Imaging and Display Applications using Fast Light. SPIE Proceedings Photonics, Vol. 7210, 2009.

Imaging and Display Applications using Fast Light

Ian McDowall*[a], Mark Bolas*[b]

[a]Fakespace Labs, 241 Polaris Ave, Mountain View, CA 94043;
[b]Dept. of Interactive Media, USC School of Cinematic Arts,

ABSTRACT

The unique qualities of the TI DLP devices have enabled a number of interesting applications. The DLP is essentially a fast binary light modulator and using the power of modern graphics processors these devices can be driven with images computed on the fly at rates of several thousand frames per second. A number of these applications have been developed by the University of Southern California where fast light is exploited to create a light field display. In another application, fast light is coupled with a synchronized high speed camera to extract the 3D shape of an object in real time.

Imaging and Display Applications using Fast Light

Ian McDowall*[a], Mark Bolas*[b]
[a]Fakespace Labs, 241 Polaris Ave, Mountain View, CA 94043;
[b]Dept. of Interactive Media, USC School of Cinematic Arts,
University Park, LUC-310B, Los Angeles, CA 90089.

ABSTRACT

The unique qualities of the TI DLP devices have enabled a number of interesting applications. The DLP is essentially a fast binary light modulator and using the power of modern graphics processors these devices can be driven with images computed on the fly at rates of several thousand frames per second. A number of these applications have been developed by the University of Southern California where fast light is exploited to create a light field display. In another application, fast light is coupled with a synchronized high speed camera to extract the 3D shape of an object in real time.

Keywords: DLP, Mule, 3D, Fast Light, Projector, Pattern Projector, Light Field

1. INTRODUCTION

The unique qualities of the TI DLP devices have enabled a number of interesting applications. The DLP is essentially a fast binary light modulator and using the power of modern graphics processors these devices can be driven with images computed on the fly at rates of several thousand frames per second.

1.1 Creating Fast Light – The Mule

The Mule is the combination of a TI DLP chip Discovery 1100 or 3000 card along with an interface created by Fakespace Labs which creates a DVI interface to the Discovery card. The DVI interface looks to a PC like a standard display and from the PC's point of view behaves like a monitor. The Mule electronics take in the DVI video frames which are composed of 24 bit pixels – 8 bits for red, 8 for green, and 8 for blue and instead of using these bits to create a color image the Mule creates a sequence of binary images in time. Each incoming image therefore encodes 24 binary images which on a traditional display are used to adjust the relative brightness of red, green, and blue pixels. The Mule however treats each of these 24 binary images equally.

Each binary image of the incoming frame is transferred to the DLP chip as an entire frame and for $1/24^{th}$ of the incoming frame time. Various control signals are also created to enable synchronization of the Mule with other devices. Each image from the DVI contains 24 binary images so the output frame rate of the Mule is 24 times that of the incoming video rate. If the incoming video is at 60Hz, the binary frame rate will be 24*60 or 1,440Hz at a vertical refresh rate of 100Hz the resulting binary image sequence from the Mule will be 2,400 frames per second. Rates of several thousand frames per second are possible with this approach.

The Mule thus interfaces to a standard DVI port. Creating the images can be accomplished in Open GL or DirectX. This may be done by taking advantage of the graphics card's ability to draw an image in one bit plane of the frame buffer at a time. The software application draws images to each of the 24 bit planes of the frame buffer in sequence and independently. The resulting frame output the composite image over DVI. This approach makes the programming and rendering of the patterns easy to change over time – they do not have to be pre-computed and can be changed with minimal programming effort and the results of the changes seen in real time.

2. PATTERNS FOR SENSORS

The binary nature of the DLP and the ability to create binary patterns quickly lend the DLP technology to creating real time patterns for various sensors. Typically these sensors are synchronized to the patterns being created on the DLP chip and the images processed along with other information such as the relative projector and camera pose in space.

*[a] ian@well.com *[b] bolas@well.com

Emerging Digital Micromirror Device Based Systems and Applications, edited by Larry J. Hornbeck, Michael R. Douglass,
Proc. of SPIE Vol. 7210, 72100H · © 2009 SPIE · CCC code: 0277-786X/09/$18 · doi: 10.1117/12.810589

2.1 Coded Spaces

In addition to cameras, single element sensors can be used – for example, the light from a projector can be used to encode the position within that space. The simplest encoding is to give every pixel in the image a unique 2D code – one way to do this is to encode columns and rows with simple binary or grey code patterns. The DLP chips are 1024 x 768 so 10 bits can be used to encode each of the x and y values across the chip. Then sensors in the light from the projector can very easily determine their location on a 2D plane illuminated by such a projector. This is a tracking approach which could work in underwater environments and as many objects as desired can locate themselves.

