

Snapshot Difference Imaging using Correlation Time-of-Flight Sensors

CLARA CALLENBERG, Institute for Computer Science II, University of Bonn

FELIX HEIDE, Stanford University

GORDON WETZSTEIN, Stanford University

MATTHIAS B. HULLIN, Institute for Computer Science II, University of Bonn




Fig. 1. We introduce a method that uses time-of-flight (ToF) imagers not for measuring scene depth, but rather as analog computational imagers that can directly measure difference values at a pixel level. We demonstrate this principle with a slightly modified ToF camera system (a), where simple reconfigurations in the optical path enable a wide range of imaging modalities. For instance, our system can directly sense temporal gradients (b), depth edges (c), direct-light-only images (d) and spatial gradients (Fig. 10) – each in a single exposure and without any additional decoding steps. We further show that the remarkable noise statistics of such imagers can be exploited to extract two color channels (here: red and blue) from a single snapshot taken under red *and* blue illumination (e). The top images in columns (b)–(e) are reference photographs of the respective scenes; the bottom ones visualize the output of our system.

Computational photography encompasses a diversity of imaging techniques, but one of the core operations performed by many of them is to compute image differences. An intuitive approach to computing such differences is to capture several images sequentially and then process them jointly. In this paper, we introduce a snapshot difference imaging approach that is directly implemented in the sensor hardware of emerging time-of-flight cameras. With a variety of examples, we demonstrate that the proposed snapshot difference imaging technique is useful for direct-global illumination separation, for direct imaging of spatial and temporal image gradients, for direct depth edge imaging, and more.

CCS Concepts: • **Computing methodologies** → **Image processing; Computational photography**;

Additional Key Words and Phrases: computational photography, ToF imaging

This work was supported by the German Research Foundation (HU-2273/2-1), the X-Rite Chair for Digital Material Appearance, a National Science Foundation CAREER award (IIS 1553333), a Terman Faculty Fellowship, and the KAUST Office of Sponsored Research through the Visual Computing Center CCF grant. We thank Nick Maggio for his help on early experiments.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Association for Computing Machinery.

0730-0301/2017/11-ART220 \$15.00

<https://doi.org/10.1145/3130800.3130885>

ACM Reference format:

Clara Callenberg, Felix Heide, Gordon Wetzstein, and Matthias B. Hullin. 2017. Snapshot Difference Imaging using Correlation Time-of-Flight Sensors. *ACM Trans. Graph.* 36, 6, Article 220 (November 2017), 10 pages. <https://doi.org/10.1145/3130800.3130885>

1 INTRODUCTION AND MOTIVATION

Over the last two decades, research in computational photography has been striving to overcome limitations of conventional imagers via a co-design of optics, sensors, algorithms, and illumination. Using this paradigm, unprecedented imaging modalities have been unlocked, such as direct-global light transport separation [Nayar et al. 2006], gradient imaging [Tumblin et al. 2005], temporal contrast imaging [Lichtsteiner et al. 2008], and direct depth edge imaging via multi-flash photography [Raskar et al. 2004]. A common operation for many of these techniques is to record two or more images and then compute the difference between them. Unfortunately, difference imaging is challenging for dynamic scenes, because motion creates misalignment between successively captured photographs which is in many cases difficult to mitigate in post-processing. In this paper, we explore a new approach to capturing difference images in a single exposure and generalize difference imaging to a variety of applications.

We propose to re-purpose time-of-flight (ToF) sensors to facilitate instantaneous difference imaging. The usual application for these sensors is depth imaging. In that context, they are operated

in conjunction with a periodically modulated light source. Light that has been reflected by the scene is demodulated by the sensor, reconstructing the shift in modulation phase and thereby the depth estimate per pixel. This functionality is achieved by a pixel architecture that employs two potential wells for photoelectrons to be stored in during the exposure, and that subtracts the charges accumulated in these two wells (Fig. 2). In other words, the core functionality of time-of-flight sensors is based on being able to take the difference of two incident signals before analog-to-digital (A/D) conversion.

Rather than computing scene depth, we demonstrate how ToF sensing technology can be used to conveniently implement a range of computational photography techniques, including direct-global separation, direct depth edge imaging, spatio-temporal gradient imaging, and more. The capabilities unlocked with snapshot difference imaging are particularly interesting for applications that require low-power, low-latency on-board processing with low bandwidth communication channels, such as internet-of-things devices. With this paper, we take first steps towards these directions.

Specifically, our contributions are the following:

- We introduce the concept of generalized difference imaging and develop an image formation and a noise model for this principle.
- We construct a prototype difference imager using a modified time-of-flight camera combined with multiple, spatio-temporally coded light sources.
- We evaluate the proposed imaging concept with several practical applications, including direct-global separation, direct depth edge as well as spatio-temporal gradient imaging.
- We demonstrate that two images can be recovered from a single difference image by exploiting characteristics of the proposed image formation model.

Overview of benefits and limitations. The proposed method has two primary benefits. First, capturing a difference image within a single exposure allows for faster time scales to be recorded than capturing two separate images and subtracting them digitally. Second, the noise properties of difference imaging before A/D conversion are shown to be favorable over digital subtraction post A/D conversion. A limitation of the proposed technique is that it relies on ToF sensors, which currently provide much lower resolution and signal quality than well-established CMOS or CCD sensors. Thus, comparing digital difference imaging with CMOS sensors and analog difference imaging with ToF sensors may not be beneficial for the latter approach. Yet, we demonstrate that our method yields superior noise performance for sensors with comparable characteristics.

2 RELATED WORK

Computational ToF imaging. This work presents a method for difference imaging by re-purposing two-bucket sensors usually used for depth imaging in lock-in ToF cameras. Lock-in time-of-flight sensors are a rapidly emerging sensor technology, with Microsoft’s Kinect for XBOX One as the highest-resolution sensor available on the market at 512×424 pixels [Bamji et al. 2015]. For technical details on lock-in ToF sensors we refer the reader to Lange et al. [1999] and Hansard et al. [2012]. A growing body of literature re-purposes these emerging sensors, in combination with computational methods, to

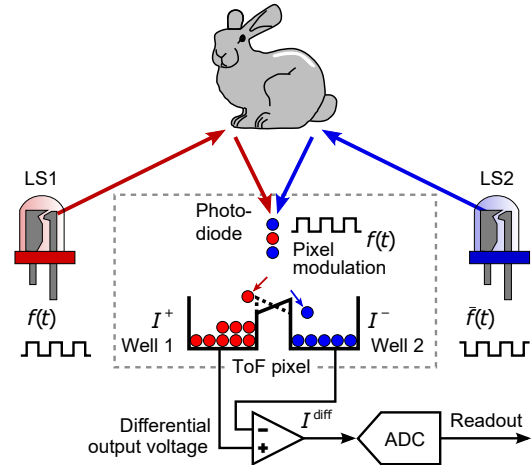


Fig. 2. Principle of operation of a time-of-flight (ToF) pixel. A light source is temporally modulated, its emitted light is reflected by the scene, and then demodulated in the pixel. To demodulate the coded illumination in the detector, two wells in each pixel collect charge carriers and an electric field oscillates at the demodulation frequency to direct incident photoelectrons into one or the other well. The sensor circuit measures the voltage difference before digitizing it by an analog-to-digital converter (ADC). Here, we illustrate how the difference between two modulated light sources can be directly measured with such a pixel architecture.

address a variety of challenging problems in imaging and vision. Kadambi et al. [2013] reduce multi-path interference by coding the modulation profiles, Heide et al. [2013] recover temporal profiles of light transport by measuring frequency sweeps, which allows for improved imaging in scattering media [Heide et al. 2014]. Recently, Tadano et al. [2015] designed depth-selective modulation functions enabling virtual-blue screening and selective back-scatter removal as applications.

