheld projector motion
- Pre-warped image remains stable



Here we introduce interactive projection, allowing a user with a handheld projector to do mouse-style interaction with projected information. This is achieved by treating a projection as conceptually having two parts – a stabilized component that is static in the display surface, and a cursor that follows any user pointing motion of the projector – effectively allowing the user to track a cursor across a projection. Accompanying mouse buttons are used to do selection. Interactive projection is not specific to a tagged environment, but the technologies meld well.

Image Quasi-Stabilization



Eliminate hand jitter using inertial sensors+camera

In creating a handheld projector, an immediate problem that arises is that hand-jitter results in jitter in the projection. The core requirement to deal with this problem is to compute the pose of the projector relative to the display surface.

Quasi-stabilization

preserves the form of the projection up to an unknown translation in the plane i.e. it preserves projection shape, size and orientation.

But the projection translates in the display plane in accordance with projector motion parallel to the plane.

Absolute Stabilization

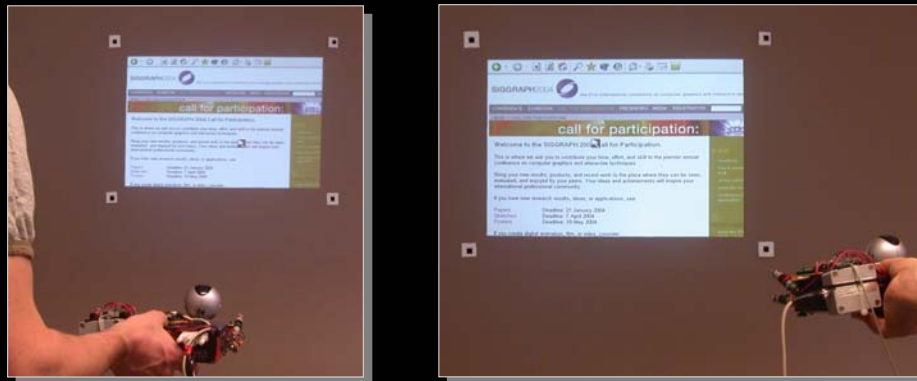


Image stays registered with world features

Absolute stabilization is possible when the camera is viewing four or more fixed surface points in general position, or the equivalent, on a planar display surface. Our goal is to find a homography H *between* the projector image plane and a fixed coordinate frame on the surface. Homography H *specifies the mapping between each pixel* on the projector image plane and the fixed surface coordinate frame. Hence we can use its inverse $\text{inv}(H)$ *to transform from a desired (fixed) projection on the surface to the projector image plane*, thereby determining the required projector image to give the fixed projection.

Image Stabilization



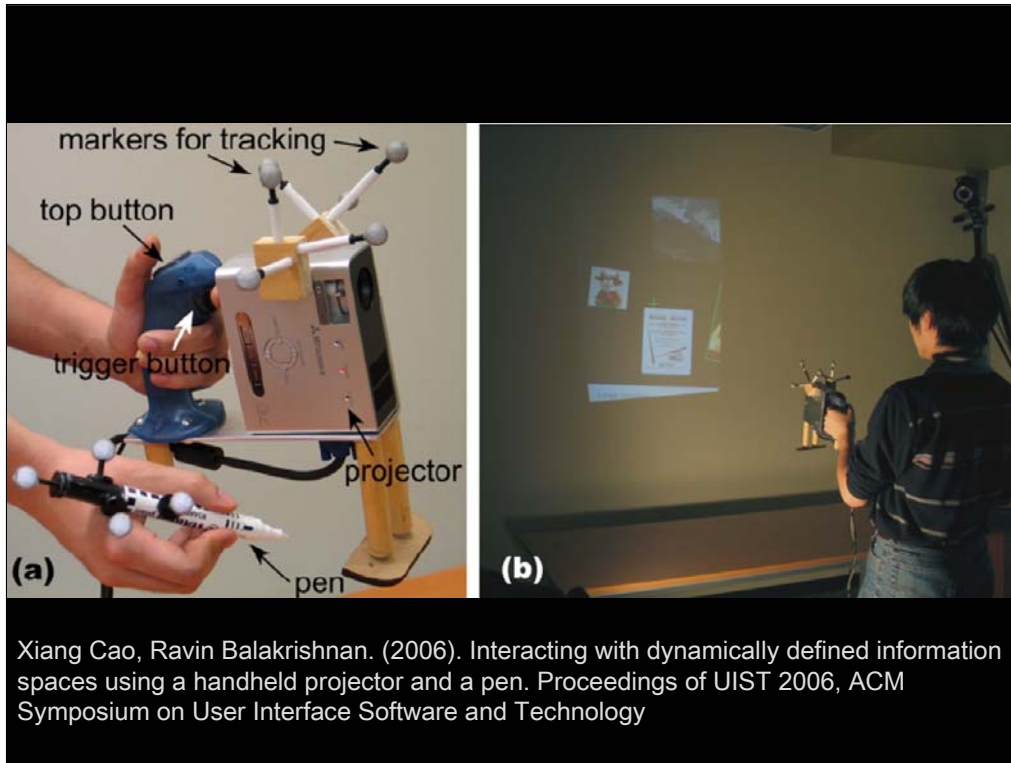
Video

A new way to play tic-tac-toe

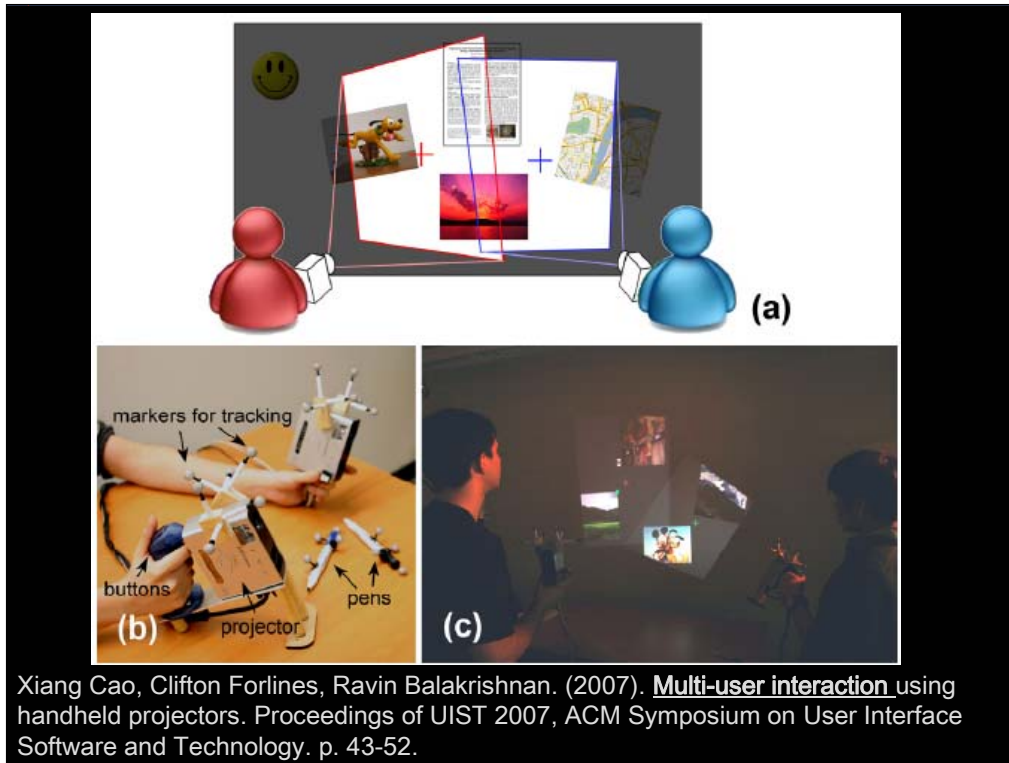


- Handheld projector.
- Grid projects to a fixed position on the wall.
- Cursor is guided by pointing the projector.

Playing game is a nice way to demonstrate various functionalities.




Handheld devices will soon have the ability to project information onto any surface, thus enabling interfaces that are not possible with current handhelds. The authors in Toronto explore the design space of dynamically defining and interacting with multiple virtual information spaces embedded in a physical environment using a handheld projector and a passive pen tracked in 3D. They develop techniques for defining and interacting with these spaces.



They extend the prior single-user research to co-located multi-user interaction using multiple handheld projectors. They present a set of interaction techniques for supporting co-located collaboration with multiple handheld projectors, and discuss application scenarios enabled by them. Handheld projectors provide interesting design challenges compared to other co-located collaborative settings such as a shared tabletop display. For example, users can create their individual displays with their projectors, allowing for easy support of personalized views, which is seldom the case in other settings.

Object Adaptive Projection

- Method
 - Passive elements
 - Identification of objects
 - Pose of projector
 - Gesture interaction




- Benefits
 - HMD: tracking issues
 - PDA: 'Last foot' problem

Let us look at object augmentation using a hand-held projector, including a technique for doing mouse-style interaction with the projected data. Common to some previous approaches, we do object recognition by means of fiducials attached to the object of interest. Our fiducials are 'piecodes', colored segmented circles, which allow thousands of distinct colorcodings.

As well as providing identity, these fiducials are used to compute camera pose (location and orientation) and hence projector pose since the system is fully calibrated. With projector pose known relative to a known object, content can be overlaid on the object as required.

AR Issues

- Preprocessing:
 - Authoring
- Runtime:
 - **Identification:** Recognition of objects
 - Using markers and visual tags
 - **Registration:** Finding relative pose of display device
 - Dynamic estimate of translation and rotation
 - Render/Warp images
 - **Interaction:**
 - Widgets, Gesture recognition, Visual feedback




The pie-codes help us solve two critical run-time components of an AR system.



Video

A hand-held projector can use various aspects of its context when projecting content onto a recognized object. We use proximity to the object to determine level-of-detail for the content. Other examples of context for content control would be gestural motion, history of use in a particular spot, or the presence of other devices for cooperative projection. The main uses of object augmentation are (a) information displays on objects, either passive display, or training applications in which instructions are displayed as part of a sequence; (b) physical indexing in which a user is guided through an environment or storage bins to a requested object; (c) indicating electronic data items which have been attached to the environment. Related work includes the Magic Lens [Bier et al. 1993], Digital Desk [Wellner 1993], computer augmented interaction with real-world environments [Rekimoto and Nagao 1995], and Hyper mask [Yotsukura et al. 2002].