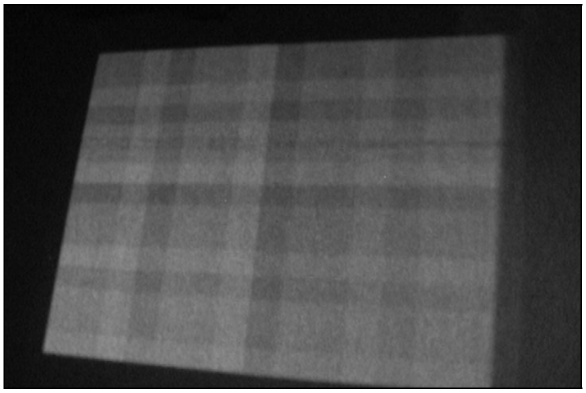

An example 2D pattern

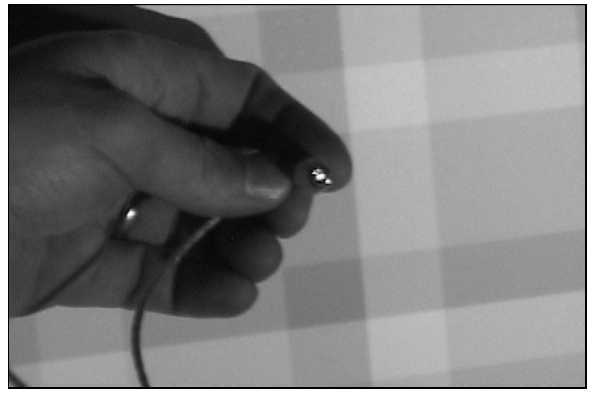

A single photo diode sensor

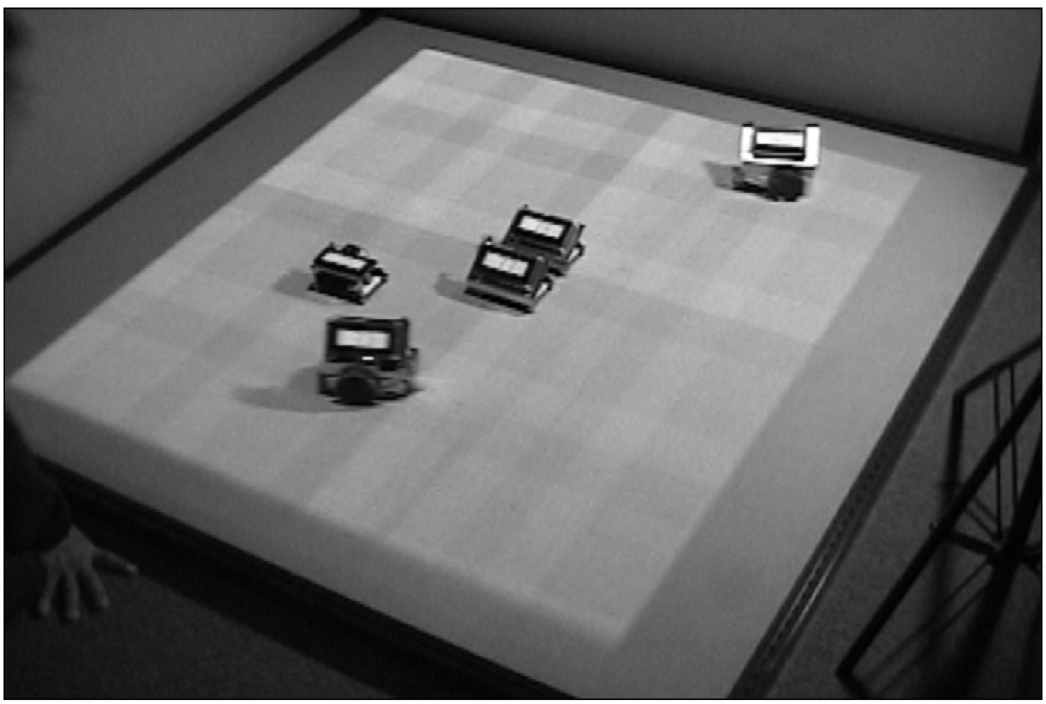

A flock of autonomous mobile devices which locate themselves on the plane and move around.