Differential-pixel sensors. The proposed difference imaging method subtracts two signals before A/D conversion by “piggybacking” on two-bucket ToF sensor technology. Wang et al. [2012] have previously proposed a custom sensor design that also performs pre-ADC subtraction for the purpose of optoelectronic filtering and light field capture. Specifically, the authors use programmable gain operational amplifiers to compute the sum and difference of pixel pairs, which is then passed on to the A/D converter. In combination with local diffractive gratings as optics on every pixel, this allows to realize filtering with positive and negative filter coefficients. In contrast to a conventional sequential capture approach, these differential-pixel sensors offer reduced bandwidth [Wang and Molnar 2012], at the cost of spatial resolution. Compared to the proposed method, the optical filters are static and prohibit the flexible modes of operation demonstrated in this work. Changing from one difference-imaging task to another would require placing a different mosaicking pattern on the sensor.

Differential optical systems. Instead of this optoelectronic approach to difference imaging, one could also imagine cameras that perform the signal subtraction purely optically, plus a DC offset to ensure positivity. Building on Zomet and Nayar’s work [2006],

Koppal et al. [2013] present an optical design consisting of a micro-aperture mask in combination with lenslets allowing to design custom optical template filters for a variety of computer vision tasks. This approach may be adopted to design optical filters that perform spatial gradient or other filter differences in a single-shot, by designing difference filters with a DC offset to ensure non-negative coefficients. In theory, this approach would require variable high-resolution aperture patterns [Zomet and Nayar 2006]. Note also, that the proposed approach would be a natural choice for suppressing the DC in such a setup by relying on the adaptive background suppression of recent ToF sensors.

Event-based sensors. A further sensor design for differential measurements are event-based dynamic vision sensors [Gottardi et al. 2009; Lichtsteiner et al. 2008], which have been demonstrated for applications in vision and robotics, such as tracking [Kim et al. 2016] and simultaneous localization and mapping (SLAM) [Weikersdorfer et al. 2014]. Each pixel in such sensors asynchronously measures temporal intensity changes and generates spike events for temporal differences with magnitude above a given activation threshold. This event-driven behavior is achieved by augmenting each pixel with its self-timed switched-capacitor differencing circuit. By reading out pixels asynchronously, the core benefit of this sensor design is the low bandwidth requirement, enabling high frame rates and low power consumption [Gottardi et al. 2009; Lichtsteiner et al. 2008]. However, similar to the differential-pixel sensors, this comes at the cost of reduced spatial resolution, when compared to conventional sequential capture. While the proposed solution, based on ToF sensors, shares limitations in resolution, temporal differencing sensors do not support the very flexible modes of operation shown in this work. For example, capturing intensity or depth images requires solving ill-posed inverse problems [Kim et al. 2016].

Split-pixel HDR sensors. Backside-illuminated split-pixel architectures have become the dominant choice for high-dynamic-range (HDR) imaging in high-speed vision cameras [Willassen, Solhusvik, Johansson, Yaghamai, Rhodes, Manabe, Mao, Lin, Yang, Cellek, et al. Willassen et al.]. Single-shot HDR capture is essential for vision-based autonomous or assisted driving systems where reacting to fast moving objects over a wide dynamic range is critical [Solhusvik et al. 2013]. A variety of HDR sensor designs for high frame rates have been proposed in the past. Skimming HDR sensors perform partial resets (draining) of the accumulated charges during integration, allowing repeated partial integration with successively shorter resets [Darmont 2012]. The repeated integration can cause motion artefacts if partial saturation are reached quickly. Split-pixel architectures eliminate this issue by dividing each pixel into multiple buckets [Nayar and Branzoi 2003; Wan et al. 2012]. Multiple exposures are captured simultaneously by implementing different-sized photosensitive areas (OmniVision OV10640, OV10650). Given the emerge of split-pixel architectures as a key vision sensor technology, we believe that the proposed two-bucket difference imaging method may have broad applications even beyond the ones in this work.

3 IMAGING PRINCIPLE

Of all technologies that can be used for time-of-flight imaging, correlation sensors are the most widespread and affordable. This is also

the type of imager we are using in this work; throughout the paper, we use the term “time-of-flight (ToF) sensor” synonymously for this particular technology.

A pixel in a ToF sensor measures the amount of correlation between the incoming temporally varying photon flux $g_i(t)$ and a sensor modulation signal $f(t) \in [0, 1]$ that also varies in time [Lange et al. 1999]. Unlike regular CCD or CMOS sensors where electrical charges generated in a photodiode are collected in a potential well, ToF sensors feature two such wells per pixel (Fig. 2). The sensor modulation $f(t)$ decides whether a charge generated at time t will tend to end up in one well or the other. At the end of the integration phase, the difference between the two wells is read out and digitized. Neglecting quantization from the A/D conversion, this results in the digital value

$$\hat{I}^{\text{diff}} = \rho \cdot \eta \cdot (\hat{I}^+ - \hat{I}^-), \quad (1)$$

where ρ is the conversion factor from electron counts to digital units, and η denotes the so-called *demodulation contrast* [Schmidt 2011]. \hat{I}^+ and \hat{I}^- are the photoelectrons collected in the two wells over the integration period $[0, T]$:

$$\begin{bmatrix} \hat{I}^+ \\ \hat{I}^- \end{bmatrix} = \int_0^T \begin{bmatrix} f(t) \\ 1 - f(t) \end{bmatrix} g_i(t) dt \quad (2)$$

The incoming photon rate $g_i(t)$ is a function of the scene and the time-varying intensity $g(t)$ of an active light source that illuminates it. In ToF imaging, $f(t)$ and $g(t)$ are periodic functions of the same high frequency, typically 20–100 MHz, and the delay of light propagating from a source to the sensor results in a relative phase shift which is measured to recover depth. In snapshot difference imaging, we introduce two modifications to this scheme. Firstly, we reduce the modulation frequency to a point (1–5 MHz) where the propagation of light through near-range scenes can be assumed to be instantaneous and $f(t)$, typically generated by a digital circuit, only assumes the values 0 and 1. Secondly, we use two light sources, one (LS1) driven using the same function $f(t)$ and the other one (LS2) with its logical negation $\bar{f}(t)$. According to Eq. 2, the photocharges collected in \hat{I}^+ will record an image of the scene as illuminated by LS1, and LS2 will fill \hat{I}^- . The pixel value \hat{I}^{diff} thus measures the difference between two images taken under different illumination conditions, an insight that forms the foundation of this work.