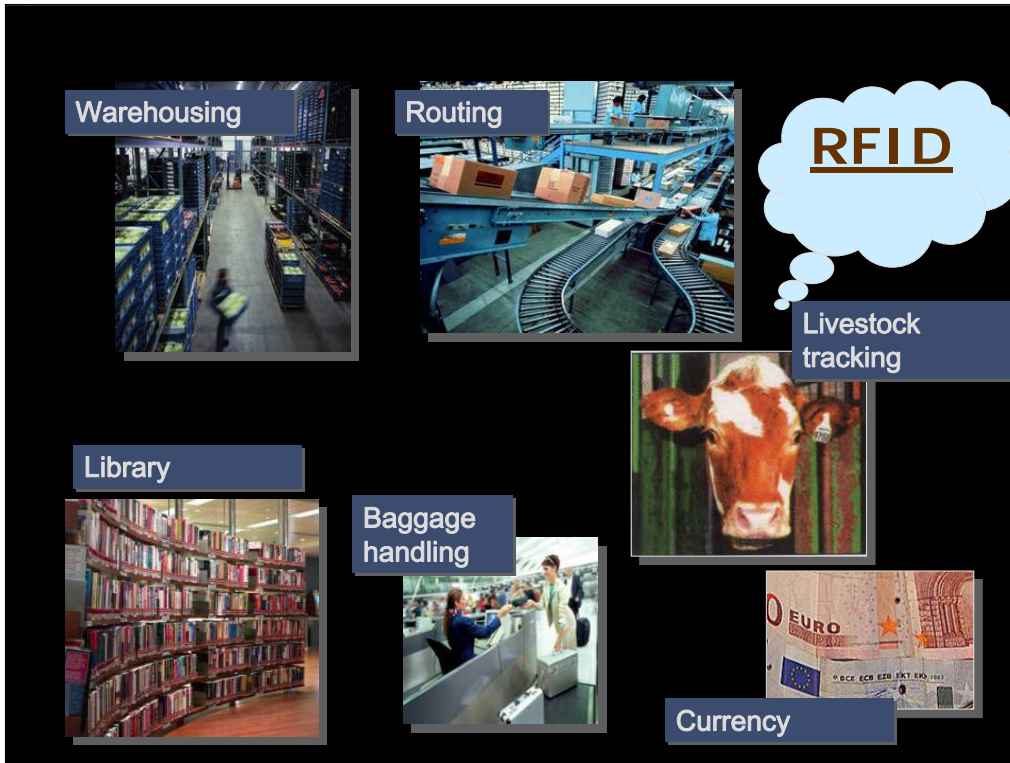
AR Issues



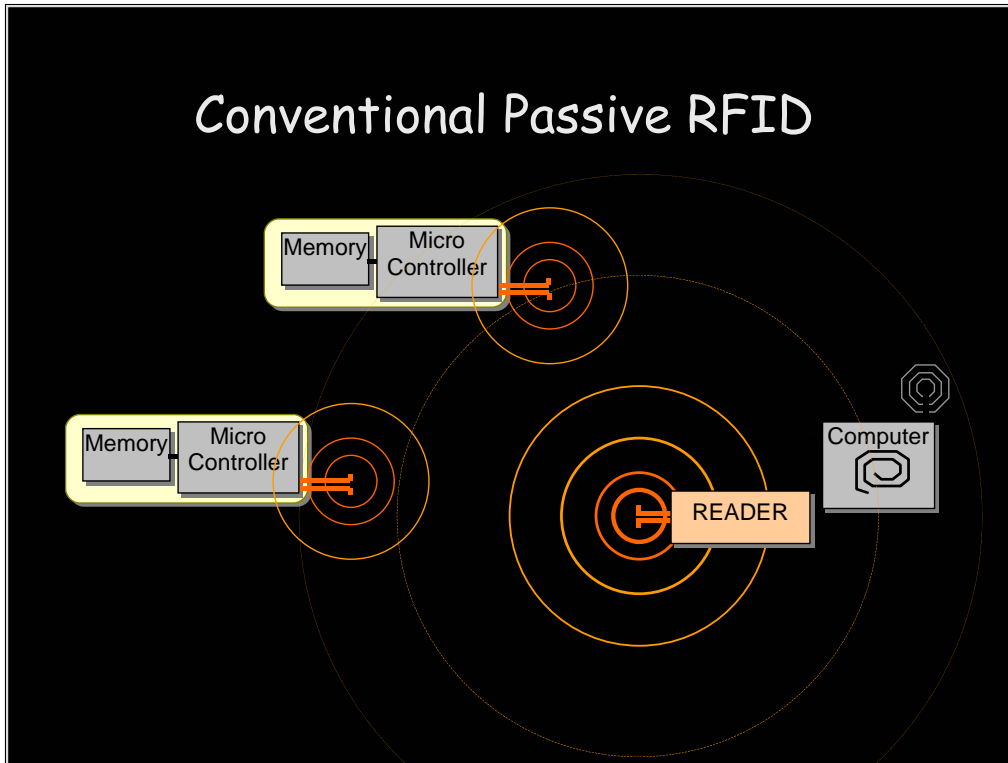
- Preprocessing:
 - Authoring
- Runtime:
 - **Identification**: Recognition of objects
 - Using markers and visual tags
 - **Registration**: Finding relative pose of display device
 - Dynamic estimate of translation and rotation
 - Render/Warp images
 - **Interaction**:
 - Widgets, Gesture recognition, Visual feedback

RFID ?

The pie-codes however can be replaced with radio frequency identification tags (RFID tags).



The first large-scale use of passive-RFID tags is expected to be for inventory control as part of logistics (a US\$900 billion industry), so we turn to a scenario in a warehouse - locating objects with a required property and annotating them.



Conventional tag communication works by broadcast from an RReader, with response from all in-range tags. Limiting the communication to a required tag is traditionally achieved using a shortrange tag-reader and close physical placement with the tag.

Powered radio-frequency tags currently use a battery that is about the size of a watch-battery, have a lifetime of a few years, and have a cost of a few dollars. In contrast, passive RFID tags are unpowered, can be as small as a grain of rice, and cost tens of cents [Want 2003]. Prices of both are dropping but the price differential will remain. The size and cost properties are such that RFID is showing signs of being adopted as a mass-deployment technology. Current commercial applications including embedding of RFID tags in packaging for inventory control, non-contact access control, and ear tags for livestock. Despite the interest in RFID, the available functionality is very limited – an RF-reader broadcasts a request, and in-range tags (collect energy from the RF pulse and) reply. The work is motivated by the observation that RFID is showing the potential to be a ground-breaking, pervasive technology, yet current functionality is limited.

Tagged Books in a Library

✓ Id : List of books in RF range



✗ No Precise Location Data

Are books in sorted order ?

RF beacons or tag transmit devices references but without the ability to point and without visual feedback.


A library scenario: Finding which books are on the shelf within the RF range is easy with a traditional (handheld) RF reader. But how can one find out which books are out of the alphabetically sorted order ? With passive photosensing RFID attached to each book we can find the exact location of each book. So, one can verify if there is a mismatch between the list of books sorted by title versus list of books sorted by position coordinates. The mismatch can also be indicated by projecting the arrows back on the shelf indicating the correct position. (Green arrows.) We can also find out if any book is placed upside down. We attach two tags, one at the top and one at the bottom. Books for which the location of the two tags is reversed is marked as upside down. This is indicated visually with red arrows.

AR Issues

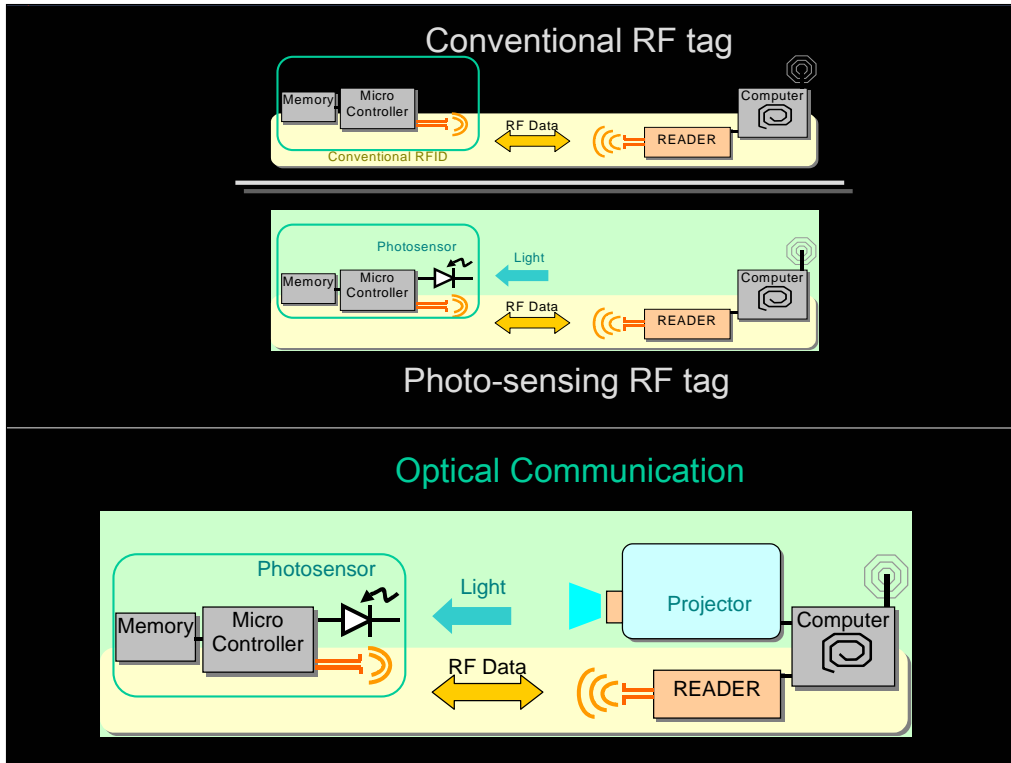
- Preprocessing:
 - Authoring
 - Runtime:
 - **Identification:** Recognition of objects
 - Using markers and visual tags
 - **Registration:** Finding relationships between the real and virtual worlds
 - Dynamic estimate of translation and rotation
 - Render/Warp images
 - **Interaction:** Supporting user interaction
 - Widgets, Gestures
-
- RFID
- Photosensing RFID
- Projector for visual feedback

Let us see how a combination of photosensing RFID tags and pocket projectors allows us to solve AR problems in a new imperceptible way.


RF Tag + Photosensor



We augment each tag with a photo-sensor to significantly extend the current functionality and support radio frequency identity and geometry (RFIG) discovery. The ability to address and wirelessly access distributed photosensors creates a unique opportunity. We recover geometric information, such as 3D location of tags or shape history of tagged objects, and exploit the associated geometric operations to bring the RF tags into the realm of computer vision and computer graphics.




The key issue in evolving our active tag system to passive tags would be power. We only allow computation and sensing consistent with the size and power levels we felt were achievable on a passive RFID system. For example, (a) tags are not photo-sensing or computing until woken up by the RF reader and (b) we do not have a light emitting diode (LED) on the tag as a visual beacon to a human or camera-based system because it would be power-hungry. Also note that the tags incorporate a photo-sensor, so a passive version could draw power not just from the RF channel, but also from the incident light. Of course, there would be significant engineering challenges in moving from active to passive RFID.



