Temporal Dithering of Illumination
for Fast Active Vision

Srinivasa G. Narasimhan[1], Sanjeev J. Koppal[1], and Shuntaro Yamazaki[2]

[1] The Robotics Institute, Carnegie Mellon University, USA
[2] National Institute of Advanced Industrial Science and Technology, Japan

Abstract. Active vision techniques use programmable light sources, such as projectors, whose intensities can be controlled over space and time. We present a broad framework for fast active vision using Digital Light Processing (DLP) projectors. The digital micromirror array (DMD) in a DLP projector is capable of switching mirrors "on" and "off" at high speeds ($10^6/s$). An off-the-shelf DLP projector, however, effectively operates at much lower rates (30-60Hz) by emitting smaller intensities that are integrated over time by a sensor (eye or camera) to produce the desired brightness value. Our key idea is to exploit this "temporal dithering" of illumination, as observed by a high-speed camera. The dithering encodes each brightness value uniquely and may be used in conjunction with virtually any active vision technique. We apply our approach to five well-known problems: (a) structured light-based range finding, (b) photometric stereo, (c) illumination de-multiplexing, (d) high frequency preserving motion-blur and (e) separation of direct and global scene components, achieving significant speedups in performance. In all our methods, the projector receives a single image as input whereas the camera acquires a sequence of frames.

1 Introduction

Projectors are commonly used as programmable light sources for a variety of active vision techniques including structured light range finding [1,2,3,4,5,6], photometry-based reconstruction [7,8], relighting [9], light transport analysis [10,11] and depth from defocus [12]. The intensity and color of the scene's illumination can be controlled over space and time depending on the task at hand. For instance, projecting a set of colored striped patterns onto a scene alleviates the problem of finding correspondences for 3D reconstruction [2].

Recently, Digital Light Processing (DLP) technology (`http://www.dlp.com`) has enabled mass production of low cost projectors with high quality. The key component of a DLP projector is the Digital Micromirror Device (DMD). Each mirror in the DMD is 14×14 microns and can switch between two orientations, $+12^o$ and -12^o [13]. In one orientation, incident light is reflected by the mirror toward the outside scene and in the other, light is reflected onto a black surface within the projector. These mirrors can switch between orientations in a few microseconds, enabling high precision control of illumination. As a result, the

DMD device has found applications in areas ranging widely from microscopy to chemistry to holographic displays [13].

The operating principle of the DMD device has also been exploited in computer vision and graphics. Nayar et al. [14] re-engineer a DLP projector into a DMD-camera and demonstrate the notion of programmable imaging for applications including adaptive dynamic range and optical filtering and matching. Based on the theory of compressive sampling, a single pixel camera has been implemented where the DMD device used to compute optical projections of scene radiance [15]. Jones et al. [16] modify a DLP projector using custom made FPGA-based circuitry to obtain 1-bit projection at 4800Hz. Using this, they generate high speed stereoscopic light field displays. McDowall and Bolas [17] use a specially re-programmed high speed projector based on Multiuse Light Engine (MULE) technology to achieve range finding at kilohertz rates.

1.1   Temporal Dithering in a DLP Projector

In order to project a desired intensity value, the DLP projector emits a series of light pulses of different time intervals [13]. A sensor aggregates the pulses of light over the duration of its integration time (say, 1/30s in a video camera) to capture the final gray-valued brightness. This Pulse-Width modulation (PWM) by the projector is unique for every input intensity and can be termed as "temporal dithering" of the illumination. As we shall show, this dithering allows us to encode scene illumination in novel ways to achieve significant speedup in the performance of virtually any active vision technique.

But how do we capture this high speed dithering? The exposure time (1/30s) of a video camera is too long to observe the temporal illumination dithering clearly. One possibility is to precisely synchronize the camera with a DLP projector and to expose the camera only for the duration of a single projected light pulse (a few microseconds). Raskar et al [18] and Cotting et al [19] use this technique to embed illumination patterns in the scene that cannot be observed with the naked eye. The focus of these works is on intelligent office applications with 30-60Hz performance requirements.

