Real-time computational photon-counting LiDAR

Matthew Edgar
Steven Johnson
David Phillips
Miles Padgett

Real-time computational photon-counting LiDAR

Matthew Edgar,[a,*] Steven Johnson,[a] David Phillips,[b] and Miles Padgett[a]
[a]University of Glasgow, Scottish Universities Physics Alliance, Optics Group, School of Physics and Astronomy, Glasgow, United Kingdom
[b]University of Exeter, School of Physics, Exeter, United Kingdom

Abstract. The availability of compact, low-cost, and high-speed MEMS-based spatial light modulators has generated widespread interest in alternative sampling strategies for imaging systems utilizing single-pixel detectors. The development of compressed sensing schemes for real-time computational imaging may have promising commercial applications for high-performance detectors, where the availability of focal plane arrays is expensive or otherwise limited. We discuss the research and development of a prototype light detection and ranging (LiDAR) system via direct time of flight, which utilizes a single high-sensitivity photon-counting detector and fast-timing electronics to recover millimeter accuracy three-dimensional images in real time. The development of low-cost real time computational LiDAR systems could have importance for applications in security, defense, and autonomous vehicles. © *2017 Society of Photo-Optical Instrumentation Engineers (SPIE)* [DOI: 10.1117/1.OE.57.3.031304]

Keywords: computational imaging; photon counting; light detection and ranging; time of flight; three-dimensional imaging.

Paper 171572SS received Oct. 3, 2017; accepted for publication Dec. 8, 2017; published online Dec. 29, 2017.

1 Introduction

Time-of-flight (TOF) three-dimensional (3-D) imaging has importance for many applications, including robotics, security, and autonomous vehicles. Long-range light detection and ranging (LiDAR) systems typically rely on illuminating a scene with a short-pulsed laser, and temporally correlating the reflected light intensity with the outgoing pulse to determine the time of flight to different points in the scene. A depth map can then be accurately determined by multiplying the time of flight with the speed of light c and a reflectivity map obtained from the amplitude of the reflected intensity.

The overall performance of a time-of-flight imaging system is determined by the choice of detector, the laser, the scanning hardware/strategy, the time-tagging electronics, and the image reconstruction algorithm. The most common approach for obtaining transverse spatial resolution from a LiDAR system uses a single time-resolving detector for determining the time of flight, one pixel at a time, by scanning the illumination and/or the detector across a field of view. Importantly, for long-range imaging applications where the reflected light intensity is low, the operation of this system demands the use of high-sensitivity single-photon counting devices and it is often desirable to be sensitive to nonvisible wavelengths.[1–3] When the detection is shot-noise limited, it is necessary to integrate the detections from many backscattered photons to improve statistical confidence. The inherent dead time of single-photon sensitive detectors, typically on the order of 10s of nanoseconds, prohibits the retrieval of short-range timing information from a single illumination pulse, leading to the use of high-repetition-rate pulsed lasers. Conventional opto-mechanical scanning technologies, for instance, a pair of galvanometer mirrors, typically have scan rates in the kHz regime. When used in combination with a raster scanning strategy, the acquisition time scales linearly with image resolution.

To acquire transverse spatial resolution, a desirable option might be to flood illuminate the scene and use an array of time-resolving detectors, such as Geiger mode single-photon avalanche detector (SPAD) arrays,[4–6] to acquire time of flight for each pixel simultaneously. Currently, these arrays have resolutions of a few thousand pixels, and can have limitations such as pixel cross talk, excessive dark noise, readout noise, and stringent requirements on readout clocks per pixel. This maturing technology is already indicating a promising future for compact LiDAR systems, but an interesting question arises when considering the fill factor, size, and overall quantum efficiency of these devices compared with a single large-area time-resolved detector when used in combination with scanning hardware/strategies and image reconstruction algorithms. It is the latter which forms the motivation for this work.

In the last few years, there have been a number of interesting demonstrations for recovering 3-D images using a single-pixel detector. One scheme scans a scene, pixel by pixel, line by line, using a pulsed illumination source and measures the reflected light using an avalanche photodiode (APD), where the first detected photon is used to recover depth and reflectivity via time of flight.[7] An alternative method makes use of structured pulsed illumination and a SPAD,[8–10] but importantly, unlike raster scanning techniques is able to benefit from reduced acquisition times by employing compressed sensing principles, which takes advantage of the sparsity in natural scenes.[11–13]

In this work, we demonstrate a 3-D imaging system, capable of recovering millimetric depth precision, using a single large-area photomultiplier, a high-speed spatial light modulator, and a short-pulsed near-infrared laser. The use of a simplified compressed sensing strategy, used in conjunction with fast streaming of the time-resolved intensity histograms and an efficient image reconstruction algorithm, overcomes the time constraints of previous works and permits continuous real time image reconstruction at up to 3 frames/s.