3.1 Noise model

Time-of-flight imagers are complex photonic devices and as such suffer from noise of various different sources [Schmidt 2011]. The differential measurement scheme, and in particular the multi-tap measurement schemes typically used in ToF operation, cancel out many of the systematic errors introduced by the hardware. None of these measures, however, are capable of removing shot noise, which is the uncertainty that occurs during the counting of photoelectrons.

If \hat{I}^\pm are the *expected* electron counts for the two wells, the *actual* number of collected electrons I^\pm in any image recorded is a Poisson-distributed random variable with mean μ^\pm and variance $(\sigma^\pm)^2$ that are both identical to the respective expected value:

$$\mu^\pm = (\sigma^\pm)^2 = \hat{I}^\pm \quad (3)$$

As the difference of two independent random variables, the final




Fig. 3. Components of our system based on the Texas Instruments OPT8241-CDK module (bottom left). An external function generator (bottom right) produces a relatively low-frequency square wave, MOD, and its negated version, $\overline{\text{MOD}}$, that correspond to $f(t)$ and $\bar{f}(t)$ and modulate the sensor and two light sources (top right). The ENABLE signal is held high by the camera during integration.

pixel value is also a random variable, and it follows a Skellam distribution [Hwang et al. 2012; Skellam 1946]. Mean μ_{diff} and variance σ_{diff}^2 relate to the means μ_{\pm} and variances σ_{\pm}^2 of I_{\pm} as

$$\mu_{\text{diff}} = \eta(\mu_+ - \mu_-) = \eta(\hat{I}_+ - \hat{I}_-) \quad (4)$$

$$\sigma_{\text{diff}}^2 = \eta^2(\sigma_+^2 + \sigma_-^2) + \sigma_{\text{read}}^2 = \eta^2(\hat{I}_+ + \hat{I}_-) + \sigma_{\text{read}}^2 \quad (5)$$

where σ_{read}^2 models additional noise sources (assumed to be zero-mean), and the device constant $\eta \in [0, 1]$ is the imager's contrast [Schmidt 2011]. In terms of a matrix-vector product:

$$\begin{pmatrix} \mu_{\text{diff}} \\ \sigma_{\text{diff}}^2 - \sigma_{\text{read}}^2 \end{pmatrix} = \underbrace{\begin{pmatrix} \eta & -\eta \\ \eta^2 & \eta^2 \end{pmatrix}}_H \begin{pmatrix} \hat{I}_+ \\ \hat{I}_- \end{pmatrix}. \quad (6)$$

Note that the uncertainty σ_{diff}^2 of the measurement I_{diff} depends not primarily on the net difference value, but rather on the latent components \hat{I}^{\pm} that are subtracted from one another. Thus, even when there is zero signal ($\hat{I}^{\text{diff}} = 0$), the actual observation I_{diff} can suffer from significant noise. This is a principal property of difference imaging and holds for all sorts of settings, applying to in- as well as post-capture differencing techniques. We call the system-specific matrix H the *Skellam mixing matrix*. In Section 6, we show how it can be used to recover two color channels from a single exposure — an insight that, to our knowledge, has not been discovered before.

4 PROTOTYPE DIFFERENCE IMAGER

We constructed snapshot difference imagers based on two different time-of-flight sensing platforms. Our first prototype (not pictured) is a recreation of Heide et al.'s system [2013] that is based on the discontinued PMD Technologies CamBoard nano. The second prototype (Figs. 1 and 3) combines the Texas Instruments (TI) OPT8241-CDK evaluation module with external modulation and light sources in a




Fig. 4. Polarization difference imaging principle: the light of two identical light sources is modulated in (+) and out (−) of phase, respectively, with the sensor. The light sources' polarization directions are mutually perpendicular, one them being aligned in parallel with the analyzer filter in front of the sensor. Light that is reflected directly off the surface of the scene preserves the polarization while the light that scatters multiple times in the scene becomes depolarized. The sensor measures the difference image $I_{\text{diff}} = I_{\parallel}^{\text{direct}} + \frac{1}{2} I_{\parallel}^{\text{rot}} - \frac{1}{2} I_{\perp}^{\text{rot}} = I_{\parallel}^{\text{direct}}$.

similar way to the system described by Shrestha et al. [2016]. Both imagers have their infrared-pass filters removed so they can sense visible light. (For the TI sensor, we carefully polished the filter coating off using a Dremel 462 rotary tool.) To enable difference imaging, the external light sources are configured to operate in and out of phase, respectively. The systems are configured to capture at an exposure time of 2000 μs (PMD) and 1000 μs (TI) and 60 frames per second.

Each of our light sources carries three OSRAM OSLOM LEDs that are switched using the same signal. To implement the different imaging modalities described in Section 5, we equipped the light sources with LEDs of different colors (red, blue, infrared), placed them in different positions and equipped LEDs and camera with polarization filters as required for each particular purpose.

In Appendix A, to lower the entry barrier for the reader, we describe an alternative modification for stock OPT8241-CDK modules that does not require any custom hardware.

Measurement procedure. To reduce fixed pattern noise, a black frame was recorded before the data acquisition with our setup, and later subtracted from each measured frame. As a result, the difference image pixels obtain negative and positive values, depending on the charge balance of the two potential wells. This applies to all results shown in this paper.

5 APPLICATIONS AND EVALUATION

5.1 Polarization-based direct-global separation

Many computational imaging techniques are based on a decomposition of light transport into a *direct* component, i.e., light that has undergone exactly one scattering event between light source and camera, and a multiply scattered *indirect*, or *global*, component. Being able to investigate these components separately has been shown to enable, for instance, more realistic digital models for human skin [Ma et al. 2007] or more robust 3D scans of objects




Fig. 5. Polarization difference images of the two different scenes seen as RGB images in column (d). Column (a) shows the difference image I^{diff} containing only directly reflected light. The images in column (b) depict I^+ (parallel polarizers), the ones in column (c) I^- (crossed polarizers). The latter contains only light that has undergone multiple scattering events. The top dataset was captured with the PMD-based system, the bottom row with the TI sensor. The images have been rescaled for better visibility.

made of challenging materials [Chen et al. 2007; O’Toole et al. 2015]. While true separation into direct and indirect components is not within reach, researchers have used common observations about light transport to derive useful heuristics: indirect light tends to be spatially low-pass filtered [Nayar et al. 2006], it generally does not fulfill the epipolar constraint [O’Toole et al. 2012] and does not preserve polarization [Wolff 1990]. Here, we use our setup to exemplarily implement the third of these heuristics in the form of a single-shot polarization-difference imager, and demonstrate its capability to isolate directly reflected light.