In contrast, our work focuses on exploiting the temporal dithering for fast active vision. For this, we use a novel combination of a high speed camera and an off-the-shelf DLP projector. Figure 1 illustrates the dithering of an 8-bit InFocus IN38 DLP projector as observed by a Photron PCI-1024 high speed camera. A calibration image composed of 5×5 pixel blocks each with a different intensity value from 0 to 255 is input to the projector. Each intensity at a pixel C in this calibration image is projected onto a flat screen using a unique temporal dithering $D_C(t)$, over discrete time frames t. The high speed camera observes the projected images at 10 kHz. Notice the significant variation in the images recorded. The plot in Figure 1(d) shows the patterns emitted by the projector for 4 input brightnesses $(165, 187, 215, 255)$, as measured over 100 camera frames. The temporal ditherings corresponding to all the 256 input intensities in the calibration image are collated into a photograph for better visualization of this principle. The temporal dithering is stable and repeatable but varies for each projector-camera system.
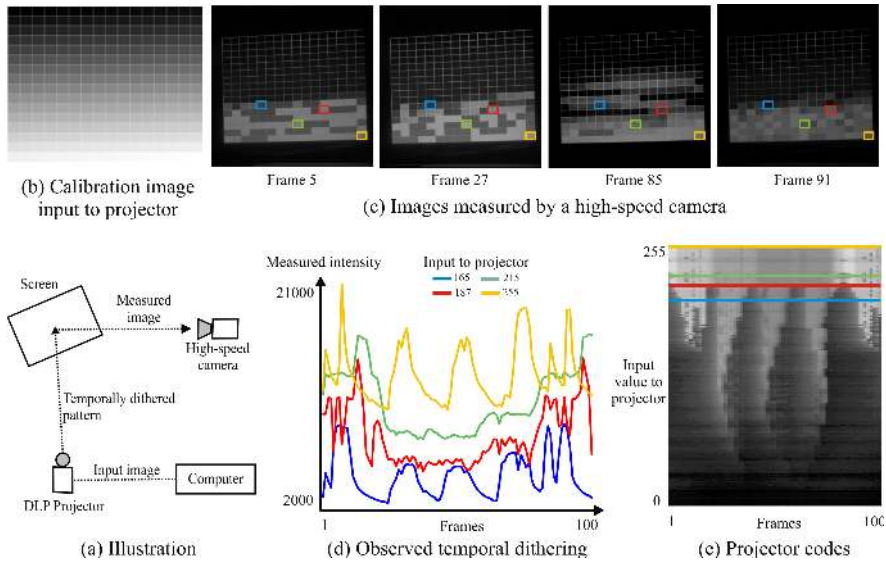
(b) Calibration image input to projector

(c) Images measured by a high-speed camera

Frame 5   Frame 27   Frame 85   Frame 91

(a) Illustration

(d) Observed temporal dithering

(e) Projector codes

Fig. 1. Reverse engineering a DLP Projector: (a) A DLP projector converts the input intensity received into a stream of light pulses that is then projected onto a screen. A high speed camera viewing the screen aggregates the brightness over the duration of its integration time. (b) A calibration image composed of 5×5 pixel blocks each with a different intensity from 0 to 255 is input to the projector. (c) The camera records the projector output at 10 kHz. In (d) we show gray-valued intensities measured over time by the high speed camera for 4 example intensities input to the projector. Notice the significant variations in the plots. In (e), the temporal dithering for all 256 projector input intensities is collated into an image. This temporal dithering is repeatable and can be used to encode illumination in a novel way, enabling fast active vision.

1.2   Implications of Temporal Dithering

The high speed illumination modulation of a DLP projector can be exploited to speed up a series of well-known active vision problems, making them applicable to dynamic scenes. For each of these problems, we select a simple existing algorithm to demonstrate our technique, although more sophisticated algorithms may be used to achieve further speed up:

(a) The unique encoding of intensities allows us to obtain camera-projector pixel correspondences allowing 3D reconstruction at high speeds.

(b) By multiplexing illumination from three projectors, we compute the surface normals using photometric stereo [8] at high speeds.

(c) We de-multiplex illumination [20] from multiple projectors to capture the appearances of a dynamic scene from different lighting directions.

(d) We demonstrate the ability to project high frequency complementary patterns to separate the direct and global components [10] in a dynamic scene.

(e) We discuss motion blurring of an object illuminated by a DLP projector and captured by a low frame rate camera (30-60 Hz). The temporal dithering

Fig. 2. Illumination and acquisition setup for structured light based 3D reconstruction: The Photron high speed camera is placed vertically above the Infocus DLP projector. A vertical plane is placed behind the scene (statue) for calibration.

preserves higher frequencies in the motion-blurred image. This is similar to the work of Raskar et al [21] who demonstrate that fast camera shutter modulation during image acquisition preserves higher frequencies.

In methods (a)-(d), the projector receives a single image as input via a computer, whereas the high speed camera acquires a sequence of frames. The effective speedup achieved depends on the task at hand and the quality of the result desired given the signal-to-noise ratio in the captured images. In addition, the intensity variation due to dithering can be observed reliably even with camera frame rates as low as 300 fps enabling applications with slower performance requirements. Unlike previous work, our techniques do not require any projector-camera synchronization, hardware modification or re-programming of the DMD device, or the knowledge of proprietary dithering coding schemes. Thus, we believe this work to be widely applicable. Better visualizations of all our results are available through our website (`http://graphics.cs.cmu.edu/projects/dlp-dithering`).

2   Projector-Camera Correspondence for Reconstruction

Structured light-based triangulation has commonly been used for 3D reconstruction [1]. A known set of spatio-temporally varying intensity patterns is projected onto a scene and the reflected images are used to find the corresponding pixels between the projector and the camera. The corresponding pixels are then triangulated spatially (or by temporal analysis [4]) to obtain 3D structure. It is assumed that the scene motion is negligible while the patterns are projected. Since projectors have been operated at 30-60Hz, most implementations achieve slower than real-time performances. Fast illumination dithering in a DLP projector enables high speed reconstruction.