For example, a warehouse employee identifies food products that are close to expiry date and annotates an instruction to trash them. If these products were items on a computer desktop, this could be done with a few clicks. Our goal is to craft a scheme in the physical world that maintains the simplicity of the computer environment. In the computer, the items are files, and the interface is via keyboard, mouse, and display. In the physical world, the items are tagged objects, and the interface uses a handheld projector and user interaction directly through the projected information.

AR with Photosensing RFID and Handheld Projector

Video



For each tag

- From light sequence, decode x and y coordinate
- Transmit back to RF reader (Id, x, y)

Visual feedback of 2D position

- a. Receive via RF $\{(\text{Id}_1, x_1, y_1), (\text{Id}_2, x_2, y_2), \dots\}$
- b. Illuminate those positions



RFID

(Radio Frequency Identification)



RFIG

(Radio Frequency Id and **Geometry**)

Projector =
Optical Communication + Display + Interaction




Find tag location using
optical communication

Stabilized projection display

Select tags using
desktop like interaction


Machine AR

Robot 'Laser' Guidance Picking and Sorting Tagged Objects




Laser-guided robot. Guiding a robot to pick a certain object in a pile of objects on a moving conveyor belt, the projector locates the RFIG-tagged object, illuminating it with an easily identifiable temporal pattern. A camera attached to the robot arm locks onto this pattern, enabling the robot to home in on the object.

High speed projector
communication




Location dependent
Audio-stream projector



Nii, H., SUGIMOTO, M., and INAMI, M. 2005. Smart Light-Ultra High Speed Projector for Spatial Multiplexing Optical Transmission. PROCAMS 2005

The system built by Hideaki Nii and Masahiko Inami projects 3 different audio streams to different icons on the screen. They made a LED array board that was embedded in the camera and can project 3 different sounds simultaneously. If a subject changes a direction of the detector to a different icon, he/she can hear a different sound.

Fiber-calibration

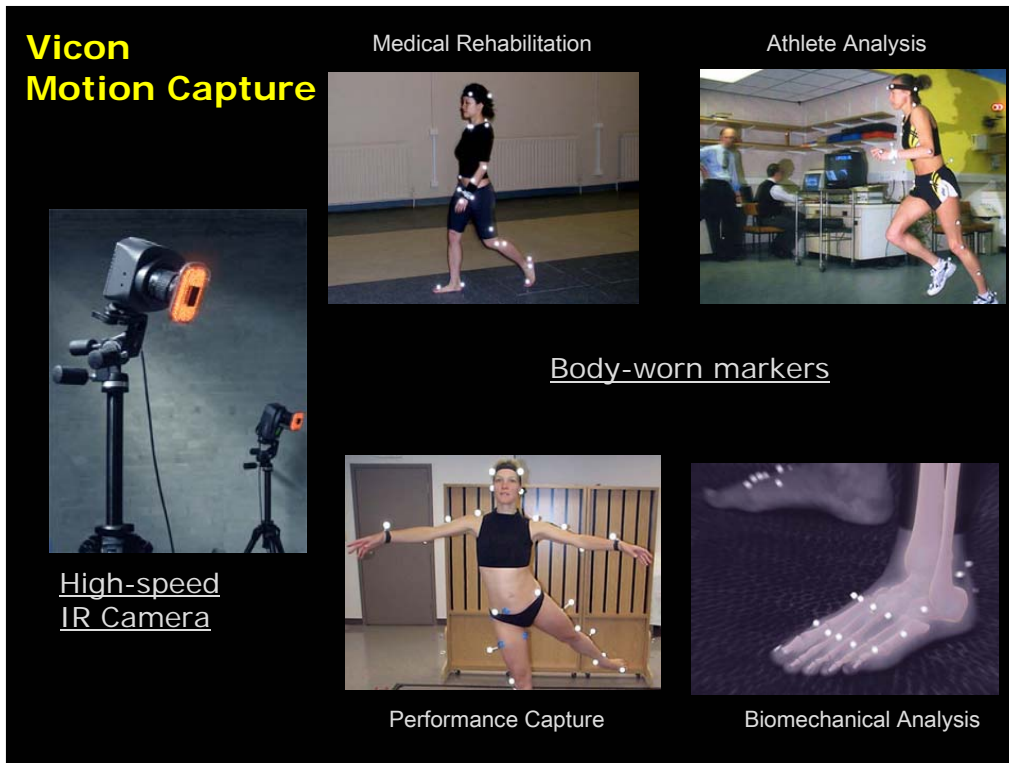


- Embed light sensors into the target surface
- optical fibers channel light energy from each corner to sensors
- USB connection to the PC
- White front surface hides fibers and acts as a light diffuser

Lee, J., Dietz, P., Raskar, R., Aminzade, D., and Hudson, S. "Automatic Projector Calibration using Embedded Light Sensors", UIST 2004.

The fundamental concept is to: 1) Embed optical sensors into the projection surface. 2) Project a series of Gray-coded binary patterns. 3) Decode the location of the sensors for use in a projected application. This video demonstrates this idea in the form of a target screen fitting application. It goes on to demonstrate how this approach can be used in multi-projector applications such as stitching (creating a large display using tiled projection) or layering (multiple versions of content on the same area for view dependent displays). Additionally, it can be used to automatically register the orientation of 3D surfaces for augmenting the appearance of physical objects.

This technique is also useful for performing automatic touch calibration of interactive whiteboards or touch-tables.



Consider how optical communication can be used to built motion capture systems.

For high speed tracking, the majority of optical motion capture systems use high speed cameras. These camera-based systems require special sensors and high bandwidth, are expensive and use a sanitized environment to maintain a high-contrast between the marker and its background. In this paper, we reverse the traditional approach. Instead of high speed cameras, we use high speed projectors to optically encode the space. Instead of retro-reflective or active light emitting diode (LED) markers, we use photosensitive tags to decode the optical signals.


'Motion' Capture ?

- Building a human model
 - Dense sampling over surface
 - Geometry with Id at every millisecond
 - Bio parameters
- Getting intimate
 - Cameras ..
 - Wearables
 - Second Skin (Sensor suit)
 - Tapping inside
- Close the loop in bio-I/O
 - Remote monitoring: Elderly care, training
 - Robot observation: learning, worker safety
 - Feedback for biomech/neuro interfaces




Unlike previous methods that employ high speed cameras or scanning lasers, we capture the scene appearance using the simplest possible optical devices – a light-emitting diode (LED) with a passive binary mask used as the transmitter and a photosensor used as the receiver. We strategically place a set of optical transmitters to spatio-temporally encode the volume of interest. Photosensors attached to scene points demultiplex the coded optical signals from multiple transmitters, allowing us to compute not only receiver location and orientation but also their incident illumination and the reflectance of the surfaces to which the photosensors are attached. We use our untethered tag system, called Prakash, to demonstrate the new methods.

Inside Modulated LED Projector




Multi-LED projectors are light transmitters for space-labeling. Each beamer is simply a LED with a passive binary film (mask) set in front. The light intensity sequencing provides a temporal modulation, and the mask provides a spatial modulation. We use a rigid array of such beamers, called projectors. The binary masks of individual beamers are carefully chosen to exploit the epipolar geometry of the complete beamer arrangement. Each beamer projects invisible (near infrared) binary patterns thousands of times per second.




Photosensing


tags determine their location by decoding the transmitted space-dependent labels.

2D Location



3D Location






For each tag

- From light sequence, decode x and y coordinate
- Transmit back to RF reader (Id , x , y)

Imperceptible Tags under clothing, tracked under ambient light



Let us look at the benefits of using multi-LED projectors for communication and photosensing markers.


Expensive high-speed cameras pose several scalability issues. Bandwidth limits resolution as well as frame-rate. Higher frame-rate (i.e., shorter exposure time) requires either brighter controlled scene lighting for passive markers or the use of power hungry active LED markers.

To robustly segment the markers from the background, these systems also use methods for increasing marker contrast. This usually involves requiring the actor to wear dark clothing under controlled lighting. The use of photosensing allows capture in natural settings.

Since the photosensors are barely discernible, they can be embedded in a wide range of natural clothing so long as the photosensing element is exposed. The power of emitters is comparable to the IR emission from TV remote controls. Instead of high-power emission, we exploit high-frequency modulation to robustly communicate with the photosensors. Similar to photosensors in TVs, our sensors will work in many lighting conditions. So, in studio settings, the actor may wear the final costume, and he/she can be shot under theatrical lighting.

Coded Illumination Motion Capture Clothing

- 500 Hz with Id for each Marker Tag
- Capture in Natural Environment
 - Visually imperceptible tags
 - Photosensing Tag can be hidden under clothes
 - Ambient lighting is ok
- Unlimited Number of Tags
 - Light sensitive fabric for dense sampling
- Non-imaging, complete privacy
- Base station and tags only a few 10's \$
- Full body scan + actions
 - Elderly, patients, athletes, performers
 - Breathing, small twists, multiple segments or people
 - Animation Analysis




Motion capture to date has been limited to exclusive and special purpose environments. Our low cost, easily portable, and visually imperceptible tracking system makes motion capture practical in a much wider application area. Interwoven tags on can help analyze patient rehabilitation progress after injuries. Tracking may be supportable in home video games and other casual virtual and augmented reality interfaces. The dream of filmmakers and game developers is 'on-set motion capture'. One of the recent examples is the motion and appearance capture in the movie 'Pirates of the Caribbean'. Our system can support this operation with unlimited number of imperceptible and interactive tags.

Recap



- All in one-solution
 - Display, AR, Mobility, Communication
- Mobility
 - Tiniest display devices
 - Desktop-like interaction
- Communication
 - Optical and RF tags
 - Temporal optical codes
 - Structured light illumination
 - Augmented Reality
 - Motion Capture



Mobile projectors are allowing new opportunities thanks to the emerging small form factor.

We can use projectors in a flexible way in everyday settings. The basic unit is a projector with sensors, computation, and networking capability. Singly or in a cluster, it can create a display that adapts to the surfaces or objects being projected on. As a hand-held, it allows projection of augmentation data onto a recognized object, plus mouse-style interaction with the projected data.

Projectors can also behave as “Smart Light,” which can provide not only images but also optical information.

The ideas provide geometric underpinnings for a new generation of projectors – autonomous devices, adaptive to their surroundings and able to optically communicate with smart devices.