*Address all correspondence to: Matthew Edgar, E-mail: matthew.edgar@glasgow.ac.uk
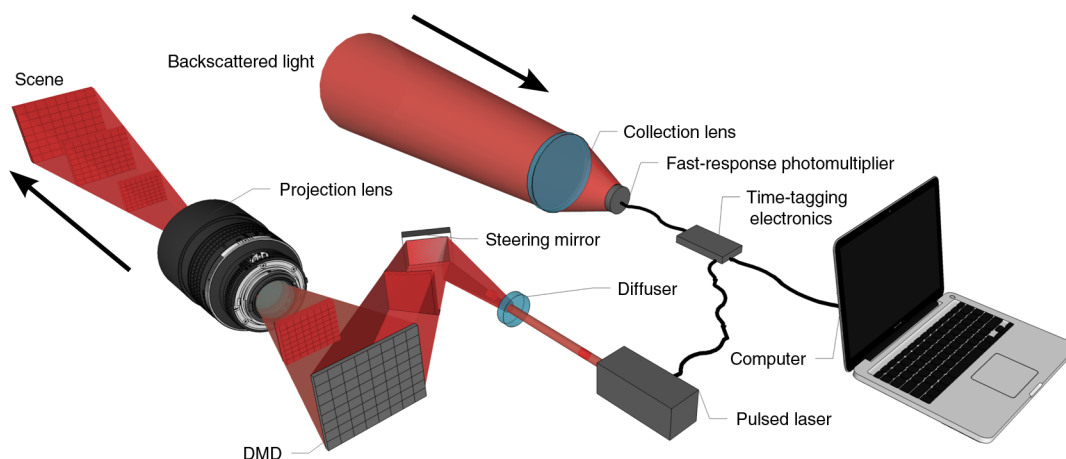
Fig. 1 A pulsed laser uniformly illuminates a DMD, used to provide structured illumination onto a scene, and the backscattered light is collected onto a high-speed photodiode. The measured light intensities are used in a computer algorithm to reconstruct a 3-D image.

2 Experimental Setup

Our single-pixel 3-D imaging system, shown in Fig. 1, consists of a 100-fs pulsed laser with a repetition rate of 100 MHz (Toptica FemtoFErb) with a center wavelength of 780 nm and a high-speed digital-micromirror-device (DMD) (Texas Instruments Discovery 4100 supplied by Vialux, model V-7001) to provide time-varying structured illumination. A fast-response Geiger-mode photomultiplier tube (PicoQuant PMA192) is used in conjunction with a 50-mm diameter collection lens to detect the first backscattered photon resulting from each pattern interacting with the scene. Time-correlated single-photon counting (TCSPC) electronics (a customized Horiba DeltaHub) record the time of arrival for each detected "event" relative to the synchronization output provided by the laser. Our choice of TCSPC electronics enables continuous streaming of up to 20,000 histograms/s, with 25-ps bin widths, allowing for real-time signal processing suitable for this investigation.

3 Sensing Strategy

For any single-pixel imaging system, which involves performing a series of measurements in time, the choice of scanning basis is an important consideration to optimize the image quality while minimizing the acquisition time. The finite modulation rate of the DMD implies a fundamental trade-off between acquisition time and image resolution. However, we note that most natural images exhibit similar characteristics, for example, sparsity in their spatial frequencies, the principle underpinning image compression techniques, which opens the possibility to compressively sense at the acquisition stage, such that fewer measurements than the number of image pixels are required.

A variety of approaches exist for performing sub-Nyquist sampling of a scene using a single-pixel camera. The most common approach in other work involves acquiring a series of measurements using a basis which is spatially incoherent with the scene properties, and utilizing a nonlinear optimization algorithm to recover an image whose spatial properties satisfy the prior knowledge of the scene characteristics.[11,12] In general, computational overhead associated with the recovery for such problems exceeds the acquisition time, which can prohibit real-time imaging applications. Nonetheless, a variety of postprocess techniques have been proposed

and demonstrated to recover video-rate performance on 2-D[13,14] and 3-D imaging systems.[8–10]

In a recent time-of-flight imaging demonstration[15] using an integrating photodetector for detection, a subset of the naturally constructed Hadamard matrices[16] was used to provide structured pulsed illumination. The corresponding backscattered intensities measured by the APD are subsequently used within a iterative reconstruction algorithm, having negligible computational overhead compared with the total acquisition time. This work was led by an understanding that for an iterative reconstruction algorithm the patterns associated with the largest signals have the most influence on final reconstructed image quality.[17] It was noted that practical demonstrations of the so-called "evolutionary" compressed sensing strategy has limitations for scenes exhibiting dynamic behavior when the spatial properties change significantly from frame to frame and within each frame.