Our goal is to obtain correspondences between the projector and camera pixels at high speeds. Consider a high speed camera viewing a dynamic scene that is illuminated by the DLP projector. A *single* image composed of a set of horizontal lines of randomly chosen colors and intensities is input to the projector via a
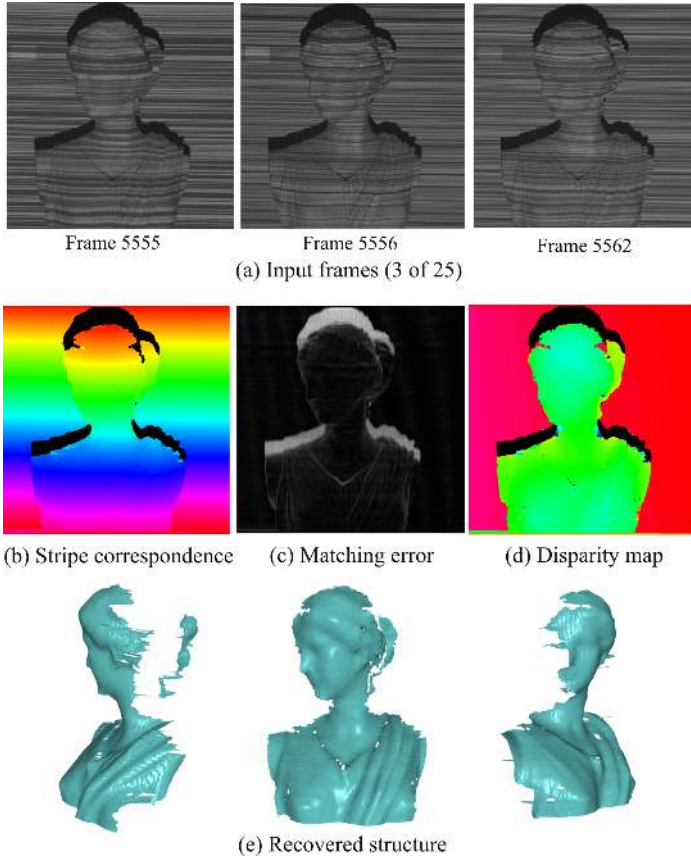
Frame 5555   Frame 5556   Frame 5562

(a) Input frames (3 of 25)

(b) Stripe correspondence   (c) Matching error   (d) Disparity map

(e) Recovered structure

Fig. 3. Results of 3D reconstruction using the DLP projector for a moving statue: (a) Three frames captured by the high speed camera illustrate the fast modulation of illumination incident on the scene. 20 continuous frames are used to match the intensity variation observed on the scene point against the normalized intensity variation observed on the vertical plane behind the object. (b) The best match finds correspondences between projector and camera pixels. The error map is shown in (c). The (d) disparity and (e) recovered shape of the object is shown from different viewpoints.

laptop. Let $I(t)$ be the vector of intensities observed, over a set of frames, at a scene point P. The normalized correlation between $I(t)$ and temporal dithering function $D_C(t)$ for each C (Section 1.1) is computed to obtain the projector pixel C corresponding to the image pixel P. But how do we synchronize the frames from the projector and the camera? One approach is to include a small planar patch in the scene where correspondence between the corners of the patch can be easily established (say, manually). This correspondence allows us to synchronize the measured intensity vector with the temporal dithering.

We performed two experiments with a rotating statue and with a cloth waved quickly in front of the high speed camera. For convenience, the camera and the
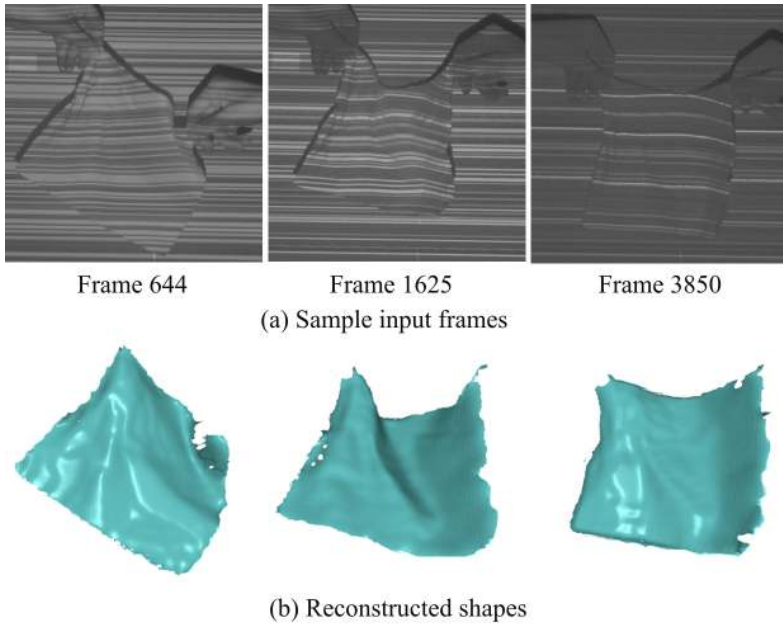
Frame 644   Frame 1625   Frame 3850

(a) Sample input frames

(b) Reconstructed shapes

Fig. 4. 3D reconstruction of a cloth waved: (a) Twenty continuous images acquired are used to reconstruct the 3D structure of a dynamically deforming cloth object. We show three sample input images here taken at different times. (b) The reconstructed cloth in different frames shows complex structure such as folds and creases that can only be captured at high speeds.

projector are placed in a fronto-parallel configuration with a vertical baseline (see Figure 2). The high speed camera is operated at 3kHz and the projector is reverse engineered at this speed as described in Section 1.1. A homogeneous vertical plane is used as the background behind the scene. The dithering $D_C(t)$ can be captured from the pixels on this plane and simultaneously with the object. Hence, in this setting, we simply correlate the normalized brightness vector $I(t)$ of a point on the object with the dithering vector $D_C(t)$ at every pixel C on the background plane and no additional synchronization is needed.