2.2 Capturing Sub-Surface Features

Computer graphics is increasingly interested in creating images which model real world light transport phenomena. One example of this is the modeling of sub surface scattering. Skin for example both reflects surface light and some portion of the illumination scatters beneath the skin. Scattering behavior changes as blood flow to the skin modulates – flushed cheeks reflect a sudden increase in blood flow to the capillaries near the surface. These scattering changes can be acquired in real time on a moving face using techniques developed in at the ICT Graphics Lab under the direction of Paul Debevec [1,2]. The technique is to project a sequence of 48 stripe patterns onto the face. Images of each pattern falling on the face are recorded. These image sequences are then processed to extract the surface reflection and the subsurface components as separate quantities. This sort of data set enables the creation of more lifelike virtual renderings in computer graphics for movies and video games.

One position of the stripe patterns on the face. The patterns walk over the face rapidly and images are captured during that process.

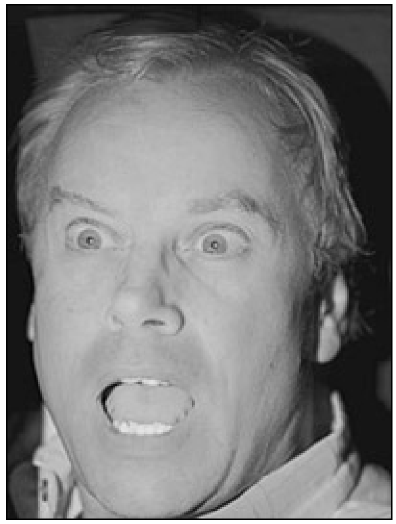

Post processed image showing the average surface reflectance. This image is a composite image of the reflectance from the middle of the white part of the stripes.

Post Processed image showing light that scattered sub-surface, this is an image composed of light reemitted between the stripes – i.e. the middle of the dark stripes.

The three images above illustrate one data capture of this work. The accurate recovery of the surface and subsurface components also requires knowledge of the surface geometry, as described in [2] where patterns are sequenced to both do the stripes above but also to get the surface geometry.

2.3 Physics – the DLP as a laser cavity

The Ginston Lab at Stanford has employed the DLP as an optical filter for creating a programmable wavelength laser. Patterns on the DLP, change the cavity resonance to adjust the center frequency of the laser [4].