A subsequent investigation[18] using a similar system but employing an integrating APD for detection made use of a visible camera to obtain 2-D images of the scene, which were used to determine the optimal subset of the basis for sampling 3-D information. Here, the 2-D image stream, provided by the visible camera, at a rate of 30 Hz, is used to continuously "simulate" the expected intensity signals for the entire basis, which are used to order the basis according to their magnitudes, from which an arbitrary subset of the basis is chosen for display on the DMD. This is an example of a stimulated or adaptive compressed sensing strategy.

In this work, the performance of the 3-D computational LiDAR system, described in Sec. 2, was evaluated by scanning the spatial properties of the scene using either the complete Hadamard basis or a subset of the Hadamard basis. The performance of the system using alternative compressive sensing strategies forms the basis of a follow-up publication.

The Hadamard basis has several important properties. Hadamard matrices, H, are an orthogonal basis set, which are $N \times N$ square matrices having values of ± 1 which satisfy

$$H^T H = H^2 = NI, \tag{1}$$

where I is the identity matrix. Therefore, the inverse of a Hadamard matrix is

$$H^{-1} = \frac{1}{N} H. \tag{2}$$

Each row from H can be reshaped and rescaled into a unique 2-D binary pattern for display on the DMD, such that an M pixel image can be fully sampled after performing all N measurements.

An important feature of naturally constructed Hadamard matrices is that in each row, the number of "pixels" in the scene that is sampled is exactly 50%. We note that the natural Hadamard patterns once reshaped in two-dimensions for display on the DMD exhibit spatial properties ranging from coarse to fine resolution. Ordering and displaying the basis on the DMD according to the scale of the pattern resolution can be advantageous if the scene exhibits predominantly lower spatial properties, since the number of patterns required to enable image reconstruction can be significantly reduced.[19] In this work, the Hadamard basis is ordered in this way, such that the spatial frequencies of the scene are effectively measured from lowest to highest, allowing the number of measurements acquired in any frame to be chosen arbitrarily by the camera operator.

It is worth noting that since the DMD is a binary modulator, and only one detector is being used in this demonstration, we are able to measure a signal corresponding to the image intensities overlapping with values in the Hadamard patterns of $+1$ and would need to estimate the intensity measured for -1 values. Instead we choose to display each pattern, consisting of $+1$ values, followed immediately by a pattern containing the -1 values (its negative). From these two measurements we obtain a differential measurement by subtracting one from the other, similar to performing heterodyne detection, which has the benefit of removing external noise arising from fluctuations in the source brightness or ambient intensity changes.

Figure 2 shows an illustration of sample Hadamard pattern pairs displayed on the DMD for a period of 0.5 μs (during which the 780-nm laser pulses $\sim 5 \times 10^4$ times) and corresponding to the histograms obtained from the TCSPC electronics. For each pair of patterns, we obtain a differential histogram, which is subsequently used to perform image reconstruction. Typically, each histogram contains ~ 1000 photon detections.

4 Real-Time Image Reconstruction

For any single-pixel imaging system, the time-varying intensity signal $S_i(t)$ associated with each projected pattern P_i (length N, pixel number) is directly proportional to the overlap between each pattern and the scene reflectivity O (also length N), given by

$$S_i(t) = P_i \cdot O, \tag{3}$$

where the temporal resolution of the discretely sampled intensity signal is determined by the bin width of the TCSPC electronics, and can be expressed in units of distance $S_i(z)$ since $z = ct/2$, where c is the speed of light and t is the time of flight. Furthermore, we perform calibration of the measured intensity to account for attenuation of the light with increasing distance from the detector, such that

$$S_{i,\mathrm{corr}}(z) = S_i(z) \left(\frac{z_0 + z}{z_0} \right)^2, \tag{4}$$

where z_0 corresponds to the range of the first depth interval in the histogram.