Here, twenty frames were used to obtain correspondences, taking $20/3000 = 1/150s$. In general, the number of frames necessary depends on the desired matching accuracy and the SNR of the acquired images. By sliding the 20 frame window across the full image sequence, 3D reconstruction can be obtained at rates as high as 3 kHz (speed of camera). However, in practice, the rate of reconstruction is lower considering the speed of the object's motion. Figures 3 and 4 show the correspondences, the disparity and the reconstructions obtained for the two scenes.

Note that the dithering can also be observed at lower frame rates and hence a lower cost camera may be used for slower reconstructions. We repeated our experiments for the same scene at four lower frame rates — 1500Hz, 600Hz,
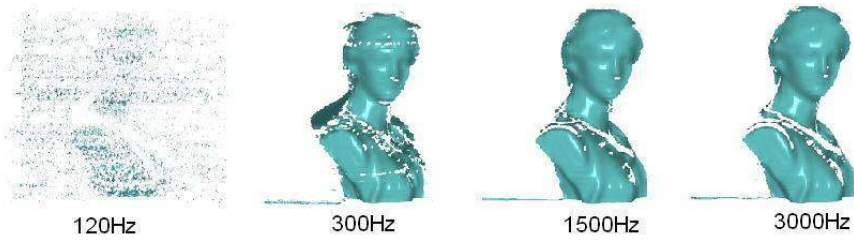
Fig. 5. Reconstructions obtained using videos captured at reduced frame rates. Even at 300Hz, the quality of the reconstruction obtained remains acceptable indicating that temporal dithering can be exploited at this frame rate.

300Hz and 120Hz. Figure 5 shows the reconstructions obtained. The frame rate of 120Hz is too low to capture the required intensity variation and hence, the projector-camera pixel correspondences are unreliable. However, at 300Hz, the reconstruction quality is still acceptable indicating that the temporal dithering can be exploited even at this frame rate.

3   Illumination De-multiplexing for Appearance Capture

Acquiring scene appearance from multiple lighting directions is necessary for image-based relighting and photometric recovery of scene properties (BRDF and 3D structure). In most works, the scene is assumed to be static and the acquisition can take any where from a few seconds to several minutes. Using DLP illumination, we capture the appearances of a dynamic scene from multiple lighting directions. For this, we draw upon the idea of illumination de-multiplexing [20], where the images of the scene are simultaneously captured from multiple source directions and de-multiplexed in software to obtain the desired images under each lighting direction. This technique increases the signal-to-noise of the captured images while keeping the number of captured images unchanged.

The difference between Schechner et al. [20] and our technique is in the coding: they use binary Hadamard codes, whereas we rely on the temporal dithering of DLP illumination. The acquisition setup consists of three DLP projectors (Infocus IN38 and LP120, and Sony XGA DataProjector) that simultaneously illuminate the scene from different directions. Since we wish to illuminate the scene uniformly, a single constant brightness image is input to each of the projectors.

The three projectors differ in their brightness and contrast ratings and dithering behaviors. The captured intensity at time instant t is written as a sum of irradiances due to the illuminations from all projectors ($k = 1 \ldots 3$):