**SIGGRAPH 2008 Course on
Projector-based Graphics**


**Projector-based Illumination for
3D Scene Modeling**

Hendrik P.A. Lensch
MPI Informatik

Overview

- Scene appearance as higher dimensional reflectance fields
- Capturing (and removing) global versus local illumination effects
- Pattern projection for 3D geometry acquisition

Computational Illumination for Scene Digitization

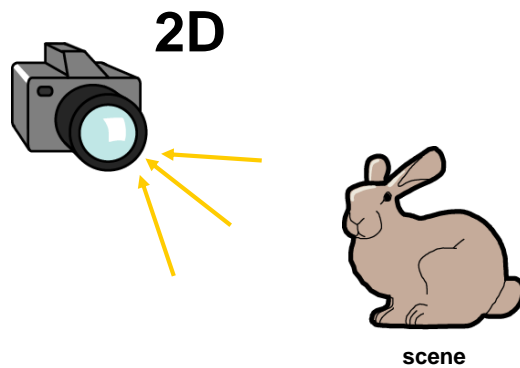


In this part of the talk we will discuss the design and the analysis of illumination patterns for scene digitization tasks, for acquiring 3D geometry to capturing the appearance of a scene. In the latter, we will discuss techniques for measuring the light transport on a ray-to-ray basis which further allows to distinguish local illumination effects, where the incident light is directly reflected at the scene surface, from global effects where light might be scattered or interreflected multiple times before arriving at the observer. We will first address the problem of appearance acquisition and then discuss the implications of a separation into direct and global components for 3D range scanning.

There are of course lots and lots of other fields where computational illumination plays a role, e.g. confocal microscopy,

Digitizing Real-World Objects

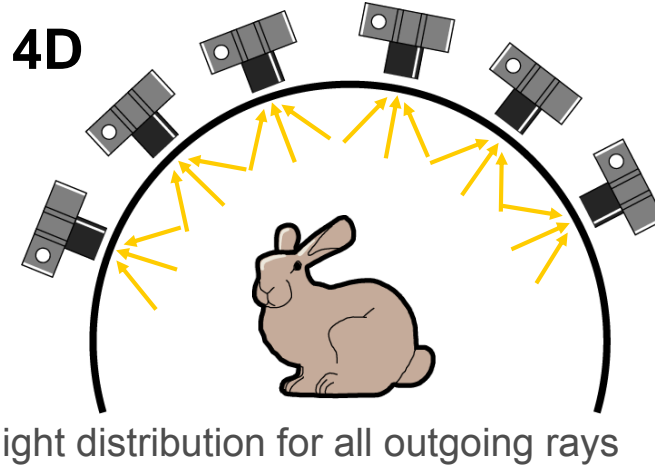
with a single picture



The simplest way to digitize a scene is to acquire a 2D photograph. Unfortunately, besides zooming in and out, the operations that can be performed on a single image are rather limited.

Light Fields

[Gortler96], [Levoy96]



Lumigraphs or Light Fields represent a collection of photographs from a set of different view points. In essence, a fully sampled light field captures the outgoing radiance for any ray leaving the surface.

Light Fields


[Gortler96], [Levoy96]



novel view synthesis but no relighting

With light fields applying view interpolation (maybe geometry assisted) arbitrary views of the object can be generated correctly as long as the virtual camera stays outside the visual hull. Direct and indirect reflections are reproduced in the rendered images, i.e. one can observe moving highlights, but the object will always be shown in the lighting that was present during acquisition.

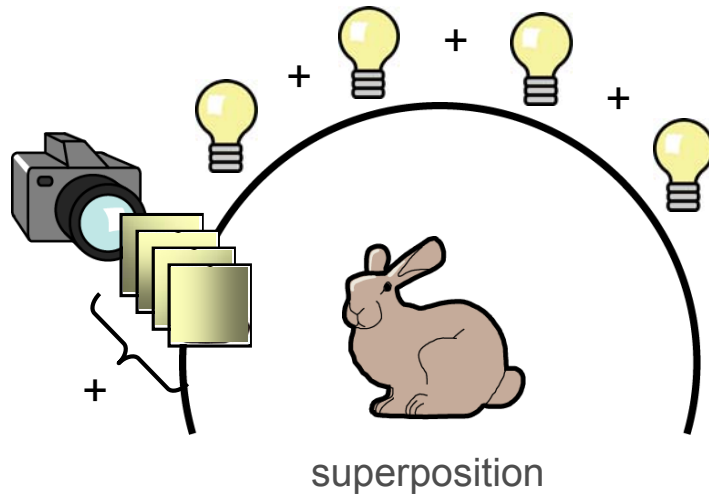
Relighting



In order to obtain renderings in a novel lighting condition, one needs to capture a data set of images in a variety of illumination conditions. The simplest is to illuminate the scene subsequently from a number of directions, e.g. from all positions on a hemisphere.

Relighting

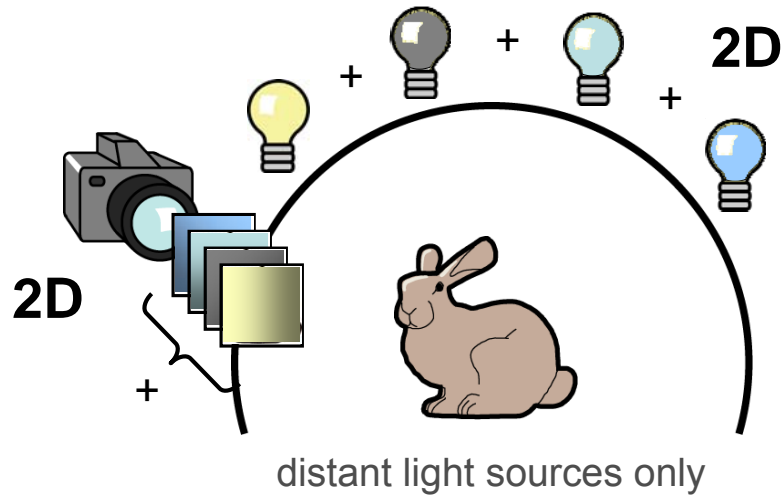
[Debevec2000]



Due to the superposition principle of light these captured image can simply be added up to produce an image of the scene as if illuminated from all light sources at the same time. The data structure is called a reflectance field.

4D Reflectance Fields

[Debevec2000]



mpii
max planck institut
informatik

Hendrik Lensch

It is even possible to assign arbitrary weight to the virtual light sources by simply multiplying the individual images with a different color before adding up. This way the appearance of the scene can be reproduced in arbitrary environments.

One restriction however is that the incident illumination is constrained to originate from the positions of the capturing light sources, for example from the sampling positions on a sphere. It is not possible to change the distance of the light source.

As all light sources are assumed to be directional light sources, reflectance fields captured this way cannot reproduce the appearance according to a spatially varying illumination pattern formed for example by a spot light or by projected shadows.

The direct reflection might be correctly reproduced, but the indirect reflections or subsurface scattering will be wrong.

Far and Near-Field Illumination




Far and Near-Field Illumination



This is demonstrated in these two slides. The projected light pattern alters the incident illumination for every point on the lamp shade. In the back on even sees how the incident light pattern is slightly blurred due to the scattering within the cloth. These effects cannot be reproduced with a far-field reflectance field.

6D Reflectance Fields

[Masselus2003]

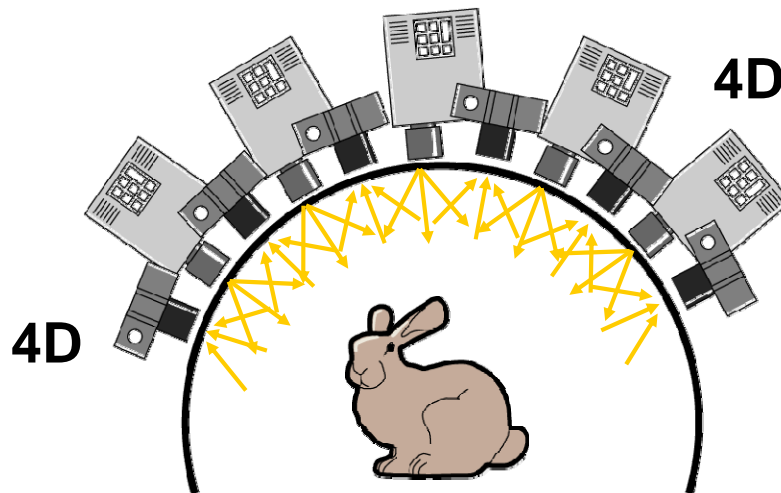


mpi
max planck institut
informatik

Hendrik Lensch

In order to cover global illumination effects correctly one needs to acquire a reflectance field for individual incident rays. Now, one image is captured for every possible incident light ray, instead of for each light position. The collection of incident rays form a 4D incident light field. Capturing one image per ray is of course expensive. Masselus et al. therefore only employed projectors with sixteen pixels. In the remainder of the course we will show how different illumination patterns can speed up this process.

8D Reflectance Fields



arbitrary view point + arbitrary illumination



Hendrik Lensch

The Holy Grail of appearance acquisition is to capture a full 8D reflectance field where the look of the object is captured from all possible direction from every ray in the incident light field. One can think of a reflectance field as an operator transforming the incident light field into the reflected light field.

Main Problem

- sampling an **8D function**
 - spending 100 samples/dimension
→ 10^{16} samples
 - hi-res 3D geometry: 10^8 vertices

- coherence in reflectance fields
→ reduced data complexity

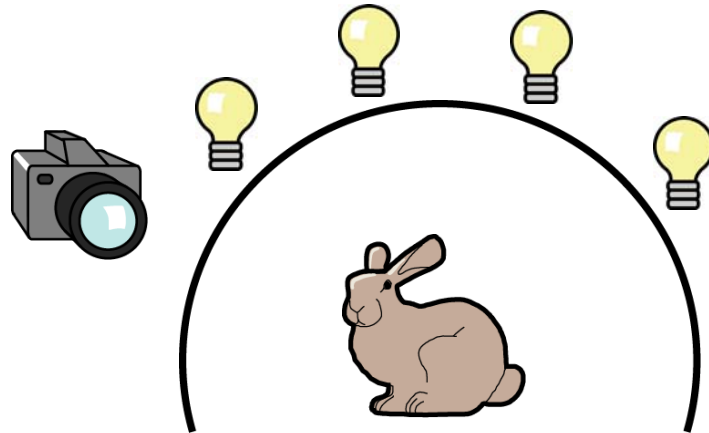
- no complete solution yet



Sampling, storing or rendering an 8D function is a tough problem. Approaches therefore often reduce the dimensionality of the problem, e.g., assuming a single view point, distant illumination, etc. In addition, using structured illumination, it is possible to capture multiple samples in the same shot.

Distant Illumination

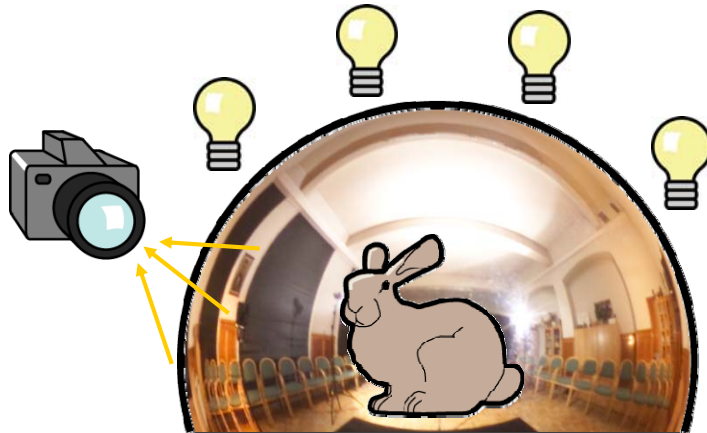
Reflectance Field



one image for each light direction

Let's start with a reflectance field for distant illumination.