Choosing to sample a scene using complete or incomplete Hadamard basis allows a simple iterative reconstruction algorithm to be employed, having negligible computational time on a standard computer processing unit. Thus after each pattern and corresponding histogram, an estimate of the 3-D image cube, I_{3D}, can be reconstructed according to
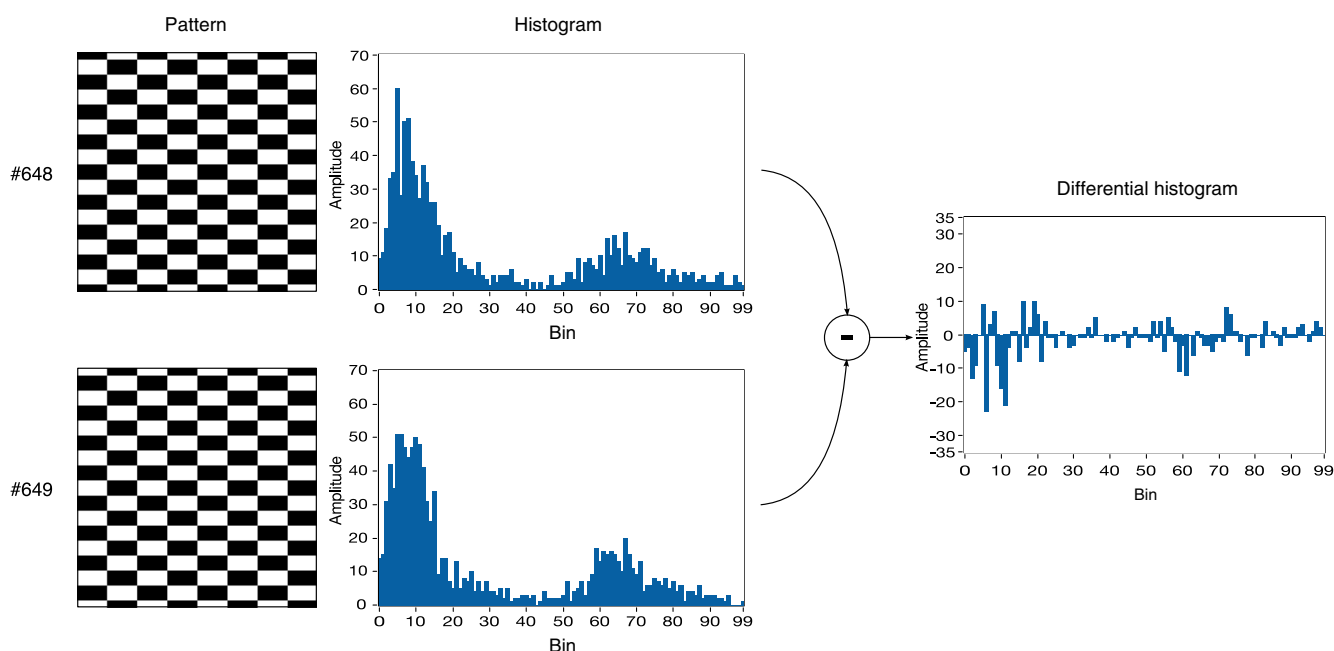
Fig. 2 Sample of patterns displayed on the DMD, constructed from reshaped rows of the Hadamard matrix, and the corresponding measured histograms resulting from accumulating single-photon detections of $\sim 5 \times 10^4$ laser pulses. It is the differential histogram that is used in subsequent 3-D image reconstruction.

$$I_{3D}(z) = \sum_{i=1}^{M} S_{i,\text{corr}}(z) \cdot P_i, \tag{5}$$

where M is the number of patterns used for sampling. The resulting image cube is a discretized array of 2-D images equally spaced in depth by 3.75 mm, determined by the 25-ps temporal resolution of the TCSPC electronics.

Operating the system in real time typically demands short integration times, thus the detected photon flux is often low. We can, however, apply "spatiotemporal" smoothing to the reconstructed image cube to help overcome the effects of Poissonian noise. In this work, we convolve the image cube with a normalized 3-D smoothing kernel $\kappa(x, y, z)$ having dimensions (3, 3, 5) given by

$$\kappa(x, y, 1) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0.033 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\kappa(x, y, 2) = \begin{bmatrix} 0 & 0.033 & 0 \\ 0.033 & 0.066 & 0.033 \\ 0 & 0.033 & 0 \end{bmatrix}$$

$$\kappa(x, y, 3) = \begin{bmatrix} 0.033 & 0.066 & 0.033 \\ 0.066 & 0.133 & 0.066 \\ 0.033 & 0.066 & 0.033 \end{bmatrix}$$

$$\kappa(x, y, 4) = \begin{bmatrix} 0 & 0.033 & 0 \\ 0.033 & 0.066 & 0.033 \\ 0 & 0.033 & 0 \end{bmatrix}$$

$$\kappa(x, y, 5) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0.033 & 0 \\ 0 & 0 & 0 \end{bmatrix} \tag{6}$$

At each pixel of the denoised 3-D reconstruction, we obtain the intensity by integrating along the z-dimension. For non-transparent objects, we can make the assumption that there is only one reflective surface at depth z for each transverse pixel location (x, y). Therefore, we estimate the depth z for each each pixel (x, y) to be that where the measured intensity is maximum.