According to Fig. 4, illuminating the scene using two light sources with crossed linear polarization and an analyzing filter on the camera, one can consider four different components of the image:

- $I_{||}^{\text{direct}}$ light that initially passed the polarization filter parallel to the analyzer and that was reflected directly in the scene, hence preserving the orientation of the polarization,
- $I_{||}^{\text{rot}}$ light that initially passed the polarization filter parallel to the analyzer and that was scattered multiple times in the scene, thus not preserving the orientation of the polarization,
- $I_{\perp}^{\text{direct}}$ light that initially passed the polarization filter perpendicular to the analyzer and that was reflected directly in the scene,
- I_{\perp}^{rot} light that initially passed the polarization filter perpendicular to the analyzer and that was scattered multiple times in the scene.

Assuming that multiple scattering in the scene completely depolarizes the light for both initial directions of polarization, the amount of the light reaching the camera after illumination with light polarized in parallel with the analyzer is the component that is in phase with $f(t)$:

$$I^+ = I_{||}^{\text{direct}} + \frac{1}{2}I_{||}^{\text{rot}},$$

while the amount of light reaching the sensor after illumination with light polarized perpendicularly to the analyzer is only the ratio of the incident light that has been depolarized and therefore passes the analyzing filter:

$$I^- = \frac{1}{2}I_{\perp}^{\text{rot}}.$$

The light sources are identical and we assume that the ratio between directly and indirectly scattered light is equal for both initial directions of polarization, hence $I_{\perp}^{\text{rot}} = I_{||}^{\text{rot}}$. Therefore, the difference image that the PMD sensor acquires in lock-in operation with the light sources is

$$I^{\text{diff}} = I^+ - I^- = I_{||}^{\text{direct}},$$

leaving an image containing only directly reflected light. Images of two sample scenes can be found in Fig. 5. Column (a) shows the difference image I^{diff} . Columns (b) and (c) show the parallel and perpendicular components $I_{||}^{\text{meas}}$ (b) and I_{\perp}^{meas} (c), respectively, where the latter visibly contains only indirectly scattered light.

5.2 Bipolar color matching functions

Being able to characterize and classify materials is important in many applications. Liu and Gu [2014] proposed to use discriminative illumination, or optimized pairs of spectro-angular illumination patterns to classify materials on a per-pixel basis. Here, we adopt the spectral aspect of this work, using our PMD setup to construct an active camera that discriminates between objects of red and blue reflectance in a single shot. By equipping L1 with red and L2 with blue LEDs, we obtain a *bipolar* color camera that measures a positive response for objects that are predominantly red, and a negative response for bluish objects. Fig. 6 shows an example measurement taken of the X-Rite ColorChecker, where the positive or negative response in the colored patches can clearly be seen. Patches that reflect red and blue to equal parts, like the greyscale, result in a response that is approximately zero.




Fig. 6. Left: Difference image of a color calibration chart, taken using PMD setup using alternating illumination of red and blue light. Reddish and bluish color patches obtain values on the opposite ends of the scale. Patches whose reflection spectra do not favor either red or blue obtain a value of approximately zero. Edges of some patches appear exaggerated due to partial shadowing of the light sources. Right: RGB scan taken of the color chart.

This example demonstrates the applicability of snapshot difference imaging to discriminative color imaging. The red–blue illumination we used is merely a simple example case. Multiple different light sources can be combined to obtain complex bipolar illumination spectra that can be individually adapted to a variety of applications. We also envision this capability to facilitate interesting and novel approaches to image segmentation and classification or to enable direct sensing of primary colors with bipolar matching functions, like the red primary $\bar{r}(\lambda)$ in the CIE 1931 RGB color space¹.

5.3 Depth edge and directional gradient imaging

It is often hard to deduce the structure and shape of three-dimensional objects from conventional photographs, as they may show low contrast between spatially distinct features of the object. Illuminating the object from two different angles, however, can unveil the depth structure of a scene and facilitate, for instance, a segmentation of the image. Similarly to Raskar et al. [2004], our setup can be used to produce directional gradient images of a scene, visualizing depth continuities as shown in Fig. 7. In this mode of operation, two identical light sources of opposite polarity are placed on opposite sides of the sensor. Whenever a depth discontinuity shadows one of the light sources, the resulting image displays positive or negative values. All other pixels obtain a value around zero. By varying the distance between the light sources, different edge widths are obtained. As the light source separation approaches the distance between scene and camera, the system records shading images like Fig. 8. Similarly to Woodham’s photometric stereo method [1980], they could be used to estimate the surface orientation of an object.

Comparison of single- and two-shot edge imaging. One of the key advantages of snapshot difference imaging is that it is immune to scene motion, whereas multi-shot techniques typically suffer from alignment issues when objects are rapidly moving. To illustrate this, we recorded two image sequences of a moving scene (bonsai tree shaken by wind) at the same frame rate of 60 frames per second. In Sequence 1, we used snapshot difference imaging with both light sources active; for Sequence 2, we alternated between LS1 and LS2, and digitally computed difference images between successive frames. As the results (provided as supplemental video) show, the single-shot difference images are significantly clearer with more consistent

¹https://commons.wikimedia.org/wiki/File:CIE1931_RGBCMF.svg




Fig. 7. Horizontal (a)/(b) and vertical (c) alignment of the light sources creates corresponding gradient images (TI camera). The distance between the light sources determines the width of the edges in the image: short distance (a) vs. larger distance (b).




Fig. 8. Directional difference image acquired with two light sources in wide vertical spacing (TI camera).

leaf shapes than the two-shot ones, and virtually free of ghosting artifacts. On the other hand, the single-shot images show a slight increase in fixed-pattern noise.

5.4 Spatio-temporal gradient imaging

A feature of difference imaging is its capability to extract essential information from heavy streams of image data. Here, we use our setup to implement cameras that selectively sense spatial or temporal changes in the input, opening use cases such as machine vision and data compression [Lichtsteiner et al. 2008].

Spatial gradient. We devise an optical setup that, in combination with a snapshot difference imager, performs edge detection in analog hardware. The key is to introduce a small spatial displacement between the images I^+ and I^- , so the net image becomes the difference between two shifted copies of the scene. While this could in principle also be done through a mechanical element in the optical path (similar to active image stabilization in photography), we only add optically passive components to our setup. In particular, we use oppositely polarized light sources as in Section 5.1. Instead of the analyzing filter on the lens, we place a birefringent crystal immediately on top of the sensor, behind the camera lens. We determined that an undoped YVO4 crystal, 1 mm thick and inclined by 20° with respect to the optical axis, causes a displacement between light of different polarization directions by about $15\ \mu\text{m}$, or one pixel of the TI sensor. For a polarization-preserving scene, this setup produces




Fig. 9. A birefringent crystal is placed between the sensor and the scene; the scene is illuminated using identical light sources that are polarized in perpendicular angles. One of them is operated in phase (+) and the other in opposite phase (-) with the sensor. In direct reflections which preserve the polarization, light from the light sources will be refracted in different angles inside the birefringent crystal and hence undergo a relative shift.