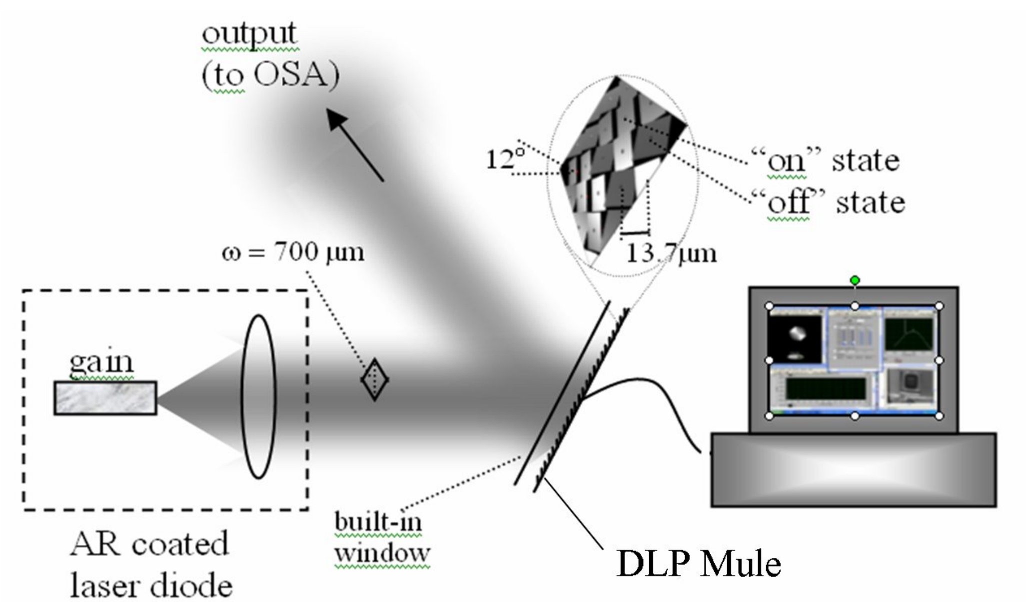

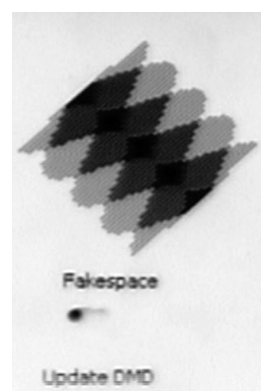

One of the patterns rendered on the DLP chip creating the optical filter; adjusting the pattern changes the laser frequency. See [4] for the physics.

3. PATTERNS FOR DISPLAY

Most DLP chips have gone into displays of various kinds including projectors, and rear projection television type products. Using the DLP as a really fast binary image creation means has enabled them to be used in a number of non traditional displays.

3.1 Multi-Person Stereo

As part of the Emerging Technologies venue at Siggraph, Fakespace Labs collaborated with artist Perry Hoberman to create a piece called Snared Illumination [3]. This artwork used a projector running at 1,440 frames per second. Walking up to the display, monochrome image is seen. When viewed through a special pair of high speed shutter glasses however, a stereoscopic image could be seen. The images shown on the projection were constructed in such a way as to hide images – an image may be hidden by projecting it briefly followed by its inverse. An integrating sensor like the eye will not see either image. In practice, the eye makes small saccades revealing the border between the image and its inverse. This effect can be masked almost entirely however by having another image which is recognizable for the eye to look at. If the masking image is brighter than the image / inverse pair, then the revealing of edges during eye saccades disappears and only the masking image is seen. In the Siggraph piece, the display showed 3 stereo pairs (6 images plus 6 inverses to hide the stereo pairs) and a masking image which was brighter than the stereo images and completely hid the stereo pairs.

A prototype showing the masking image and one eye through a synchronized shutter. The relative brightness has been adjusted. The delightful artworks by artist Perry Hoberman for the actual piece were developed after this photograph and supported three independent stereo views simultaneously.

3.2 Light Field Display

Researchers at the ICT Graphics Lab have used the DLP based fast projectors to create a light field display [6]. This display uses a high speed projection device, and a special spinning surface. The light field created reproduces, with certain limitations, the light emanating from a small scene. Objects shown on the display exhibit correct self occlusion as observers moves relative to the object.

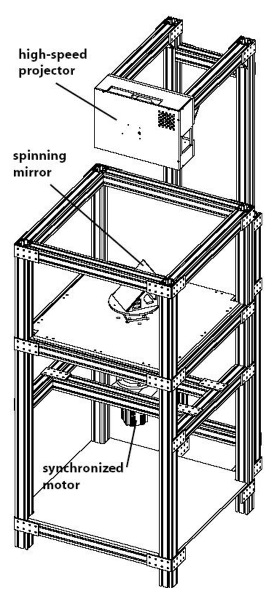

The systems design of the display is as follows. A PC renders stacks of 24 frames and composites them into 24 bit images at 180Hz and these images are delivered to a Mule Projector resulting in a frame rate of 4,320 binary frames per second. The projector projects onto a tilted spinning reflective anisotropic diffuser. The reflective anisotropic diffuser takes an incoming ray of light from the projector and spreads that pencil of rays vertically but not significantly in the horizontal. The spinning anisotropic diffuser is carefully synchronized to the projector.

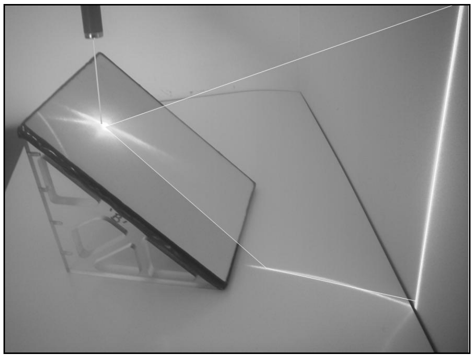

The behavior of the anisotropic diffuser
to incoming rays.