It has been shown that the ranging precision can be enhanced beyond the limits of the system hardware by performing a variety of techniques on the measured intensity signals, such as parametric deconvolution,[20] curve fitting or interpolation.[15] However, as a result of technical difficulties experienced in this investigation, we are unable to report on such improvements to the ranging at this time.
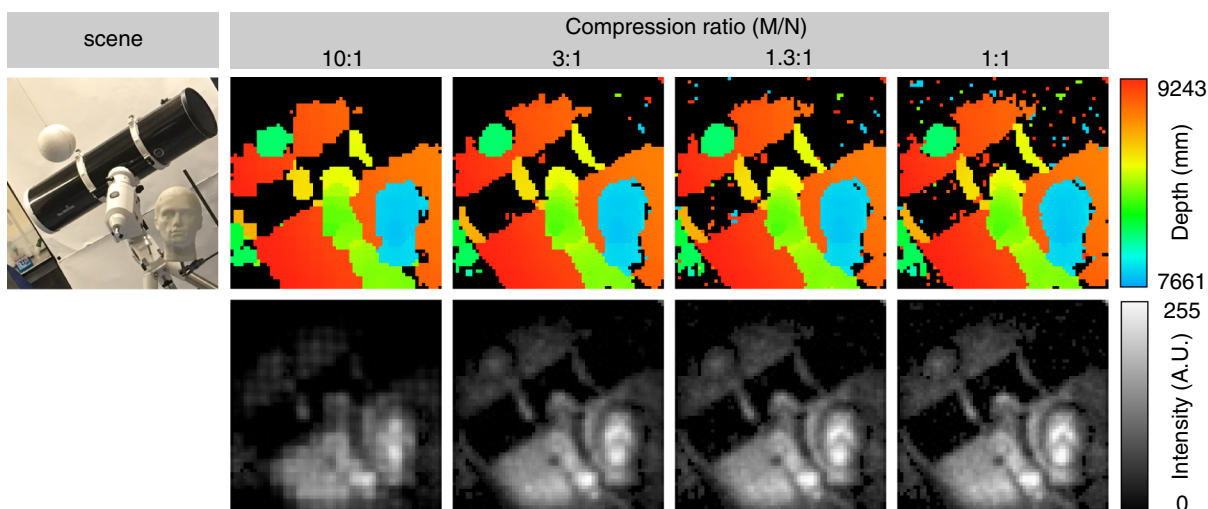
5 Results

The signal processing and image reconstruction algorithm that were used to obtain the results presented here were designed and implemented entirely in the LabVIEW software development environment. An outline of the structure of the LabVIEW software used, summarizing the main functions, is provided in Appendix A.

In one experiment, a scene consisting of a mannequin head located at a distance of ~3 m from the camera system was imaged using an 85-mm focal length lens, and signal processing was performed in real time at a rate of 3 fps. To achieve this, the first 333 Hadamard patterns of 1024 from the ordered sequence were displayed one after another continuously on the DMD. A histogramming range of 100 bins on the TCSPC electronics was chosen for subsequent image processing as described in Sec. 4. The modulation rate of the DMD was 2052 Hz, that is each pattern was displayed for ~0.5 ms, in which time the laser pulses 50,000 times resulting in ~1000 "events" added to each histogram over a temporal range of 2.5 ns. As the repetition rate of the laser was 100 MHz, the maximum imaging range equates to ~6 m. A sample of the reconstructed intensity and depth
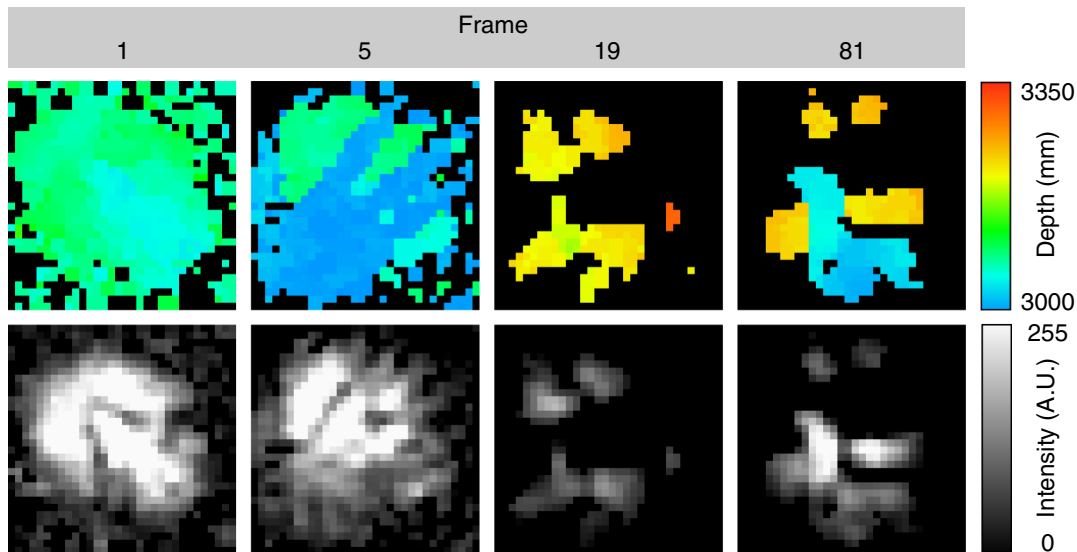
Fig. 3 Sample of video frames showing the reconstructed intensity and depth for a dynamically changing scene. The system operates in real time at a rate of 3 Hz. Each frame is reconstructed from 333 patterns equivalent to a compression ratio of 3:1. Throughout the video, the changing scene contains a mannequin head, a real waving hand, a real head wearing safety-glasses, and a "thumb's-up" gesture. The video is available online in the supplementary materials. These frames have been taken from Video 1 (Video 1, MOV, 0.13 MB [URL: http://dx.doi.org/10.1117/1.OE.57.3.031304.1]).