$$I(t) = \sum_{k=1}^{3} D_k(t) E_k(t) \tag{1}$$

where, $D_k(t)$ is the dithering intensity of the projector k at time t and $E_k(t)$ is the irradiance due to the scene as if illuminated only from projector k but with
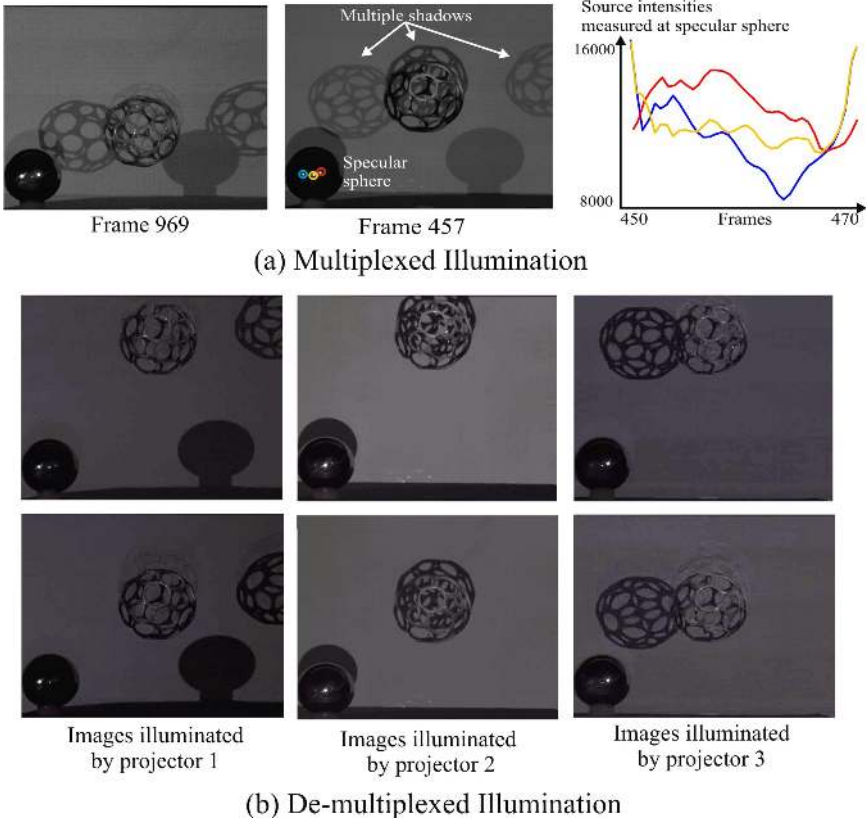
Fig. 6. Demultiplexing illumination from three projectors to create appearances under each lighting direction: The scene consists of a wiry polyhedral ball falling vertically. Notice the mixing of the shadows in the multiplexed images in (a). A mirror sphere is placed in the scene to measure the dithering intensities D_k from the three projectors. (b) The results of demultiplexing for two frames in the video sequence. A total of 10 frames are used to demultiplex. Some ghosting artifacts are due to the noisy estimation of source intensities D_k. The variation in the contrast and brightness of the three demultiplexed images are due to the quality of the three projectors. Projector 2 (InFocus IN38) is the most recent one and has the highest quality.

unit intensity. The intensities $D_k(t)$ can be obtained by observing a stationary mirror sphere placed in the scene. The observed irradiances $I(t)$ over time form a linear system which is solved to obtain the appearances $E_k(t)$ of the scene from each individual lighting direction. In practice, since the projectors are not synchronized when they illuminate the scene, the dithering intensities D_k vary significantly over time, and hence the linear system is well-conditioned.

Figure 6 shows the results of applying the above approach to a scene with a falling wiry ball. Notice the 3 shadows of the ball and the mirror sphere that

appear mixed in the multiplexed image $I(t)$. For robustness, we use 10 frames to solve the above linear system. Notice separation of the shadows in the de-multiplexed images. As before, the effective rate of demultiplexing depends on the SNR in the high speed camera. We have thus far ignored color information, however, when the three DLP projectors emit intensities in different spectral bands, the de-multiplexing algorithm can be used to colorize the acquired high speed gray-scale video.

4   Illumination Multiplexing for Photometric Stereo

Photometric stereo is a widely used method to recover the surface normals and albedos of objects that are photographed under different lighting directions. There are many variants of this approach and we chose the one by Hertzmann and Seitz [8] for its simplicity. In their work, the appearance of the scene under varying lighting is matched with that of an example sphere made of the same material (same BRDF) as the scene. The point on the sphere that produces the best match is the normal of the scene point. We will extend this approach for fast moving scenes that are simultaneously illuminated from different directions.

The scene in our experiments consists of a sphere and a falling pear both painted in the same manner (Figure 7) and illuminated by three DLP projectors simultaneously from different directions. The projectors and camera are far enough away from the scene to assume orthographic viewing and distant lighting. Since each projector must uniformly illuminate the scene, we provide a single constant brightness image as input to each projector (with different brightness values). The high speed camera records images at 3 kHz.

The projectors are de-synchronized and hence, the "multiplexed illumination" results in significant variation in the observed intensities. The normalized intensities at a scene point are compared to those observed on the sphere. The surface normal of the scene point is that of the point on the sphere which produced the best match. A matching length of 10 frames achieved robust results. A sliding window of 10 frames can be used to generate the normals up to a rate of 3 kHz. As before, the speed of the object determines the effective performance rate. Figure 7 shows the normals of the pear as it falls and bounces on a table.

5   Complementary Patterns for Direct-Global Separation

The radiance of a scene point can be divided into two components - (a) the direct component L_d, due to the direct illumination from the light source and (b) the global component L_g due to the illumination indirectly reaching the scene point from other locations in the scene [10]. The global component L_g includes effects like interreflections, subsurface and volumetric scattering and translucency. Nayar et al [10] demonstrated that using high frequency illumination, it is possible to separate the two components and obtain novel visualizations of the components for the first time. A particular choice for high frequency illumination is
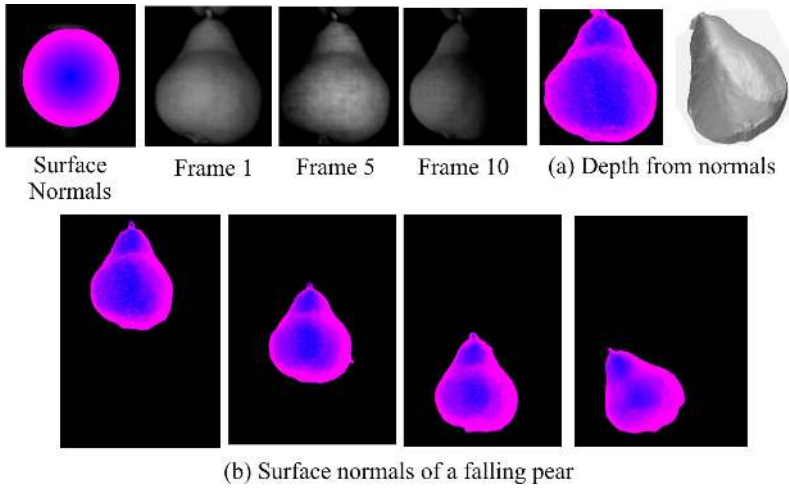
Surface Normals   Frame 1   Frame 5   Frame 10   (a) Depth from normals