4D Reflectance Field



As already said, one can reproduce the appearance of an object in the illumination of an environment map, assuming infinitely far away light sources.

Light Stage



- [Debevec2000]
- single view point
- assumes distant light sources
- video

The light stage is a device that is able to capture this. A dome of switchable light sources illuminates the actor in sequence. The video demonstrates how the collected images can then be combined to reproduce the appearance in arbitrary environments.

I would like to mention two properties of this particular setup:


-The light source images are captured one at a time, i.e. one image per light source. This corresponds to scanning through the space of light sources and is relatively slow. In this setup, it has been accelerated using a high-speed camera.

A positive point of this scanning illumination pattern is that the images can be directly used for rendering. In principle, no further analysis is necessary.

-The second issue is that the sphere of position is relatively sparsely sampled from a set of fixed positions. As with any sampling process, one might observe sampling artifacts due to undersampling

4D Reflectance Field

typical: a sparse sampling of light directions



Generating the appearance for a non-captured light source position therefore requires a modified acquisition system or additional processing in the form of light source interpolation.

Upsampling Reflectance Fields



230 input images

3547 images after upsampling

Too few directional samples lead to artifacts most noticeable at high frequency effects such as specular highlights or shadows. Employing a non-linear upsampling scheme it is possible to obtain a much clearer reflection of the environment in the sphere.

4D Reflectance Fields




no problem at glossy surfaces



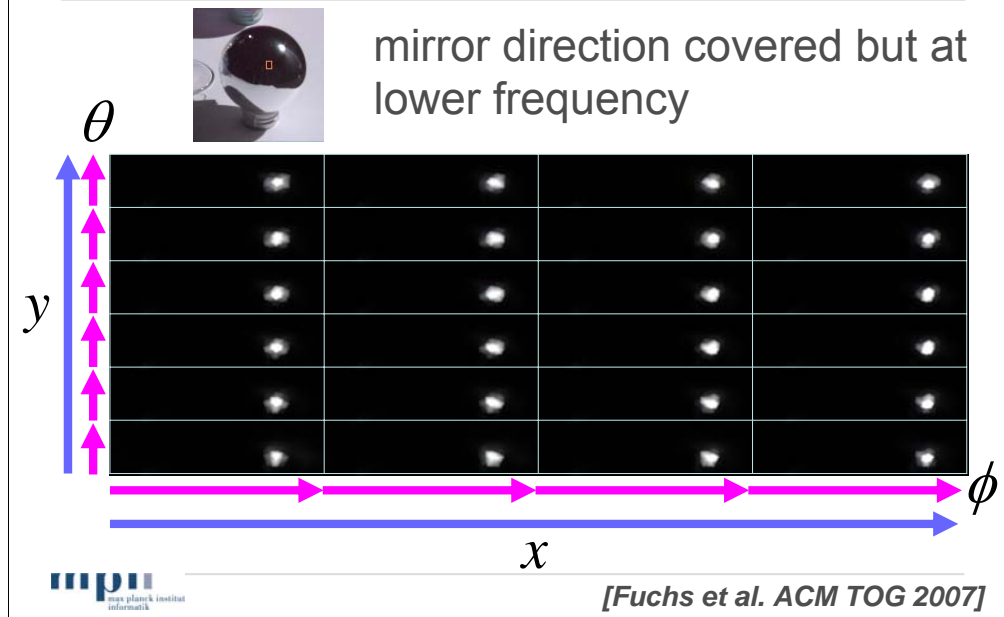
The problem gets clearer when taking a close up view on the scene: For a point on a glossy surface the reflectance function, the dependenc on the incident light ray direction (θ, ϕ) is rather smooth. A coarse sampling does not lead to any problems.

Undersampled Highlight Region



On a mirroring sphere however, the reflection function changes drastically with the incident or the viewing angle. In a neighborhood of pixels on a sphere, it might be that only for a very view pixels, the coarse sampling locations actually contain a mirror direction. A highlight generated by a moving light source would generate an incorrect intensity variation.

Prefiltering with Extended Light Sources



One can mitigate this problem by prefiltering the incident illumination, limiting the maximum frequency in the environment map. This has the effect the acquisition is done with extended light sources rather than point light sources. For each surface point a highlight will be captured.

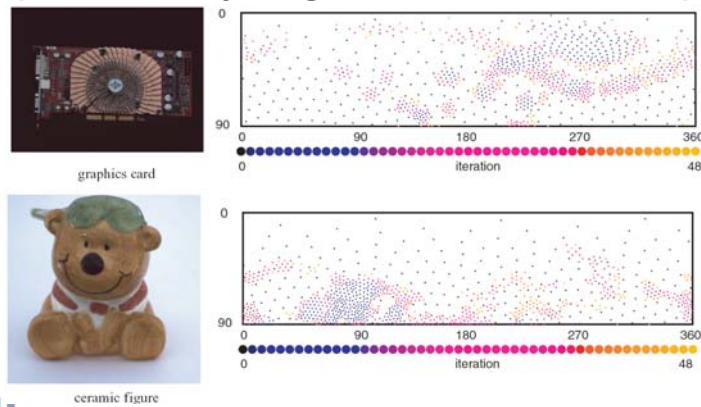
Prefiltering with Extended Light Sources



This can for example be achieved using an indirect light stage where disco spot lights project an extended spot onto the wall of a tent indirectly illuminating the scene in the middle.

Adaptive Acquisition

- choose sampling density and size depending on scene
- requires analyzing intermediate samples



mpi
max planck institut
informatik


[Fuchs et al. ACM TOG 2007]

This can for example be achieved using an indirect light stage where disco spot lights project an extended spot onto the wall of a tent indirectly illuminating the scene in the middle. With this setup, one can even run an adaptive sampling pattern where the size of the light sources and the sampling density is adapted to meet different sampling requirements for different section of the illuminating sphere.

This adaptive scheme provides a better sampling only at places where required. Compared to a sampling at full resolution a lot of sampling position can be well approximated from a sparser sampling. However, this adaptivity comes at the cost of a more complex acquisition setup, where the size and the location of the light sources need to be controllable. In addition, after each acquisition step the captured images need to be analyzed in order to predict, which samples to acquire next.

Subdivision Hierarchy

- take few input pictures
- analyze the reflectance field
- perform pair-wise upsampling, subdividing the mesh of light directions



Another way is to upsample the acquired data by performing non-linear interpolation between the originally captured samples.

Processing Pipeline

- separate in different effects:
 - highlights
 - shadows
 - low frequency effects
- separate upsampling
- refine result with texture priors



We apply different upsampling schemes for highlights, shadows and low frequency effects.

Result Videos

- continuous highlights
- smooth movement of reflections



linear



superresolution

The benefit of this approach is that from the same set of input images much for faithful renderings can be produced.

The environment is nicely reflected in the silver sphere. Shadows, highlights and even caustics move quite smoothly

Result Videos

- artificial glare
- some problems with grazing angles



linear



superresolution



[Fuchs et al. EG 2007]

Here the performance on shadows is demonstrated.


Note the correct effect of refraction in the glass. Some artifacts remain though.

The choice of sampling pattern for the incident illumination depends on the scene properties.

Pre-filtering, adaptive sampling or an advanced interpolation scheme can significantly improve over a sparse sampling.

Environment Matting

for dense sampling



one image for each monitor pixel



[Zongker et al. SIGGRAPH 99]

Environment mattes are designed to solve this task by densely sampling the light source positions. Especially geared to capture specular reflections or refractions. In principle, one could scan through all monitor pixels, but this would amount to 1 million images.

Environment Matting

- Extension of Alpha Matting capable of capturing transparent and specular objects for one view.
- Allows for reproduction with arbitrary backdrops.
- A high-resolution 4D reflectance field.



Traditional Alpha Matte

$$C = F + (1 - \alpha)B$$

- Composite color C
 - Foreground color F
 - Background color B
 - Pixel coverage α
-
- Acquired by blue/green screening
 - missing: dependence on light direction

Environment Matting - Definition

- Add reflected and refracted rays
(sum over backdrops)

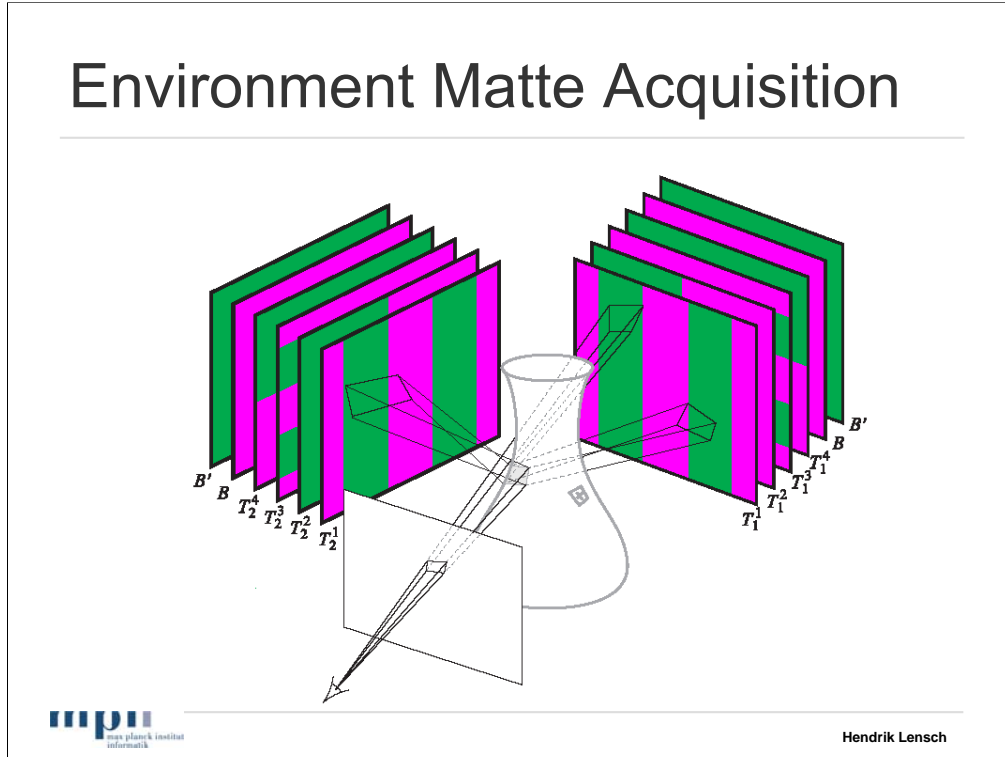
$$C = F + (1 - \alpha)B + \sum_{i=1}^m R_i M(T_i, A_i)$$

- Reflectance R
- Texture T
- Axis-aligned area A
- Averaging operator $M(T, A)$

Environment matting extends traditional blue screen matting to incorporate a directionally dependent part. In this model, it is expressed as the dependence on a rectangular patch of the environment $M(T, A)$.