Fig. 4 A photograph of the scene (left) consisting of a mannequin head, ball, telescope, and angled wall behind. The reconstructed intensity (middle) and depth map (right) are shown. Inspection of the results reveals that the black metallic telescope is a poor scatterer of the pulsed illumination; therefore, the number of backscattered photons is below the detection limits of the system.

maps is shown in Fig. 3. Throughout the acquisition, the scene is dynamically changing to obtain a polystyrene mannequin head, a waving hand, a head wearing safety-glasses, and a hand giving a "thumb's-up" gesture. The video can be found in the online supplementary materials.

In a separate experiment, the source was replaced with a 6-ps pulsed supercontinuum laser with a repetition rate of 60 MHz (Fianium femtopower) spectrally filtered 635 ± 3.5 nm, and a 50-mm focal length lens was used. The repetition rate of this source provides a maximum imaging range of ~10 m. A scene consisting of a telescope, a suspended polystyrene ball, a polystyrene mannequin head, and an angled wall was arranged at a range of 8 m from the camera. All 4096 Hadamard patterns were used to reconstruct images at 64×64 pixel resolution. A histogramming range of 500

bins on the TCSPC electronics was chosen for this experiment and the modulation rate of the DMD was reduced to 50 Hz, i.e., each pattern was displayed for ~20 ms, in which time the laser pulses 1.2×10^6 times. For this scene, the mean number of detected "events" for each pattern was 2.3×10^4, added to the corresponding histogram over a temporal range of 12.5 ns. The reconstructed intensity and depth maps are shown in Fig. 4, where the number of patterns used for reconstruction is $M = 410, 1365, 2730,$ and 4096, equivalent to a compression ratio of 10:1, 3:1, 1.3:1, and 1:1, respectively.

We can observe that for this scene, the intensity, and in particular the depth map, is well reconstructed at a compression ratio of 3:1, which enables the system to operate 3× faster than when sampling with the complete basis.
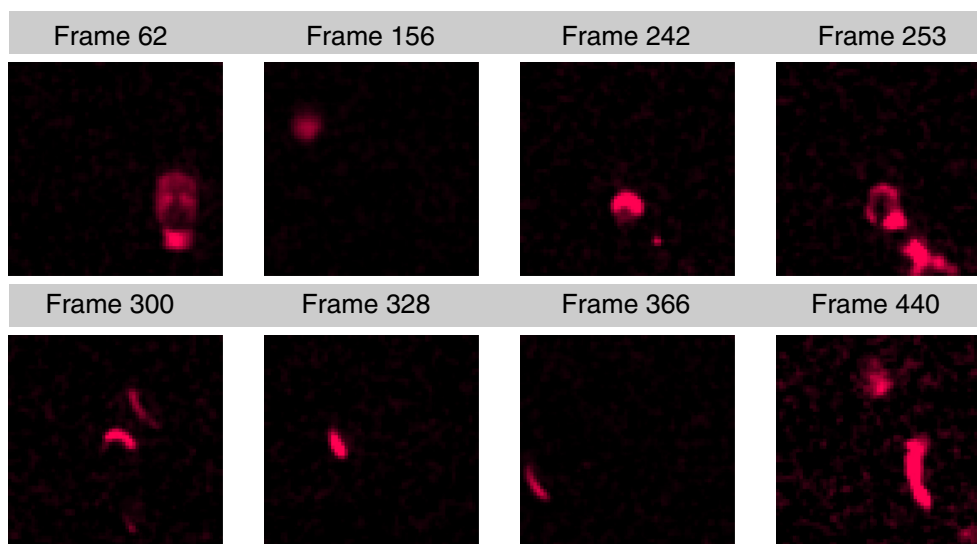
Fig. 5 A sample of intensity frames (or planes) from the 3-D image cube, which have been taken from Video 2 demonstrating that the system effectively captures "40 billion fps high-speed video" of a light pulse as it propagates throughout the scene. Inspection of the frames we observe the mannequin head appearing in frame 62, the suspended ball in frame 156, the telescope and tripod in frames 242, 300, 328, and 366 and finally the angled rear wall in frame 440 (Video 2, MOV, 1.3 MB [URL: http://dx.doi.org/10.1117/1.OE.57.3.031304.2]).