(b) Surface normals of a falling pear

Fig. 7. Photometric stereo by example: The scene consists of a fast moving pear and a sphere that are both painted similarly. Three DLP projectors simultaneously illuminate the scene and the camera operates at 3000Hz. The projectors and camera are far enough away from the scene to assume orthographic viewing and distant lighting. The surface normal at a point on the falling pear is computed by matching the normalized observed intensities to those at the points on the sphere. Since the projectors are not synchronized, the variation in multiplexed illumination from the 3 projectors is significant enough to obtain good matches for surface normals. A matching length of 10 frames achieved robust results.

a checker board pattern and its complement (with alternate bright and dark squares), both of which are projected sequentially for separation.

We exploit illumination dithering to obtain separation at video rates. However, in our setup, it is possible to input only one image to the DLP projector in 1/60s and we have no control over the temporal dithering. So, how do we project complementary patterns much faster than 1/60s? We selected two specific input brightnesses 113 and 116 whose dithered patterns are shown in the plot of Figure 8. Notice how the two patterns "flip" from bright to dark and vice versa over time. Hence, a checker pattern with these two brightnesses are input to the projector. The dithering ensures that the two complementary patterns occur at high speeds. Let the observed temporally dithered values for input values 113 and 116 be a and b, respectively, and the fraction of pixels that correspond to the value a be α (0.5 in our experiments). The two captured images are [10]:

$$L^+(x,y) = aL_d + [(1-\alpha)b + \alpha a]L_g$$
$$L^-(x,y) = bL_d + [(1-\alpha)a + \alpha b]L_g \, . \tag{2}$$

To solve the above equations, we need to know a and b in every frame. For this, we place a white planar diffuse surface behind the scene of interest. For points on this plane, $L_g = 0$ and L_d is a constant. This allows us to estimate a and b up to a

(a) Co-located projector-camera setup

Frame 2719 with
direct and global components

Measured
intensity
17000

Frame 2727
Frame 2730

3000

2721   Frames   2821

Frame 2727   Frame 2730   Direct component   Global component

Frame 11731   Frame 11733   Direct component   Global component

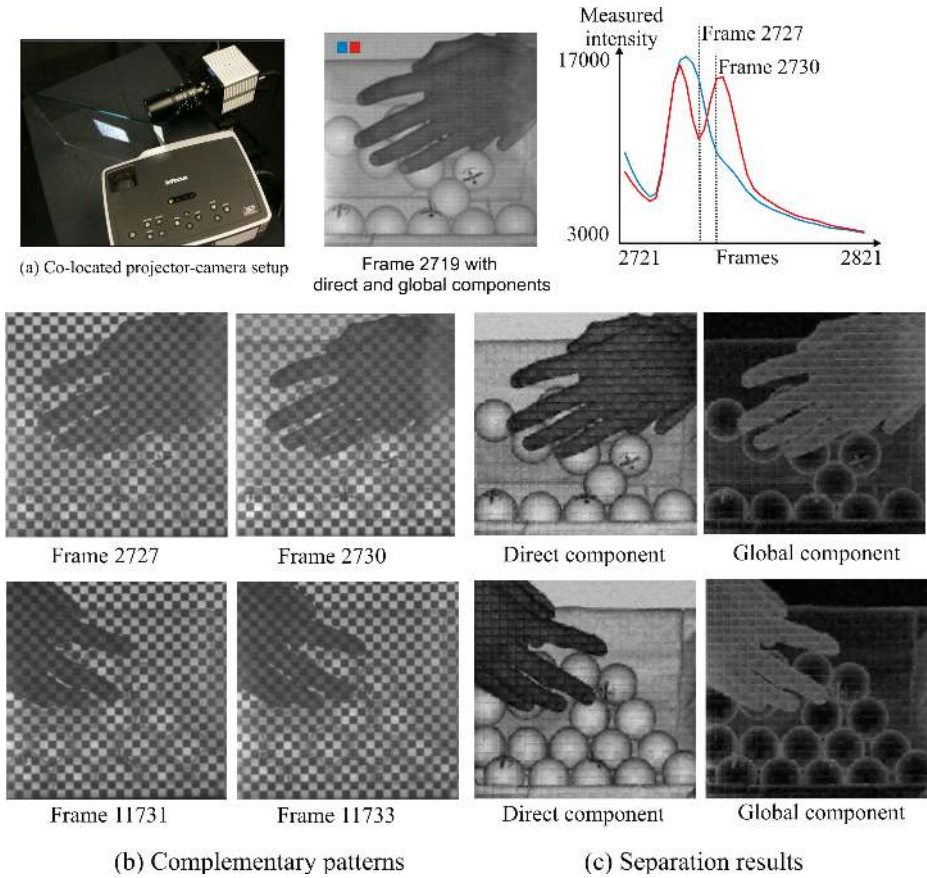(b) Complementary patterns   (c) Separation results