In order to acquire an environment matte one needs to determine the size and the location of the region of the backdrop/sides that influences each camera pixel.

Environment Matte Acquisition




The process is visualized here. On each monitor side a set of hierarchical stripe patterns is displayed. From the recorded images for each individual displayed pattern the algorithm determines the direction of the sub-cone that contributes to the pixel's color.

In addition, the algorithm determines the spread of the cone.

Compared to scanning with individual pixels these hierarchical patterns are less precise as only an approximation of the actual beam is estimated. However, the acquisition time is significantly reduced.

Environment Matting Acquisition

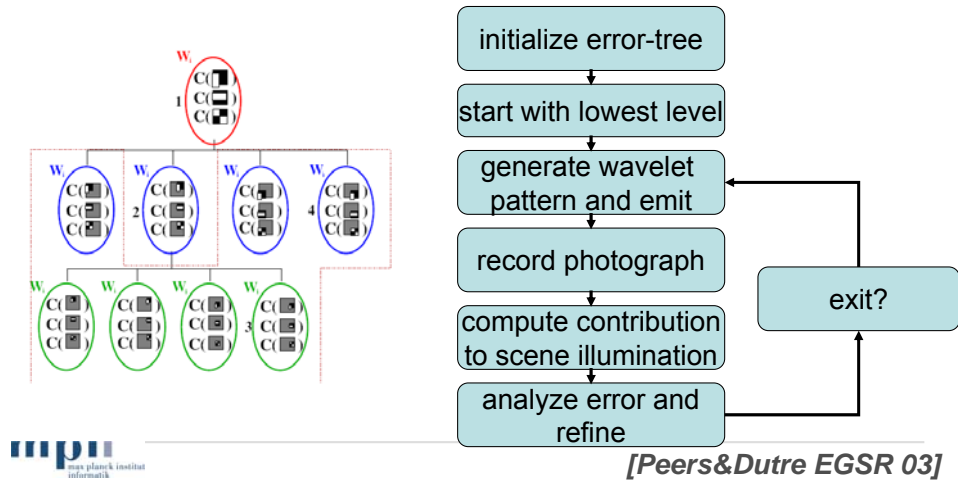


Environment Matte Extensions

- Gaussian Filter kernel [Chuang et al. 2000]
- Real-time acquisition [Chuang et al. 2000]
- Wavelets in acquisition [Peers et al. 2003]
- Multiple View Points (Opacity hulls)
[Matusik 2002]

Wavelet Environment Matting


- hierarchical wavelet basis $C = \sum_{i=1}^M a_i C_i$
- measure only where resolution is required



Peers and Dutre suggested an adaptive approach based on wavelets. The illumination domain is only refined if there are pixels whose reflection show high frequency responses in this region. For rather diffuse reflection, sub-division is not required and a lot of images can be saved. On the other hand, very sharp reflections can be correctly reproduced as well.

In principle, the output quality is as good as with a scanning approach, the number of images required is however drastically reduced. As can be seen in the flow diagram, the adaptive process requires to interleave pattern generation, acquisition and analysis.

Wavelet Noise



- randomly mix all wavelet bases for illumination (generate about 600 patterns)
- infer per-pixel reflection $T(x, y)$
- measurements $C_i(x, y) = T(x, y)W_i$
- minimize $E = \sum_i |C_i(x, y) - T(x, y)W_i|_2$
subject to a sparse $T(x, y)$
- (compare to compressed sampling)

In their wavelet noise approach Peers and Dutre no longer illuminate with individual wavelet basis but rather combine a random collection (with random weights) together in order to form a fixed set of illumination patterns.

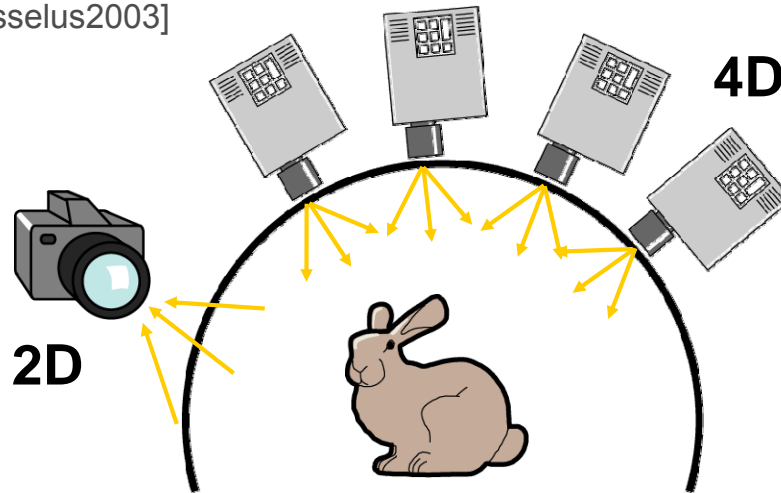
From this set of patterns the reflection functions now can only be inferred. They are no longer measured directly but need to be determined based on an optimization process. In this particular case the reflectance function of each pixel is hierarchically estimated minimizing the error between the current prediction and the measured samples while keeping the structure of the wavelet tree as simple as possible.

This process bears some resemblance with compressed sampling approaches.

The benefit of this technique is, that one obtains a hierarchical basis representation from a fixed set of patterns. No explicit sampling or adaptation is necessary. The acquisition is significantly simplified but the analysis at the end is rather complicated.

6D Reflectance Fields

[Masselus2003]




mpii
max planck institut
informatik

Hendrik Lensch

Now let's look at illumination pattern for a different type of reflectance fields where the illumination is recorded not for individual light source positions but rather for individual pairs of rays.

Pixel-to-Pixel Transport




Focusing on a single camera projector pair we need to determine for each projector pixel (p,q) what is the resulting camera image. The fourth-order tensor T stores the reflection coefficient for every pair of projector/camera pixel.

It fully describes the scenes reflection properties.


Pixel-to-Pixel Transport

- linear transport: $C = T L$
- measurement of the impulse response
- effort: one image per projector pixel = $O(N)$



This matrix can again be acquired by scanning, i.e. turning on each projector pixel individually, recording one image per projector pixel.

This is of course extremely expensive, but doable.



Here, one sees a corresponding measurement setup. Since a single projector pixel is not very bright compared to the projector's black level, we apply a laser projector with galvanic mirrors. The laser projector provides a significantly better contrast compared to video projectors.

The laser beam sweeps over the surface and a HDR video camera records every sample.



When the laser beam hits a translucent surface, the area surrounding the laser point lights up due to subsurface scattering. The footprint varies drastically with the incident location.

Video

1.000.000 images, 22 hours → model - 800MB



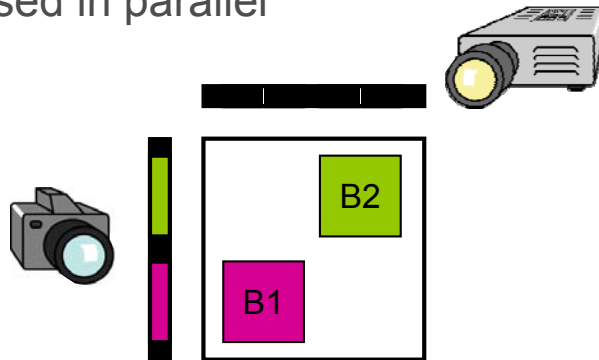
After one day and after capturing about one million images the entire matrix has been acquired. One can apply some compression for interactive rendering. Hole-filling is applied.

The data set now allows for relighting with arbitrary light patterns from arbitrary directions. The very specific properties of the light transport in translucent objects is correctly captured and reproduced.

One sees for example the drastic difference between front and back illumination as well as the light bleeding into shadowed regions.


Adaptive Parallel Acquisition

- assumption: sparse matrix, hierarchical basis
- radiometrically independent blocks can be sensed in parallel



Adaptive Parallel Acquisition

- assumption: sparse matrix, hierarchical basis
- radiometrically independent blocks can be sensed in parallel

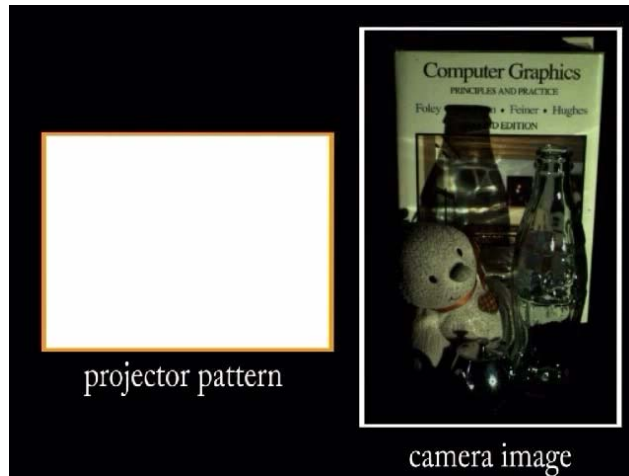


We now apply an adaptive scheme to this problem, again using a hierarchical wavelet bases, or equivalent hierarchically nested blocks.

In the case that the transport matrix is only sparsely populated, we can even parallelize the acquisition of multiple blocks. As long as they influence well separated camera regions only, they can be illuminated at the same time, and we can afterwards determine, based on the pixel location, which block in the reflectance tensor has been measured. The parallelization drastically reduces the acquisition time compared to a sequential adaptive scheme.

Adaptive Parallel Acquisition


parallelized acquisition of regions which do not overlap in the camera image



Here you can see the acquisition sequence for a quite complex scene. Initially the subdivided regions are sequentially acquired. At some point the algorithm detects that the corresponding footprints (areas of influence) in the camera image are well separated. All subsequent measurements of radiometrically independent blocks can further be captured in parallel. This process starts at some specific level of subdivision.

Adaptive Parallel Acquisition

parallelized acquisition of regions which do not overlap in the camera image $O(\log N)$



At the pixel-level, lots of samples can be acquired in parallel, especially in areas where direct reflections dominate the light transport. If the light transport is more complicated, as for example within the bottle, where refraction spreads the extent of the incident light rays, less samples can be captured at the same time.

Still the acceleration is dramatic. For most scenes, the acquisition time is in the order of $O(\log N)$ for N projector pixels


Relighting with Arbitrary Patterns

1.200 images, 2 hours → model - 220MB




The recorded reflectance field can now be relit with arbitrary illumination patterns. The tensor captured all light transport paths: one can see direct diffuse and specular reflections, interreflections. The glass bottle furthermore features refraction and caustics.

Global Light Transport



In the Cornell box we demonstrate even diffuse color bleeding and mirror reflections.