Fig. 10. Spatial gradient image of an aluminum resolution chart (TI camera). In horizontal direction, the markers show black and white edges. The markers reflect more light than the surrounding area, which is why the noise in the difference image is higher inside the markers than outside.

two identical images on the sensor area, displaced by one pixel and with opposite polarity. Uniform areas in the image cancel out in this difference image, while edges are detected as non-zero response (positive or negative depending on the direction). Figure 10 shows a gradient image of a planar aluminum resolution chart, recorded in a single shot using the TI setup.

Temporal gradient. We conclude with an example for our difference imaging approach that can even be used without active illumination. So far, we modulated the sensor with a high-frequency square wave at 50% duty cycle, which effectively made the sensor insensitive to ambient light. We now introduce a bias by choosing an asymmetric modulation pattern (Fig. 11). Light that arrives at the beginning of the exposure will now contribute to I^- , and light that arrives near the end will contribute more to I^+ . In doing so, we make the camera sense temporal changes of intensity, similar to the works by Wei et al. [?] and Hontani et al. [?]: pixels that receive more light during the second half of the exposure than during the first half appear as positive pixel values and vice versa. Fig. 12 shows an image thus captured of a rotating fan. From the temporal gradient image, the direction of rotation can be identified by the




Fig. 11. Example filters (modulation patterns) for use with our system. By default, the sensor is modulated at high frequency (top row). For the analog computation of temporal gradients, we use the 50/50 pattern with only one transition per exposure interval. This is implemented by modulating the sensor with a delayed version of the camera’s ENABLE signal (cf. Fig. 3).




Fig. 12. TI camera temporal gradient image (left) of a rotating fan (RGB image on the right). From the color gradient of blade edges the rotation direction of the fan is identifiable as clockwise.

black and white edges of the blades.

Another example is shown in Fig. 1(b), where white pellets are shown falling on the ground. The direction of motion is visible in the temporal gradient image: Pellets falling to the ground feature positive values (red) on the bottom end and negative values (blue) on the top end. Those that have bounced off the ground and fly back up (as seen in the right part of the image) have reversed shading. Pellets lying still on the ground are barely visible. Please see the real-time motion in the supplementary video.

The exposure time used for this method is 1 ms, thus transferring this method to a conventional camera would correspond to a required frame rate of 2000 fps.

5.5 Quantitative Noise Analysis

In contrast to conventional cameras, our snapshot difference imaging approach performs only one read-out operation in the process of obtaining a difference image, since the differencing operation is performed before the readout. Hence, assuming shot noise and read noise as the main contributions to the measurement uncertainty, a pre- and post-ADC difference image are expected to suffer from different noise levels:

$$\begin{aligned} \left(\sigma_{\text{diff}}^{\text{post}}\right)^2 &= \eta^2 \sigma_+^2 + \sigma_{\text{read}}^2 + \eta^2 \sigma_-^2 + \sigma_{\text{read}}^2, \\ \left(\sigma_{\text{diff}}^{\text{pre}}\right)^2 &= \eta^2 \sigma_+^2 + \eta^2 \sigma_-^2 + \sigma_{\text{read}}^2. \end{aligned} \quad (7)$$

To compare the relative performance of snapshot difference imaging with two-shot, post-capture difference imaging under otherwise identical conditions, we acquired three image sequences of a still scene, each N frames long: one sequence with both light sources activated (also shown in Fig. 5(a)), and two more with only LS1 or LS2 turned on, respectively (Figs. 5(b) and 5(c)). We then used the data acquired with separate light sources to compute another set of difference frames. As a measure for the signal quality of the




Fig. 13. Histograms of the absolute pixel variance in the analog and digital difference images of the scene depicted in the bottom row of Fig. 5.

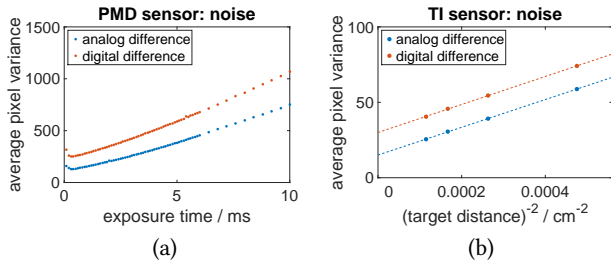


Fig. 14. Variance of each pixel, averaged over the whole difference image, for the PMD sensor (a) and the TI sensor (b). The values were obtained from 400 frames each. The lines in (b) are linear fits to the data.

difference frames, we computed for each pixel the variance across the N recorded frames, and plotted the values for all pixels in a histogram. As Fig. 13 illustrates for a case with $N = 100$ frames, the noise in the snapshot difference image is significantly lower than in the post-capture difference image.

Fig. 14 shows the pixel variance (averaged over all image pixels) in dependence of the intensity of the incident light. For the PMD camera, we varied the exposure time of otherwise identical shots using red and blue illumination. Since the TI sensor does not allow adjustment of the exposure time, we placed a white, homogeneous target in different distances from the camera and light sources in order to obtain different intensities. In both cases, the post-ADC difference image consistently shows higher noise. For low light intensities (short exposure times or large target distances, respectively), the shot noise of both pixel buckets tends to zero, so in Eq. 7, only the read noise terms remain and one has

$$\left(\sigma_{\text{diff}}^{\text{post}}\right)^2 = 2 \left(\sigma_{\text{diff}}^{\text{pre}}\right)^2.$$

The data depicted in Fig. 14 supports this expectation as the lowest measured variance values for the PMD sensor are

$$\left(\sigma_{\text{diff}}^{\text{post}}\right)_{\text{min}}^2 = 251.37, \quad \left(\sigma_{\text{diff}}^{\text{pre}}\right)_{\text{min}}^2 = 127.90$$

(ratio 1.97) and the extrapolated (via a linear fit) lowest variance values for the TI sensor are

$$\left(\sigma_{\text{diff}}^{\text{post}}\right)_{\text{min}}^2 = 29.99 \pm 0.30, \quad \left(\sigma_{\text{diff}}^{\text{pre}}\right)_{\text{min}}^2 = 15.05 \pm 0.22$$

(ratio 1.99 ± 0.37), which is in good agreement with a ratio of 2.²

²The variance values for pre- and post-ADC differencing in the PMD sensor diverge with higher exposure times while the difference between the two measurements appears constant in the TI sensor, which can be explained by dark current noise which is not explicitly accounted for in Eq. 7. Since dark current noise vanishes with the exposure time tending to zero, it does not affect the analysis of the PMD sensor, while its impact on the TI sensor analysis is unknown.