Fig. 8. Direct-Global Separation using DLP Dithering: (a) The DLP projector and the camera are co-located using a beam splitter. A single checker pattern with two intensities 113 and 116 are input to the projector. The plot shows how the input intensities are dithered by the projector over time. Notice that at certain time instants, the patterns flip between bright and dark. Thus, the projector emits complementary checker patterns as in (b) onto the scene that are used to separate the direct and global components (c). The flip occurs once in 1/100s.

single scale factor. Then, the above linear system can be solved at every pixel to obtain the separation. There is one additional complication in our setup beyond the method in [10]: it is hard to find out whether a scene point receives intensity a or intensity b from just the observed appearance of the scene. To address this problem, we co-locate the projector and the camera using a beam-splitter as shown in Figure 8. The pixels of the projector are automatically corresponded with those of the camera.

The scene in our experiment consists of a set of white ping-pong balls dropped from a hand. The ping-pong balls are mostly diffuse. Notice that the direct
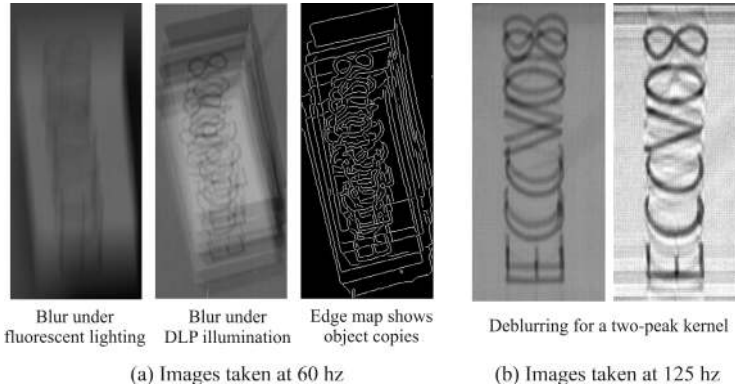
Blur under fluorescent lighting | Blur under DLP illumination | Edge map shows object copies

(a) Images taken at 60 hz

Deblurring for a two-peak kernel

(b) Images taken at 125 hz

Fig. 9. Motion blurring under DLP illumination and fluorescent illumination: The scene consists of a heavy brick falling rapidly and an image is captured with exposures 1/60s (a) and 1/125s (b). Under fluorescent illumination, the motion blur appears as a smear across the image losing high frequencies. The temporal dithering in DLP projectors acts as a high frequency modulator that convolves with the moving object. The motion-blurred image still preserves some of the high spatial frequencies. Six copies of the text "ECCV08" in (a) and 2 copies in (b) are clearly visible.

component for each ball looks like the shading on a sphere (with dark edges) and the indirect component includes the interreflections between the balls (notice the bright edges). For the hand, the direct component is only due to reflection by the oils near the skin surface and is dark. The indirect component includes the effect of subsurface scattering and dominates the intensity. The checker pattern "flips" once in approximately 1/100s and hence we achieve separation at 100Hz. Due to finite resolution of the camera and the narrow depth of field of the projector, a 1-pixel blur is seen at the edges of the checker pattern. This results in the grid artifacts seen in the results.

6   Flutter Flash: Motion-Blur under DLP Illumination

Motion-blur occurs when the scene moves more than a pixel within the integration time of a camera. The blur is computed as the convolution of the scene motion with a box filter of width equal to the camera integration time. Thus, images captured of fast moving objects cause a smear across the pixels losing significant high frequencies. Deblurring images is a challenging task that many works have addressed with limited success. A recent approach by Raskar et al. [21] uses an electronically controlled shutter in front of the camera to modulate the incoming irradiance at speeds far greater than the motion of the object. In other words, the box filter is replaced by a series of short pulses of different widths. The new convolution between the object motion and the series of short pulses results in images that preserve more high frequencies as compared to the box filter. This "Flutter Shutter" approach helps in making the problem better

conditioned. Our approach is similar in spirit to [21] with one difference: the fast shutter is simulated by the temporal dithering of the DLP illumination. Note that the DLP illumination dithering is significantly faster than mechanical shutters[1].

Figure 9 shows the images captured by with 1/60s exposure. The scene consists of a brick with the writing "ECCV08" falling vertically. When illuminated by a fluorescent source, the resulting motion-blur appears like a smear across the image. On the other hand, when the scene is illuminated using a DLP projector, we see 6 distinct copies of the text that are translated downward. A Canny edge detector is applied to the captured image to illustrate the copies. If we knew the extent of motion in the image, the locations of strong edges can be used as a train of delta signals that can be used for deblurring the image. In 9(b), we show an example of deblurring the image captured with 1/125s exposure. As in the deblurred images obtained the flutter shutter case, the DLP illumination preserves more high frequencies in the motion-blurred image.