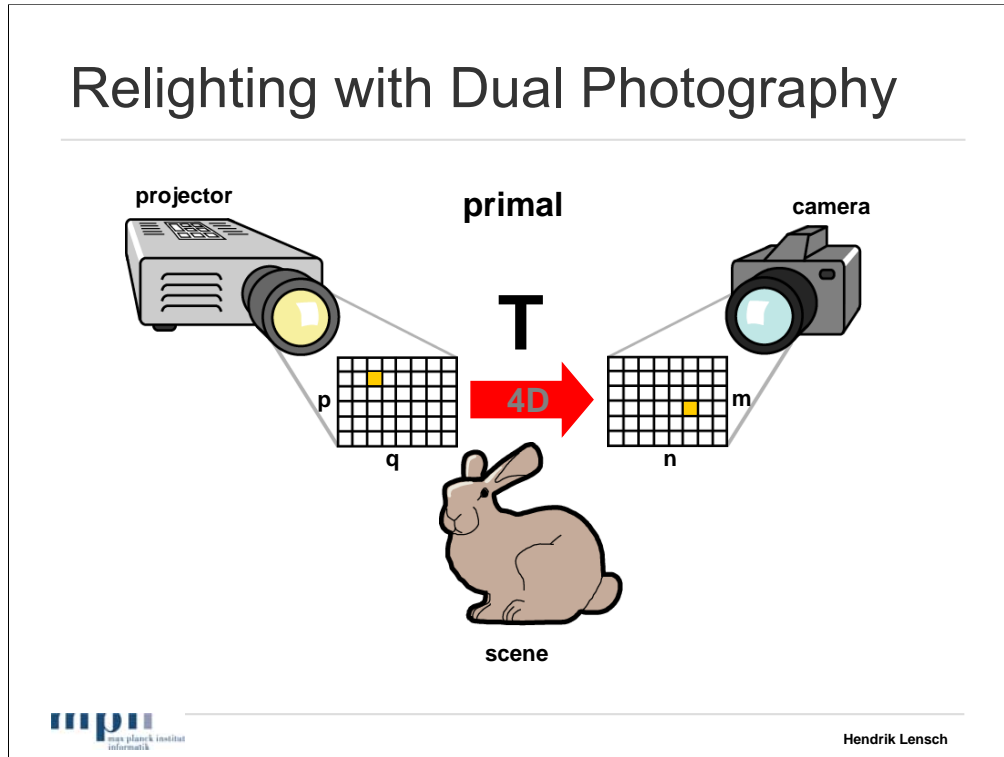
Helmholtz Reciprocity



Given the reflectance measured on a ray-to-ray basis, i.e. along one path, we can apply Helmholtz reciprocity.


It states that the reflectance along a path is the same no matter in which direction the path is followed. We can follow the light from the source to the receiver or we can swap their location and follow exactly the same path in the other direction, the observed reflectance coefficient will be exactly the same.

Relighting with Dual Photography



We can apply Helmholtz reciprocity has the tensor T captures the reflectance for any path between the projector and the camera. Helmholtz reciprocity holds for any collection of paths as long as the light source and the camera are within the same optical medium.


Relighting with Dual Photography




We can now swap the role of the camera and the projector, turning the original projector into a virtual camera, the original camera into a virtual projector. Given T one can easily compute an image from the point of view of the original projector as if illuminated from the camera. Here, one can even apply arbitrary patterns from the virtual camera. The process is surprisingly simple as the reflectance coefficients stay exactly the same. They just need to be reorder in order to account for the novel ray configuration. This can be done by simply transposing the captured tensor T .

Dual Photography

**photograph
from camera**



**dual image
from projector**



We can now generate different views from the same transport tensor, one from the original camera, illuminated from the original projector, and one from the projector illuminated from the camera.

All light transport paths are correctly incorporated, leading to a faithful reproduction of the bottles refractions and the caustics of the teapot.

Due to the duality, objects and shadows typically change place in the dual image.

Famous Examples



primal

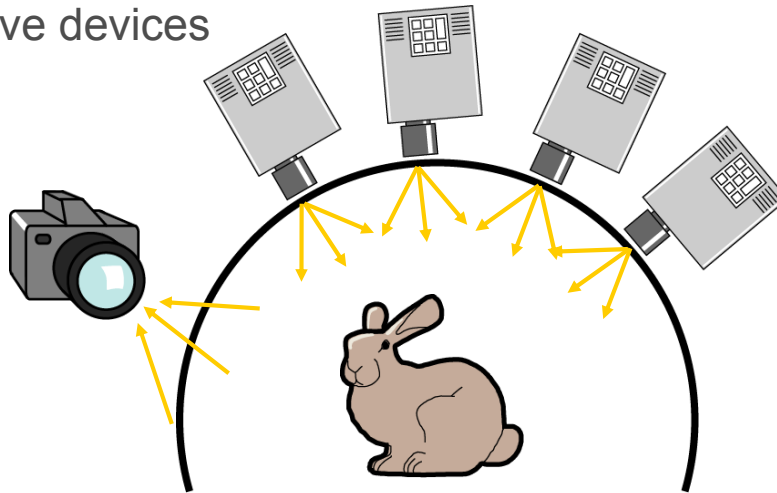


dual

Another set of rather famous objects.

Acquisition of 6D Reflectance Fields

active devices




Besides being a nice 'trick', Dual photography can be used to accelerate the acquisition of higher order reflectance fields.

In this configuration, a set of projectors produce an incident 4D light field. In order to measure the corresponding reflectance field, the reflectance field of each projector needs to be captured in sequence. Otherwise the illumination patterns of multiple projectors would overlap. It would be rather difficult to disentangle the contributions of the individual projectors.

Dual Acquisition Process

parallel acquisition by passive devices




With dual photography we can swap the role of cameras and projectors both during acquisition and during rendering. Since cameras are passive devices they can easily operate in parallel without interfering with each other.

This way, the acquisition process of a 6D reflectance field for relighting with 4D incident light fields can be accelerated to the time cost of acquiring the reflectance field between a single camera/projector pair.

Smooth Interpolation

100.000 images, 12 hours → model - 4.5GB




mpi
max planck institut
informatik

[Chen&Lensch VMV2005]

Thus it is possible, to acquire a rather dense reflectance fields on a ray-to-ray basis in significantly reduced time. It allows for relighting with arbitrary patterns from arbitrary directions.

Application: Virtual Photography

relighting arbitrarily complex scenes



novel illumination

original

acquisition pattern



[Garg et al. EGSR 2006]

Slightly extending the previous adaptive parallel acquisition approach the reflectance fields of arbitrarily complex scenes can be acquired at moderate time cost. Here you see the complex light pattern formed in a glass of gummy bears illuminated from two different directions with a high frequency pattern.

One can acquire reflectance field that reproduce correctly all local and global illumination effects.

Local vs. Global Reflections



Looking at the structure of the resulting light transport tensors the difference between global and local illumination effects are clearly revealed.

In this case, we just concentrate on the slice which is formed by illuminating every point along the white line and looking at how much light is reflected from all points along the line.

Local vs. Global Reflections




In the left case, the reflections are due to the direct reflection off a planar, textured sheet of paper.

As one can see, light will only be reflected at the place where the sheet is illuminated, leading to the (band) diagonal structure of the matrix. The tensor does not contain any contribution in off-diagonal elements.

In the right case, the light transport between the two pages of an open book are captured.

The transport is still dominated by the direct reflection, i.e. the entries along the diagonal. However, there is additional indirect light transport from the left to the right side of the book. as well as tertiary reflections. Those populate the off-diagonal elements in the matrix.

Application: Getting Rid of Global Effects



- remove off-diagonal components
- diagonal entries might still contain global components.

With this knowledge at hand, one can now try to separate direct and global illumination effects.

Simply removing off-diagonal elements from the transport tensor removes most of the indirect reflections. However, even the will to some extent be influenced by global illumination effects, e.g. 3rd order reflections back to the point of incidence.

Fast Separation of Direct and Global Effects

- main idea: global illumination effects dampen high frequencies
 - illuminate with shifted high frequency patterns
 - only the local illumination will change
 - global illumination will be invariant to phase shifts



A principled method for a fast separation of direct from global illumination effects in images captured from a single view point has been proposed by Nayar et al.


In this example, the direct component contains all specular highlights. The global component is strongest for the translucent surfaces and for the interreflections between the walls.

For this problem again, structured illumination is applied.

The idea is that only the direct reflection component will respond to high frequency variation in the incident illumination while the global component will remain fix.

Fast Separation of Direct and Global Effects

- shifted periodic patterns (e.g. checker board)
- record per-pixel minimum and maximum
- approximation: $L_g = L_{\min}$ $L_d = L_{\max} - L_{\min}$



This effect can be observed when illuminating the scene with a set of shifted periodic high frequency patterns.


In the recorded image sequence one just need to determine the minimum and maximum per pixel in order to quickly compute the direct and global component.

In the paper a more precise approximation based on the projector contrast and black-level can be found.

This tool of quickly removing any global illumination effect from measurements performed with structured light can help to make computer vision tasks more robust.

Traditional 3D Scanning

- illumination with a swept point/line $O(N)$
- detection of the brightest dot on surface
- triangulation based on laser and camera coordinates



For example, 3D range scanning, where structured illumination has been initially been applied.

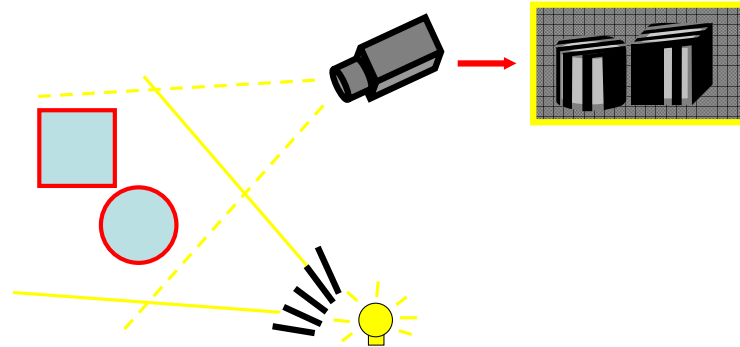
Here, we see the principle of a traditional laser range scanner. Sweeping a line or a dot over the scene the camera records the brightest point along the scan line.

This detection can be performed with sub-pixel precision.

Given the location of the brightest point in the camera image and the rotation angle of the laser, one can now easily determine the distance between the camera and the surface point by triangulation.

Structured Illumination

- reduced number of acquired images $O(\log N)$
- analyze the captured images to determine projector coordinates



Instead of sweeping a line, structured hierarchical codes can reduce the acquisition effort to $O(\log N)$. The recorded camera images now need to be analyzed in order to determine the projector row that illuminated the visible surface point in each camera pixel.

Binary Encoded Stripes

- determine the corresponding column in $\log(n)$ time.

wwbbbb**bw**



0 2 4 6 8 10 12 14
1 3 5 7 9 11 13 15



Hendrik Lensch

observations:

wwbb ~ col. 3

wbbw ~ col. 6

bbbw ~ col. 14

The easiest code simply encodes the column number in a set of bit planes, producing one illumination pattern for the most significant bit and so forth.