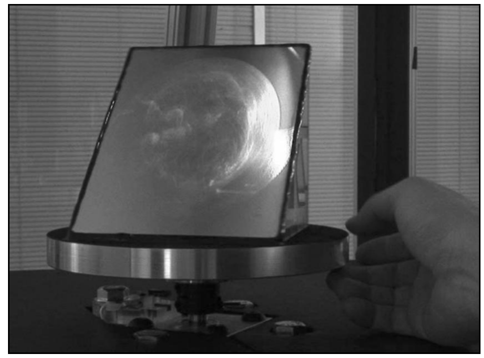

The reflective diffuser on its flywheel is
driven by a motor, not seen here

As the mirror spins, images hitting the reflective diffuser are reflected out into the space around the display. The viewer's eye sees a sequence of different frames from slightly different mirror positions. The first approach explored for creating images for the display was the usual projection type of rendering for all the various viewpoints which have to be rendered. This did create images however, they contained geometric artifacts. These artifacts were corrected by more thoroughly understanding the geometry of the display and realizing that a multi-center of projection rendering approach was required. When implemented, this resulted in more geometrically correct images.

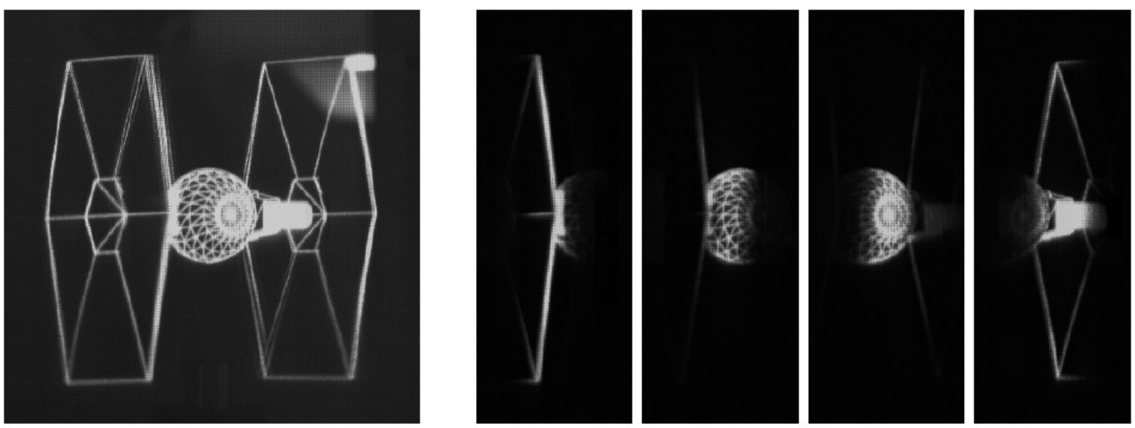

The image on the left is created by the four slices on the left. These four slices are projected in under a millisecond.

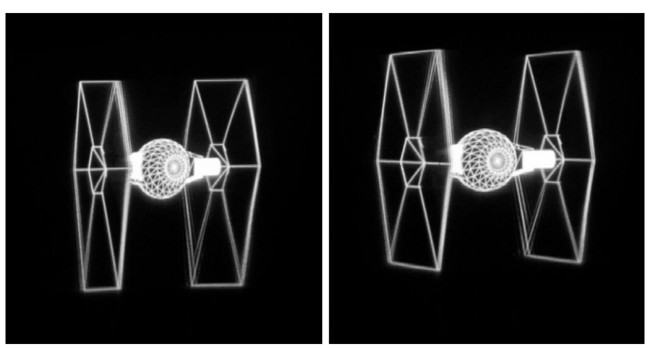

The left image is rendered with a standard perspective model. The right uses
the more geometrically correct multiple center of projection method.

To capture and render light fields of real objects, some additional computation is required to compute which viewing ray should be used to query the light field database. That computation is done as follows for the pixels that make up the image. A ray from a projector pixel intersects the mirror and defines a plane [this is a slight assumption]. This plane is intersected with a circle which defines the height of the viewer at the present time. That intersection defines a point and a ray is then traced back to the mirror and defines the direction used to query the light field database.