Since it is not possible to read out the pixel buckets separately, we furthermore cannot exclude the possibility that what we model as σ_{read}^2 is partly constituted of noise that




Fig. 15. Pixel variance in difference images of a color chart recorded with a TI ToF camera (a) and a Point Grey Flea3 camera (b). The variances were calculated from 400 and 137 frames, respectively.

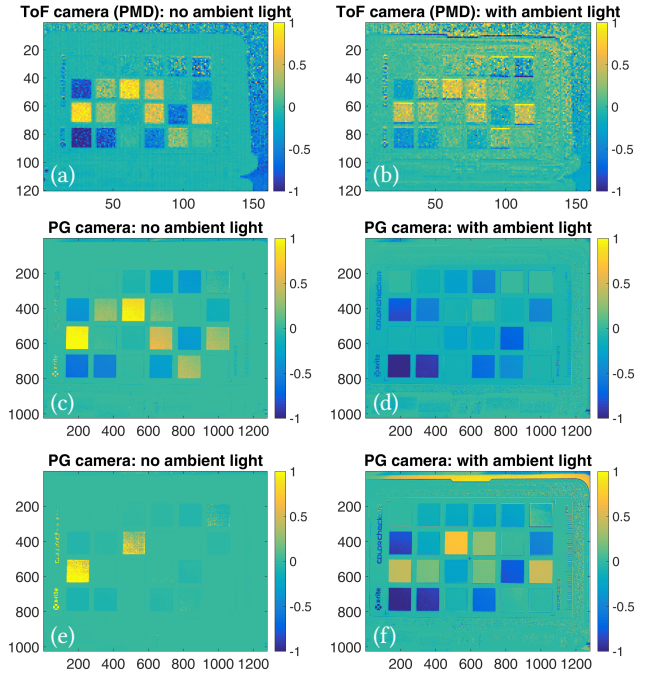


Fig. 16. Difference images without (left) and with (right) ambient illumination, acquired with a PMD ToF camera (top row) and a conventional PointGrey Flea3 camera with larger (center row) and smaller (bottom row) aperture setting.

In order to embed our setup into the context of existing camera hardware, we compared the noise level of the TI ToF camera to a conventional camera of type Point Grey Flea3 (“PG”). Fig. 15 shows the variance of each pixel in a series of difference images of a color chart taken with both cameras. For the PG camera, two sets of images have been recorded and subtracted digitally. In order to make this comparison as fair as possible, we used the same LED light sources and a shutter time of 1 ms for the TI camera and 0.5 ms for the two separate images taken with the PG camera. Evidently, the ToF camera shows higher noise and lower resolution than the conventional camera, which is expected due to the much longer development history of conventional image sensors compared to ToF sensors.

is introduced by the process of taking the difference voltage of both pixel buckets. This would reduce the factor between post- and pre-ADC-differencing in a setup with an otherwise identical conventional sensor to a value between 1 and 2.

The special architecture of ToF sensors, however, allows to neglect photoelectric charges in the pixel buckets that have been produced by light that is not correlated with the sensor modulation, i.e. ambient light. Figure 16 shows difference images of a color chart obtained with a PMD ToF sensor and a PG camera with and without ambient illumination of (several magnitudes) higher intensity than the active illumination. As shown in the upper row of Fig. 16, the additional light decreases contrast and increases noise in the difference image of the PMD sensor. With the PG camera, the sensor’s dynamic range does not allow to find a setting that captures both the target setup with and without ambient light correctly. Thus, in subfigure 16d, several patches of the color chart obtain pixel values of zero, because two saturated pixels have been subtracted from each other. If the camera parameters are adjusted such that the image with ambient illumination is properly exposed, the image without ambient light is too dark to show meaningful values (bottom row of Fig. 16). While in both cases the image quality for the properly exposed images is clearly better with the PG camera, the PMD camera shows higher variability and adaptability in terms of ambient illumination. Since the ability to suppress ambient illumination is unique to the ToF sensor type, we suppose that it has the potential of increasing the range of possible applications for the snapshot difference imaging approach in contrast to conventional camera setups in the future.

6 RECOVERING TWO IMAGES FROM A SINGLE DIFFERENCE IMAGE

In this section, we document an interesting side observation that falls directly out of the proposed image formation model for difference imaging. We can recover the two original images from a single difference image by exploiting the noise characteristics of both photon limited signals. According to Eq. 6, the noise in each pixel of a difference image is dependent on the amount of charges stored in the individual wells, rather than the resulting difference value. Therefore, we can calculate the separate values I^+ and I^- from the noise statistics (mean and variance of each pixel) of the difference image:

$$\begin{pmatrix} I^+ \\ I^- \end{pmatrix} = H^{-1} \begin{pmatrix} \mu \\ \sigma^2 \end{pmatrix}. \quad (8)$$

We propose three methods (M1–M3) to estimate these quantities:

M1. Analysis of a sequence of input image frames $I_{1..N}^{\vec{x}}$ taken under identical conditions:

$$\mu^{\vec{x}} = \frac{1}{N} \sum_{i=1}^N I_i^{\vec{x}} \quad \text{and} \quad (\sigma^2)^{\vec{x}} = \frac{1}{N-1} \sum_{i=1}^N (I_i^{\vec{x}} - \mu^{\vec{x}})^2,$$

where $I_i^{\vec{x}}$ denotes the pixel value at location $\vec{x} = (x, y)$ in the i^{th} frame.

M2. Patch-based analysis of a single pre-segmented image:

$$\mu^{\vec{x}} = \frac{1}{|P^{\vec{x}}|} \sum_{\vec{x}' \in P^{\vec{x}}} I^{\vec{x}'} \quad \text{and} \quad (\sigma^2)^{\vec{x}} = \frac{1}{|P^{\vec{x}}|-1} \sum_{\vec{x}' \in P^{\vec{x}}} (I^{\vec{x}'} - \mu^{\vec{x}})^2,$$

where $P^{\vec{x}}$ denotes the set of pixels belonging to the same image segment (patch) as pixel $\vec{x} = (x, y)$.

M3. Analysis of a single image using a bilateral filter:

$$\mu^{\vec{x}} = \frac{1}{\sum_{\vec{x}'} w} \sum_{\vec{x}'} I^{\vec{x}'} w \quad \text{and} \quad (\sigma^2)^{\vec{x}} = \frac{1}{\sum_{\vec{x}'} w} \sum_{\vec{x}'} (I^{\vec{x}'} - \mu^{\vec{x}})^2 w$$

with the bilateral weight w [Tomasi and Manduchi 1998]

$$w(\vec{x}', \vec{x}, I^{\vec{x}'}, I^{\vec{x}}) = e^{-(I^{\vec{x}'} - I^{\vec{x}})^2 / 2\sigma_{\text{range}}^2} e^{-(\vec{x}' - \vec{x})^2 / 2\sigma_{\text{domain}}^2}.$$

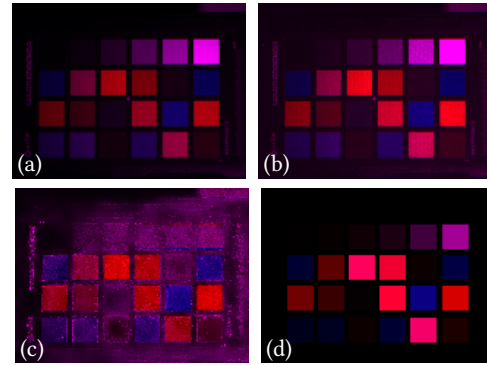


Fig. 17. Reconstruction of two color channels from a single exposure by exploiting photon statistics. (a) Ground-truth image combined from isolated measurements of red and blue illumination; (b) Reconstruction from 1000 difference images (Method 1); (c) Reconstruction from one single difference image (Method 3); (d) Reconstruction from one manually segmented image (Method 2). All images were acquired with the PMD sensor.