We note that this method of reconstructing the 3-D image cube provides access to the pulsed illumination as it is propagated throughout the scene, equivalent to having a high-speed camera capturing at 40 billion frames/s. A sample of frames from this high-speed video is shown in Fig. 5 depicting the sheet of light as it reflects from the scene at different times.

6 Conclusions and Future Work

We have demonstrated a photon-counting computational LiDAR system, utilizing short-pulsed structured illumination and a fast-response photomultiplier tube, for reconstructing 3-D scenes at up to 3 fps. We demonstrated results obtained when applying a very simple compressive sensing strategy using a subset of the Hadamard basis ordered according to spatial frequencies. Importantly, this method allows for intensity and depth image reconstruction in less time than the acquisition which enables continuous real-time operation. A variety of alternative sampling strategies can also be employed, which may yield improved performance, for instance, the use of microscanning,[21] spatially varying sampling strategies[22] or utilizing deep learning,[23] all of which are the subjects of ongoing work and follow-up publications. It is worth pointing out that this work has been carried out in a laboratory with control of the ambient lighting. Performing similar demonstrations in other scenarios, such as outdoors, would require consideration for the sensitivity of the photon-counting detector technology used (e.g., PMT). For instance, a narrowband spectral filter matching the output of the illumination source would be necessary on the detection channel to reduce the background count rate. Moreoever, operating at 780 nm the background count rate from solar activity may add excessive noise and significantly reduce overall performance and image quality from this system. Importantly, the operational spectra of the DMD (400 to 2500 nm) make them good candidates for extending these techniques to longer wavelengths, such as the short-wave infrared region, where there are several operational advantages, such as higher power eye-safe lasers, enhanced visibility at long range due to reduced atmospheric scattering, and significantly reduced solar background.

Appendix: Outline of Software Implementation

The LabVIEW program developed for this experimental demonstration can be summarized by the following key operations:

1. Initializing and setup
   - Construct and order the Hadamard patterns for displaying on DMD.
   - Initialize the DMD and mode of operation (in this demonstration master mode).
   - Initialize the Horiba DeltaHub histogramming electronics and mode of operation (in this demonstration slave mode).
   - Upload complete basis to the available RAM onboard the DMD controller.

2. Pattern display on DMD

   - An independent while loop running continuously pending a user input to stop.
   - User defined control of the pattern subset size and/or pattern order for subsequent display on the DMD (in this demonstration, the size was 333 patterns and the order was unchanged).
   - For each frame, the associated series of patterns were displayed at 2052 Hz, uninterrupted by computer communication.

3. Histogram data acquisition

   - An independent while loop running continuously pending a user input to stop.
   - Real-time streaming up to 20 kHz of histograms containing a maximum of 512 time bins, each of 25-ps bin width.
   - Each histogram streamed to the computer is synchronized using a TTL output from the DMD.

4. Signal processing and image reconstruction

   - An independent while loop running continuously pending a user input to stop.
   - Preprocessing of histogram data to first ensure no data have been lost or corrupted.
   - Recover differential intensity signals from pairs of patterns displayed on DMD (corresponding to positive/negative values in the Hadamard matrix).
   - Calibrate the differential intensity signals to account for attenuation and speed of light propagation, thereby yielding depth of flight.
   - Reconstruct 3-D image cube via iterative sum of known patterns weighted by histogram bin intensity.
   - Perform spatiotemporal smoothing to 3-D image cube and apply thresholding to reduce noise
   - Recover intensity image by integrating each pixel (x, y) in the image cube along the z-(depth) dimension.
   - Recover depth image by finding bin where intensity is maximum.

References

1. A. McCarthy et al., "Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting," *Appl. Opt.* **48**(32), 6241–6251 (2009).