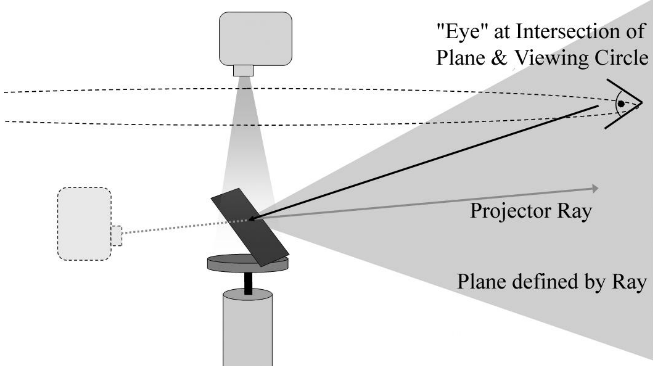

The light field database is a collection of images of a real [or computed] object. These images are taken from "around" the object. At first the display was used with a fixed viewer height, requiring a light field acquisition with a single set of images around an object. This was extended however to facilitate the use of a head tracker with the display to support vertical viewer motion as well as supporting the 360 degrees in the horizontal. To accommodate this, the light field of the object is acquired in sequence with the camera height varying between rings of images.

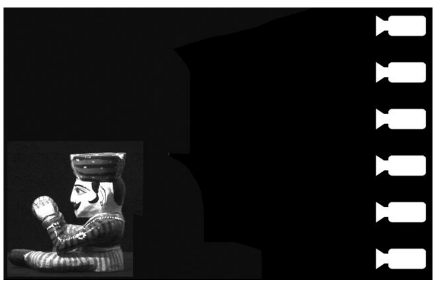

This acquired data set provides images from a regular spaced grid of points on a cylinder around the object, in this example, a small seated figure. Instead of moving an array of cameras, the object is rotated on a small turntable and after each rotation, the camera is moved to the next position and the turntable is rotated again. Note that the lighting used to illuminate the object would ideally be consistent in object space so the illumination is attached to the turntable used to rotate the seated figure.

To display the light field on the display we use the query method described previously and interpolate between cameras to accommodate for the viewer's height.

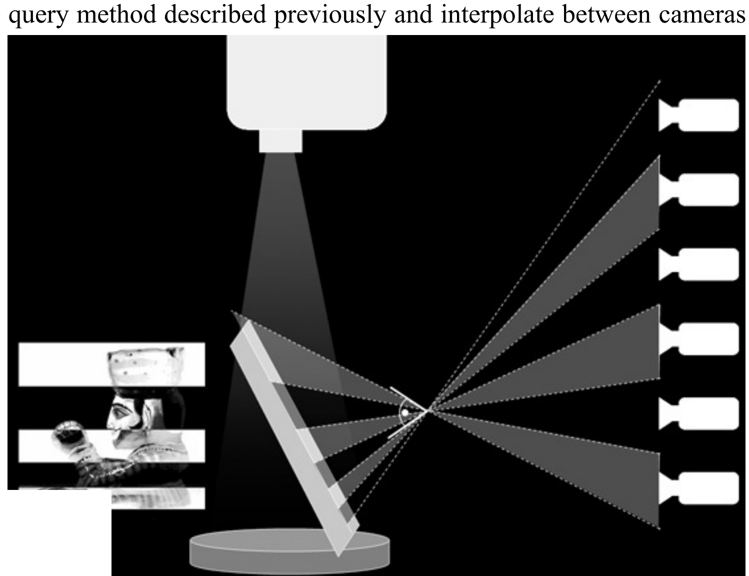

As shown on the right: for the eye position shown, rays in the light field which best match those of the eye come from each of the cameras. The contribution of every other camera has been shaded in light grey. Rays from each camera contribute to reconstructing the correct view. Note that the eye position is not collocated with any of the acquisition cameras. The seated figure is also shown with inverted stripes to illustrate which portion of the view came from which camera. The top view of this arrangement shows interpolation of the two closest camera images views around the seated figure.

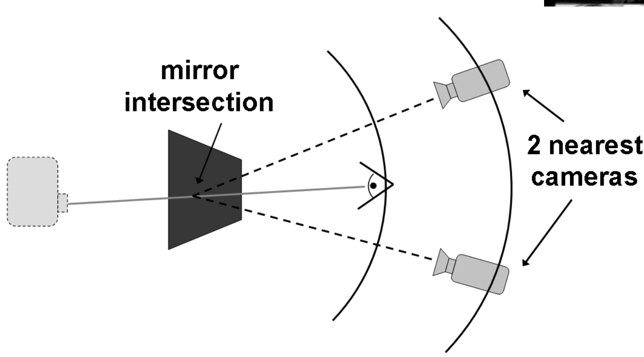

There are two additional steps which have to take place in the above sequence of events. First, the images from the camera capture are perspective images and so that the seated figure appears correctly, those images need to be re-sampled to incorporate multiple centers of projection. Second, as the display is monochrome, the images need to be dithered and this was done using the method described in [7]. With these additional steps, the real object and the virtual reproduction of it behave the same.