Fig. 17 shows the reconstructions of the individual blue and red channels from the difference image shown in Fig. 6a (for acquisition details, see Section 5.2), obtained using M1–M3 without the read noise term ($\sigma_{\text{read}}^2 := 0$). 100 dark frames were acquired, averaged and subtracted from the difference images before performing the reconstruction. M1, here using $N = 1000$ frames, delivers the best result. M2 and M3 sacrifice quality to separate the sources from one single difference image, which makes them suitable for fast moving target scenes. M2 yields the next-best reconstruction regarding color quality and particularly the gray scale (top row of patches), but it requires flat homogeneous image regions (pre-segmented by hand). M3 uses a bilateral filter to weight down dissimilar pixels when computing mean and variance. This reduces the overall estimated variance and introduces bias; nevertheless, this algorithm would be simple enough for real-time applications.

As another example of this differential recovery method, Fig. 18 shows reconstructions of the source images via M1 from a series of $N = 400$ difference images from the direct-global separation application described in Section 5.1 with and without correction for read noise, as well as ground truth. As expected, polarization-preserving reflections such as specular highlights appear in the “parallel” channel only, while sub-surface scattered (depolarized) light contributes to both channels.

Since we exploit basic properties of the Skellam distribution, this method also enables source separation for traditional difference imaging. As this approach essentially enables high-speed spatial multiplexed capture without spatial separation on the sensor, we envision a variety of applications beyond the two presented above.

7 DISCUSSION

In summary, we propose a new imaging system for direct recording of image differences in a snapshot. The proposed technique directly maps to the emerging technology of time-of-flight sensors and will therefore continue to benefit from the ongoing technological development in that area. The primary benefits of snapshot difference imaging include high video framerates that are only limited by the

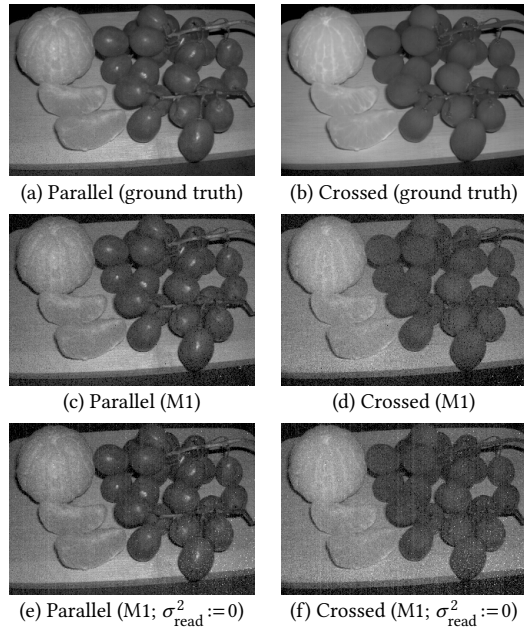


Fig. 18. Source images of the polarization difference image shown in Fig. 5a (bottom), reconstructed from the statistics of 400 difference images, with (c,d) and without (e,f) pre-calibrated read noise term. Subfigures (a) and (b) show ground truth images for comparison.

readout interface as well as lower noise and reduced alignment artifacts as compared to sequential, digital difference imaging. Finally, we devise an algorithm that is capable of encoding and extracting two different images from the mean and variance of a single photograph captured with the proposed method.

Limitations. Similar to range imaging, most of the demonstrated applications of snapshot difference imaging with time-of-flight sensors require active illumination. Joint coding and precise synchronization between the light sources and the sensor are required. The power of the employed light sources limits the range of distance within which the proposed method would function.

Future work. In the future, we would like to explore passive implementations of the proposed method, for example when using them with the natural flicker rates of existing indoor lighting. We would like to explore more sophisticated temporal coding strategies that may be able to separate direct and global illumination based on their temporal characteristics rather than their polarization properties. We would also like to explore spatio-temporal coding strategies that would allow the light sources to be used as temporally-coded projectors rather than isotropic emitters [O’Toole et al. 2015]. Finally, we would like to extend the application to mitigating multi-path interference for time-of-flight cameras and other tasks that may benefit from gradient cameras, such as Visual SLAM [Kim et al. 2016] and 3D scanning [Matsuda et al. 2015].

REFERENCES

Anonymous, . 2017. Snapshot Difference Imaging using Time-of-Flight Sensors. (2017). arXiv:arXiv:???, non-peer-reviewed prepublication by the authors.
 Bamji, C., P. O’Connor, T. Elkhatib, S. Mehta, B. Thompson, L. Prather, D. Snow, O. Akkaya, A. Daniel, A. Payne, T. Perry, M. Fenton, and V.-H. Chan. 2015. A 0.13 um