By interpreting the recorded on/off patterns as a binary number one directly can read of the projector column again for each pixel.


One limitation of this approach is however, that the encoding is constant within each bar of the highest resolution.

Stripe Patterns - Extensions

- Use Gray-code for more reliable detection
- Different pattern for interactive scanning [Rusinkiewicz '02]
- Color patterns for faster, texture invariant acquisition

Various extensions have been proposed to make the encoding more robust or more efficient.

Phase Shifting



- precision of binary encoding limited by highest resolution
- use shifted sine pattern to obtain subpixel precision
- requires precise gray levels
- period needs to be determined using coarser levels (phase unwrapping)

One way to overcome the problem of the discrete locations it to apply sinusoidal functions shifted in phase for the last few iterations. Given some photometric calibration, the exact phase of a surface point in these shifted pattern can be reconstructed with subpixel precision. The output is the phase within one period of the highest frequency. The binary codes of coarser resolution are necessary to locate the period. Here, the constant location is sufficient.

Robust 3D Range Scanning


- use patterns that inherently filter out global illumination effects
 - high frequency sinusoids do this
- use additional polarization – multiply scattered light is typically unpolarized

With the shifted sinusoidal patterns we can do two things at the same time:

- 3D range scanning using phase shifting
- separation into direct and global reflections

This produces a 3D range scanning approach that is robust to global illumination effects.

Robust Sine Patterns



- illuminate with phase shifted sinusoids: e.g.

$$I = \sin(x + \delta_i), \quad \delta_i \in \{-2\pi/3, 0, 2\pi/3\}$$

- estimate phase -> 3D position:

$$\Phi = \tan^{-1} \left(\frac{-\sum L_i \sin \delta_i}{\sum L_i \cos \delta_i} \right)$$

- or perform separation

$$L_g = \frac{2}{3}(L_0 + L_1 + L_2) - L_d, \quad L_d = \frac{2}{3} \sqrt{3(L_0 - L_2)^2 + (2L_1 - L_0 - L_2)^2}$$

The equations for obtaining the phase and for performing the separation are rather simple.

It is worth noting that estimating the direct component using sine pattern using the above equations is much more robust than detecting the minimum and maximum from a sequence of shifted binary patterns.

Using sinusoids all measurements are incorporated.

Another way to do descattering is using polarization.

There exists a vast literature on polarization for separation of reflection components.

Rather than the algorithmic approaches presented early that use high frequency patterns


polarization filtering makes use of the physical properties of scattering events. Each scattering event or reflection changes the polarization of the reflected light. Multiple scattering and subsurface scattering decorelate the resulting polarization state from the incoming orientation while direct reflections typically keep them.

These images show the comparison in 3D reconstruction.

The underlying structures, here the straws, destroy the reconstruction if pure phase-shifting is used. Its descattering capabilities are not strong enough in this case.

Using polarization difference imaging on top of the shifted sine patterns we can obtain a clear 3D scan of the first surface.

Application: 3D Scanning



mpi
max planck institut
informatik

[Chen et al. CVPR 2007]



A commercial laser range scanner has severe difficulties in scanning translucent objects. The subsurface scattering offsets the location of the peak, resulting in a corrupted 3D scan.

Removing the global effects a significantly better 3D scan can be obtained.

Modulated Patterns in 2D

- use sinusoids in both dimensions
 - results in descattering in both dimensions

$$I = \sin(x + \delta_i) \cdot \sin(y + \delta_j)$$



[Chen et al. CVPR 2008]

Modulating a horizontal sine pattern with a vertical one and capturing $M \times N$ phase shifts for both directions, the separation capabilities can be drastically increased. Modulation and shifting in two dimensions outperforms even polarization difference imaging.

Illumination Patterns and Codes

- scan
- periodic
- hierarchical
- multiplexed / Noise

- modified by the properties of light
 - wavelength
 - polarization
 - time of flight / phase (not discussed here)

In the different fields we have seen the same codes and illumination patterns over and over again.

Scanning

- moving a pixel or a line
- simple
- slow
- low light efficiency, consider black frame



$O(N^2)$




$O(N)$


Scanning is the slowest but delivers quite robust measurements. The benefit is that the acquired images can be used directly without any further processing.

Periodic Patterns


- $O(1)$ – constant number of phase shifts
- requires some analysis
- information within one period only
- some robustness wrt. global illumination
[Nayar et al. 2006, Chen et al. 2007/2008]



stripes



2D grid



sinusoids


Periodic patterns can be acquired much faster. As the number of phase shifts is fixed the acquisition effort is $O(1)$. Of course, periodic patterns are not always applicable, e.g. not for measuring reflection properties.

They require some analysis which is typically rather simple.

The second benefit of periodic patterns is that they are inherently robust against global illumination effects if analyzed appropriately.

Hierarchical Bases


- $O(\log N)$
- deliver global localization
- stripes or sinusoids (phase unwrapping), spherical harmonics



Hierarchical bases are used to compactly localize features globally. Furthermore, when acquiring reflectance fields or environment mattes they can be used to directly sample into a hierarchical representation of the otherwise large data structures.

Adaptive Patterns


- patterns adapted to scene content / signal
- based on a hierarchical basis
 - determine if subdivision is necessary
- optimal set of patterns
- requires analysis after each acquisition step




Combined with an adaptive acquisition scheme, hierarchical bases can be used to accelerate the acquisition. Sub-trees that do not require further subdivision will be culled already during acquisition, unnecessary acquisition steps are saved.

Multiplexed Illumination

- random distribution of basis elements
- same (or less) effort than scanning
- each sample is obtained by combining multiple measurements – increased SNR
- requires analysis
 - multiplication with inverse measurement matrix
 - optimization (compressed sampling)



Hadamard



wavelet noise

Multiplexing multiple basis functions, it is possible to either improve the SNR slightly or to accelerate the acquisition sacrificing some quality.

Multiplexed illumination requires some additional analysis to obtain the measurements for individual bases.

Summary

- computational illumination used for measuring
 - 3D shape
 - object appearance
- choice of ***codes***, ***acquisition*** and ***analysis*** are always interlinked

Challenges

- robust acquisition under strong ambient illumination
- dynamic scenes
- large scenes
- uncontrolled environments

SIGGRAPH 2008 Course on
Projector-based Graphics
Projector-based Illumination for 3D Scene Modeling
Hendrik P. A. Lensch

References

- [1] Billy Chen and Hendrik P. A. Lensch. Light source interpolation for sparsely sampled reflectance fields. In Günther Greiner, Joachim Hornegger, Heinrich Niemann, and Marc Stamminger, editors, *Vision, Modeling, and Visualization 2005 (VMV'05)*, pages 461–469, Erlangen, Germany, November 2005. Aka.
- [2] Tongbo Chen, Hendrik P. A. Lensch, Christian Fuchs, and Hans-Peter Seidel. Polarization and phase-shifting for 3D scanning of translucent objects. In *Proceedings of CVPR*, 2007.
- [3] Tongbo Chen, Hans-Peter Seidel, and Hendrik P. A. Lensch. Modulated phase-shifting for 3D scanning. In *Proceedings of CVPR*, 2008.
- [4] Yung-Yu Chuang, Douglas E. Zongker, Joel Hindorff, Brian Curless, David H. Salesin, and Richard Szeliski. Environment matting extensions: Towards higher accuracy and real-time capture. In *Proceedings of ACM SIGGRAPH 2000*, pages 121–130, 2000.
- [5] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *Proc. SIGGRAPH*, pages 145–156, July 2000. ISBN 1-58113-208-5.
- [6] Martin Fuchs, Volker Blanz, Hendrik P.A. Lensch, and Hans-Peter Seidel. Adaptive sampling of reflectance fields. *ACM Transactions on Graphics*, 26(2):10:1–10:18, June 2007.
- [7] Martin Fuchs, Hendrik P. A. Lensch, Volker Blanz, and Hans-Peter Seidel. Superresolution reflectance fields: Synthesizing images for intermediate light directions. *Computer Graphics Forum*, 26(3):447–456, September 2007.

- [8] Gaurav Garg, Eino-Ville Talvala, Marc Levoy, and Hendrik P. A. Lensch. Symmetric photography: Exploiting data-sparseness in reflectance fields. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, pages 251–262, June 2006.
- [9] Michael Goesele, Hendrik P. A. Lensch, Jochen Lang, Christian Fuchs, and Hans-Peter Seidel. DISCO – Acquisition of Translucent Objects. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2004)*, 23(3), 2004.
- [10] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The lumigraph. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 43–54, August 1996.
- [11] Marc Levoy and Patrick M. Hanrahan. Light field rendering. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, pages 31–42, August 1996.
- [12] Vincent Masselus, Pieter Peers, Philip Dutré, and Yves D. Willems. Relighting with 4d incident light fields. *ACM Transactions on Graphics*, 22(3):613–620, July 2003.
- [13] Wojciech Matusik, Hanspeter Pfister, Addy Ngan, Paul Beardsley, Remo Ziegler, and Leonard McMillan. Image-based 3d photography using opacity hulls. *ACM Transactions on Graphics*, 21(3):427–437, July 2002.
- [14] Shree K. Nayar, Gurunandan Krishnan, Michael D. Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Transactions on Graphics*, 25(3):935–944, July 2006.
- [15] Pieter Peers and Philip Dutré. Wavelet environment matting. In *Proceedings of the 14th Eurographics Symposium on Rendering*, pages 157–166, 2003.
- [16] Pieter Peers and Philip Dutré. Inferring reflectance functions from wavelet noise. In *Rendering Techniques 2005: 16th Eurographics Workshop on Rendering*, pages 173–182, June 2005.
- [17] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In Eugene Fiume, editor, *Proceedings of SIGGRAPH 2001*, Computer Graphics Proceedings, Annual Conference Series, pages

117–128. ACM Press / ACM SIGGRAPH, August 2001. ISBN 1-58113-292-1.

- [18] Nenanel Ratner and Yoav Y. Schechner. Illumination multiplexing within fundamental limits. In *Proceedings of CVPR*, 2007.
- [19] Szymon Rusinkiewicz, Olaf Hall-Holt, and Marc Levoy. Real-time 3d model acquisition. *ACM Transactions on Graphics*, 21(3):438–446, July 2002.
- [20] Yoav Y. Schechner, Shree K. Nayar, and Peter N. Belhumeur. A theory of multiplexed illumination. In *Proceedings of IEEE ICCV*, volume 2, pages 808–815, 2003.
- [21] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R. Marschner, Mark Horowitz, Marc Levoy, and Hendrik P. A. Lensch. Dual photography. *ACM Transactions on Graphics*, 24(3):745–755, August 2005.
- [22] Douglas E. Zongker, Dawn M. Werner, Brian Curless, and David H. Salesin. Environment matting and compositing. In *Proceedings of ACM SIGGRAPH 1999*, pages 205–214, 1999.