This technique can also be used with moving datasets. One of the action sequences to be shown on the display was captured in a large Light Stage. This data set is of a person jogging and is reproduced on the display. One can walk all the way around the little figure running in space. Below are eight frames from the resulting animation.

This work has recently been extended to support live interaction between two people. Shown at the Army Science Conference in 2008, the demonstration included the real time capture of the front of a person's face and the real time display of the 3D face on the display [5]. Audio was also conveyed between the two people. This presented a system in which one is talking to a real 3D image of a face without head tracking. The virtual head can turn to face the person they are talking to for example.

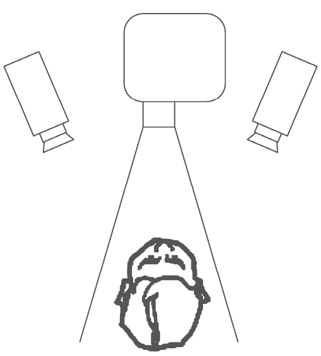

The "capture" side of the system is comprised of a pair of video cameras and a projector. The projector is creating stripe patterns and the cameras are acquiring images in such a way as to enable the 3D reconstruction of the 3D geometry of the face. The images from the cameras also provide texture information. Images are processed for the 3D display and a light field suitable for use with the display is calculated on the fly and sent to the graphics card in real time. This configuration only allows the reconstruction of the front side of the face and this is displayed as a textured convex hull on the 3D display.

The person being "captured" can see the visitor on a standard 2D display and thus has an experience similar to a video conference.

The 3D Teleconference system demonstrated at the Army Science Conference 2008.

4. ACKNOWLEDGEMENTS

This work is the product of many people. The original high speed projector work was supported by a grant from the Office of Naval Research under the direction of Ralph Wachter and Larry Rosenblum. The researchers at the ICT Graphics Lab are inspired, accomplished, and dedicated; they include Paul Debevec, Andrew Jones,

Magnus Lang, Xueming Yu, Naho Inamoto, Bruce Lamond, Charles-Felix Chabert, Pieter Peers, Tim Hawkins, Brian Miller, Bruce Lamond, Jacki Morie, Sean Bouchard, Dell Lunceford, Tom Pereira, , Bill Swartout, Randy Hill, and Randolph Hall. USC's work is sponsored by the U.S. Army Research, Development, and Engineering Command (RDECOM) and the University of Southern California Office of the Provost. Jeff Fisher and Sovann Neak created many of the mechanical and electronic implementations. Artist Perry Hoberman worked on the Snared Illumination piece and created a very compelling experience.

REFERENCES

[1] Sarah Tariq, Andrew Gardner, Ignacio Llamas, Andrew Jones, Paul Debevec, Greg Turk; *Efficient Estimation of Spatially Varying Subsurface Scattering Parameter;* 11th Int'l Fall Workshop on Vision, Modeling, and Visualzation 2006 Aachen, Germany

[2] Sarah Tariq, Andrew Gardner, Ignacio Llamas, Andrew Jones, Paul Debevec and Greg Turk, *Efficient Estimation of Spatially Varying Subsurface Scattering;* USC ICT Technical Report ICT-TR-01-2006, July 27, 2006

[3] Mark Bolas, Ian McDowall, Perry Hoberman, *Snared Illumination*, Emerging Technologies, Siggraph 2004

[4] Ruslan Belikov, Olav Solgaard, *Optical wavelength filtering by diffraction from a surface relief*; Opt. Lett. 28, 2003

[5] Andrew Jones,Magnus Lang,Xueming Yu, Ian McDowall, Mark Bolas, Paul Debevec; *Live 3D Teleconferencing* ; Army Science Conference 2008

[6] Andrew Jones, Ian McDowall, Hideshi Yamada, Mark Bolas, Paul Debevec; *Rendering for an Interactive 360º Light Field Display*; SIGGRAPH 2007 Papers Proceedings and SIGGRAPH 2007 Emerging Technologies

[7] Ostromoukhov, V. *A simple and efficient error-diffusion algorithm*. In Proceedings of ACM SIGGRAPH 2001