- CMOS System-on-Chip for a 512 x 424 Time-of-Flight Image Sensor With Multi-Frequency Photo-Demodulation up to 130 MHz and 2 GS/s ADC. *IEEE Journal of Solid-State Circuits* 50, 1 (2015), 303–319.
 Chen, T., H. P. Lensch, C. Fuchs, and H.-P. Seidel. 2007. Polarization and phase-shifting for 3D scanning of translucent objects. In *Proc. IEEE CVPR*. IEEE, 1–8.
 Darmont, A. 2012. High Dynamic Range Imaging: Sensors and Architectures. SPIE.
 Gottardi, M., N. Massari, and S. A. Jawed. 2009. A 100 μ W 128 \times 64 Pixels Contrast-Based Asynchronous Binary Vision Sensor for Sensor Networks Applications. *IEEE Journal of Solid-State Circuits* 44, 5 (2009), 1582–1592.
 Hansard, M., S. Lee, O. Choi, and R. Horaud. 2012. *Time of Flight Cameras: Principles, Methods, and Applications*. Springer.
 Heide, F., M. B. Hullin, J. Gregson, and W. Heidrich. 2013. Low-Budget Transient Imaging using Photonic Mixer Devices. *ACM Trans. Graph. (Proc. SIGGRAPH)* 32, 4 (2013), 45:1–45:10.
 Heide, F., L. Xiao, A. Kolb, M. B. Hullin, and W. Heidrich. 2014. Imaging in scattering media using correlation image sensors and sparse convolutional coding. *Opt. Express* 22, 21 (Oct 2014), 26338–26350. <https://doi.org/10.1364/OE.22.026338>
 Hwang, Y., J.-S. Kim, and I. S. Kweon. 2012. Difference-based image noise modeling using Skellam distribution. *IEEE Trans. PAMI* 34, 7 (2012), 1329–1341.
 Kadambi, A., R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar. 2013. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 167.
 Kim, H., S. Leutenegger, and A. J. Davison. 2016. Real-time 3D reconstruction and 6-DoF tracking with an event camera. In *European Conference on Computer Vision*. Springer, 349–364.
 Koppal, S. J., I. Gkioulekas, T. Young, H. Park, K. B. Crozier, G. L. Barrows, and T. E. Zickler. 2013. Toward Wide-Angle Microvision Sensors. *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (2013), 2982–2996.
 Lange, R., P. Seitz, A. Biber, and R. Schwarte. 1999. Time-of-flight range imaging with a custom solid state image sensor. In *Industrial Lasers and Inspection (EUROPTO Series)*. International Society for Optics and Photonics, 180–191.
 Lichtsteiner, P., C. Posch, and T. Delbruck. 2008. A 128 \times 128 120 dB 15 μ s latency asynchronous temporal contrast vision sensor. *Solid-State Circuits, IEEE Journal of* 43, 2 (2008), 566–576.
 Liu, C. and J. Gu. 2014. Discriminative illumination: Per-pixel classification of raw materials based on optimal projections of spectral BRDF. *IEEE Trans. PAMI* 36, 1 (2014), 86–98.
 Ma, W.-C., T. Hawkins, P. Peers, C.-F. Chabert, M. Weiss, and P. Debevec. 2007. Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proc. EGSR*. Eurographics Association, 183–194.
 Matsuda, N., O. Cossairt, and M. Gupta. 2015. MC3D: Motion Contrast 3D Scanning. In *Proc. ICCP*.
 Nayar, S. K. and V. Branzoi. 2003. Adaptive Dynamic Range Imaging: Optical Control of Pixel Exposures Over Space and Time. In *JCVV*. 1168–1175.
 Nayar, S. K., G. Krishnan, M. D. Grossberg, and R. Raskar. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph. (Proc. SIGGRAPH)* 25, 3 (2006), 935–944.
 O’Toole, M., S. Achar, S. G. Narasimhan, and K. N. Kutulakos. 2015. Homogeneous codes for energy-efficient illumination and imaging. *ACM Trans. Graph. (Proc. SIGGRAPH)* 34, 4 (2015), 35.
 O’Toole, M., R. Raskar, and K. N. Kutulakos. 2012. Primal-dual coding to probe light transport. *ACM Trans. Graph. (Proc. SIGGRAPH)* 31, 4 (2012), 39.
 Raskar, R., K.-H. Tan, R. Feris, J. Yu, and M. Turk. 2004. Non-photorealistic Camera: Depth Edge Detection and Stylized Rendering Using Multi-flash Imaging. *ACM Trans. Graph. (Proc. SIGGRAPH)* 23, 3 (2004), 679–688.
 Schmidt, M. 2011. *Analysis, modeling and dynamic optimization of 3D time-of-flight imaging systems*. Ph.D. Dissertation, University of Heidelberg.
 Shrestha, S., F. Heide, W. Heidrich, and G. Wetzstein. 2016. Computational Imaging with Multi-Camera Time-of-Flight Systems. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2016).
 Skellam, J. G. 1946. The Frequency Distribution of the Difference Between Two Poisson Variates Belonging to Different Populations. *Journal of the Royal Statistical Society* 109, 3 (1946), 296–296. <http://www.jstor.org/stable/2981372>
 Solhusvik, J., J. Kuang, Z. Lin, S. Manabe, J. Lyu, H. Rhodes, et al. 2013. A comparison of high dynamic range CIS technologies for automotive applications. In *Proc. 2013 Int. Image Sensor Workshop (IISW)*.
 Tadanu, R., A. K. Pediredla, and A. Veeraraghavan. 2015. Depth Selective Camera: A Direct, On-Chip, Programmable Technique for Depth Selectivity in Photography. In *2015 IEEE International Conference on Computer Vision (ICCV)*. IEEE. <https://doi.org/10.1109/iccv.2015.410>
 Tomasi, C. and R. Manduchi. 1998. Bilateral filtering for gray and color images. In *Proc. IEEE ICCV*. IEEE, 839–846.
 Tumblin, J., A. Agrawal, and R. Raskar. 2005. Why I want a gradient camera. In *Proc. IEEE CVPR*, Vol. 1. 103–110.
 Wan, G., X. Li, G. Agranov, M. Levoy, and M. Horowitz. 2012. CMOS image sensors

- with multi-bucket pixels for computational photography. *Solid-State Circuits, IEEE Journal of* 47, 4 (2012), 1031–1042.
- Wang, A. and A. Molnar. 2012. A Light-Field Image Sensor in 180 nm CMOS. *IEEE Journal of Solid-State Circuits* 47, 1 (jan 2012), 257–271. <https://doi.org/10.1109/jssc.2011.2164669>
- Wang, A., S. Sivaramakrishnan, and A. Molnar. 2012. A 180nm cmos image sensor with on-chip optoelectronic image compression. In *Custom Integrated Circuits Conference (CICC), 2012 IEEE*. IEEE, 1–4.
- Weikersdorfer, D., D. B. Adrian, D. Cremers, and J. Conradt. 2014. Event-based 3D SLAM with a depth-augmented dynamic vision sensor. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 359–364.
- Willassen, T., J. Solhusvik, R. Johansson, S. Yaghami, H. Rhodes, S. Manabe, D. Mao, Z. Lin, D. Yang, O. Cellek, et al. A 1280× 1080 4.2 μm split-diode pixel HDR sensor in 110nm BSI CMOS process.
- Wolff, L. B. 1990. Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 11 (1990), 1059–1071.
- Woodham, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical engineering* 19, 1 (1980), 191139–191139.
- Zomet, A. and S. K. Nayar. 2006. Lensless imaging with a controllable aperture. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, Vol. 1. IEEE, 339–346.

A MINIMAL SETUP

A basic implementation of a snapshot difference imager for infrared only can be obtained by connecting a TI OPT8241-CDK sensor board with two of its original infrared light sources via an extended ribbon cable (Fig.19). The modulation polarity is reversed (wires #10 and #12 swapped) and wire #16 is cut for the second light source. Note, however, that the OPT8241-CDK board by itself cannot generate modulation signals below 10 MHz. Since the camera captures groups of four phase-shifted sub-frames (in 90° steps), the effective frame rate is reduced by a factor of 4 compared to our system.




Fig. 19. Modified ribbon cable with (top to bottom) connectors for sensor board, LS1 and LS2.