7   Discussion

Speed vs. accuracy trade-off. One limitation of our approach is the requirement of a high speed camera. The acquisition speed of the camera and the effective speed of performance achieved depend on the task at hand and the signal-to-noise ratio of the captured images. For instance, the decision to use 10 frames for demultiplexing illumination or photometric stereo, or to use 20 frames for structured light, was mainly influenced by the noise characteristics of the camera. A more scientific exploration of this trade-off is required to better understand the benefits of our approach to each technique. A future avenue of research is to design 2D spatial intensity patterns that create temporal dithering codes that are optimal for the task at hand.

Issues in reverse engineering. The images shown in Figure 1 are dark for the input brightness range of 0 to 90. Despite the claim from manufacturers that the projector displays 8-bits of information, only about 160 patterns are usable for our experiments. To compensate for this, the projector performs spatial dithering in addition to temporal dithering in a few pixel blocks. This is an almost random effect that is not possible to reverse engineer without proprietary information from the manufacturers. We simply average a small neighborhood or discard such neighborhoods from our processing.

Other active vision techniques and illumination modulations. We believe that the temporal illumination dithering can be applied to a broader range of methods including pixel-wise optical flow estimation and tracking, projector defocus compensation and depth from defocus [12] and spectral de-multiplexing. While we exploit the temporal dithering already built-in to the projector, we do not have a way of controlling it explicitly. Better control is obtained by using a more expensive and special high speed MULE projector [17]. Finally, strobe lighting, fast LED [22] and flash modulation are also effective in temporally varying (not dithering) the illumination.

---

[1] Faster shutters can be realized by electronically triggering the camera.

Acknowledgements

References

1. Will, P.M., Pennington, K.S.: Grid coding: A preprocessing technique for robot and machine vision. AI 2 (1971)
2. Zhang, L., Curless, B., Seitz, S.M.: Rapid shape acquisition using color structured light and multi-pass dynamic programming. 3DPVT (2002)
3. Davis, J., Nehab, D., Ramamoothi, R., Rusinkiewicz, S.: Spacetime stereo: A unifying framework for depth from triangulation. In: IEEE CVPR (2003)
4. Curless, B., Levoy, M.: Better optical triangulation through spacetime analysis. In: ICCV (1995)
5. Young, M., Beeson, E., Davis, J., Rusinkiewicz, S., Ramamoorthi, R.: Viewpoint-coded structured light. In: IEEE CVPR (2007)
6. Scharstein, D., Szeliski, R.: High-accuracy stereo depth maps using structured light. In: CVPR (2003)
7. Zickler, T., Belhumeur, P., Kriegman, D.J.: Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In: Heyden, A., Sparr, G., Nielsen, M., Johansen, P. (eds.) ECCV 2002. LNCS, vol. 2352, pp. 869–884. Springer, Heidelberg (2002)
8. Hertzmann, A., Seitz, S.M.: Shape and materials by example: A photometric stereo approach. In: IEEE CVPR (2003)
9. Wenger, A., Gardner, A., Tchou, C., Unger, J., Hawkins, T., Debevec, P.: Performance relighting and reflectance transformation with time-multiplexed illumination. ACM SIGGRAPH (2005)
10. Nayar, S.K., Krishnan, G., Grossberg, M.D., Raskar, R.: Fast separation of direct and global components of a scene using high frequency illumination. ACM SIGGRAPH (2006)
11. Sen, P., Chen, B., Garg, G., Marschner, S.R., Horowitz, M., Levoy, M., Lensch, H.P.A.: Dual photography. ACM SIGGRAPH (2005)
12. Zhang, L., Nayar, S.K.: Projection defocus analysis for scene capture and image display. ACM SIGGRAPH (2006)
13. Dudley, D., Duncan, W., Slaughter, J.: Emerging digital micromirror device (dmd) applications. In: Proc. of SPIE, vol. 4985 (2003)
14. Nayar, S.K., Branzoi, V., Boult, T.: Programmable imaging using a digital micromirror array. In: IEEE CVPR (2004)
15. Takhar, D., Laska, J., Wakin, M., Duarte, M., Baron, D., Sarvotham, S., Kelly, K., Baraniuk, R.: A new compressive imaging camera architecture using optical-domain compression. Computational Imaging IV at SPIE Electronic Imaging (2006)
16. Jones, A., McDowall, I., Yamada, H., Bolas, M., Debevec, P.: Rendering for an interactive 360 degree light field display. ACM SIGGRAPH (2007)
17. McDowall, I., Bolas, M.: Fast light for display, sensing and control applications. In: IEEE VR Workshop on Emerging Display Technologies (2005)
18. 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. ACM SIGGRAPH (1998)

19. Cotting, D., Naef, M., Gross, M., Fuchs, H.: Embedding imperceptible patterns into projected images for simultaneous acquisition and display. In: ISMAR (2004)
20. Schechner, Y.Y., Nayar, S.K., Belhumeur, P.N.: A theory of multiplexed illumination. In: ICCV (2003)
21. Raskar, R., Agrawal, A., Tumblin, J.: Coded exposure photography: Motion deblurring using fluttered shutter. ACM SIGGRAPH (2006)
22. Nii, H., Sugimoto, M., Inami, M.: Smart light-ultra high speed projector for spatial multiplexing optical transmission. In: IEEE PROCAMS (2005)