2. A. McCarthy et al., "Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector," *Opt. Express* **21**, 22098–22113 (2013).
3. A. M. Pawlikowska et al., "Single-photon three-dimensional imaging at up to 10 kilometers range," *Opt. Express* **25**, 11919–11931 (2017).
4. A. Giudice et al., "High-rate photon counting and picosecond timing with silicon-SPAD based compact detector modules," *J. Mod. Opt.* **54**(2–3), 225–237 (2007).
5. M. Entwistle et al., "Geiger-mode APD camera system for single-photon 3D LADAR imaging," *Proc. SPIE* **8375**, 83750D (2012).
6. N. Krstajić et al., "0.5 billion events per second time correlated single photon counting using CMOS SPAD arrays," *Opt. Lett.* **40**(18), 4305–4308 (2015).
7. A. Kirmani et al., "First-photon imaging," *Science* **343**(6166), 58–61 (2014).
8. G. A. Howland, P. B. Dixon, and J. C. Howell, "Photon-counting compressive sensing laser radar for 3D imaging," *Appl. Opt.* **50**(31), 5917–5920 (2011).
9. J. C. Howell, "Compressive depth map acquisition using a single photon-counting detector: parametric signal processing meets sparsity," in *Proc. of the 2012 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR 2012)*, pp. 96–102, IEEE Computer Society, Washington, D.C. (2012).
10. G. A. Howland et al., "Photon counting compressive depth mapping," *Opt. Express* **21**(20), 23822–23837 (2013).
11. D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory* **52**(4), 1289–1306 (2006).
12. E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory* **52**(2), 489–509 (2006).
13. M. F. Duarte et al., "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**(2), 83–91 (2008).
14. A. C. Sankaranarayanan, C. Studer, and R. G. Baraniuk, "CS-MUVI: video compressive sensing for spatial-multiplexing cameras," in *IEEE Int. Conf. on Computational Photography (ICCP '12)*, pp. 1–10 (2012).
15. M.-J. Sun et al., "Single-pixel three-dimensional imaging with time-based depth resolution," *Nat. Commun.* **7**, 12010 (2016).
16. W. K. Pratt, J. Kane, and H. C. Andrews, "Hadamard transform image coding," *Proc. IEEE* **57**(1), 58–68 (1969).
17. N. Radwell et al., "Single-pixel infrared and visible microscope," *Optica* **1**(5), 285–289 (2014).
18. M. P. Edgar et al., "Real-time 3D video utilizing a compressed sensing time-of-flight single-pixel camera," *Proc. SPIE* **9922**, 99221B (2016).
19. M.-J. Sun et al., "A Russian dolls ordering of the Hadamard basis for compressive single-pixel imaging," *Sci. Rep.* **7**(1), 3464 (2017).
20. A. Kirmani et al., "Exploiting sparsity in time-of-flight range acquisition using a single time-resolved sensor," *Opt. Express* **19**(22), 21485–21507 (2011).
21. M.-J. Sun et al., "Improving the signal-to-noise ratio of single-pixel imaging using digital microscanning," *Opt. Express* **24**(10), 10476–10485 (2016).
22. D. B. Phillips et al., "Adaptive foveated single-pixel imaging with dynamic supersampling," *Sci. Adv.* **3**(4), e1601782 (2017).
23. C. Higham et al., "Deep learning for real-time single-pixel video," (submitted) (2017).

Matthew Edgar is a research associate in the Optics Group at the University of Glasgow. He received his BSc and PhD degrees in physics and astronomy in 2007 and 2011, respectively. He started his research career in the Institute for Gravitational Research at Glasgow, developing advanced interferometric techniques to enhance the sensitivity of long-baseline gravitational wave detectors. Since joining the Optics Group at Glasgow, he has been investigating the use of camera technology to perform fundamental tests of quantum mechanics and more recently has been developing low-cost computational imaging systems for applications in methane imaging and 3-D imaging.

Steven Johnson is a research associate at the University of Glasgow working in the Optics Group. He received a MSci in theoretical physics and a PhD in ultracold atoms from the University of Birmingham in 2012. He has previously worked in industry specialising in CMOS imaging sensors. His research interests involve applications of temporally resolved computational imaging for measuring very fast phenomena.

David Phillips is a Royal Academy of Engineering research fellow in the Physics Department, University of Exeter. Since graduating with a physics degree in 2004, he has spent a few years working as a systems engineer in industry, a few months working in science policy within UK Parliament, and the rest of the time studying nanophysics and optics in academia. His current research interests are focused on computational imaging in scattering environments.

Miles Padgett holds the Kelvin Chair of Natural Philosophy at the University of Glasgow. He is fascinated by light both classical and quantum–specifically light's momentum. In 2001, he was elected to fellowship of the Royal Society of Edinburgh and in 2014, the Royal Society, the UK's National Academy. In 2009, with Les Allen, he won the IoP Young Medal, in 2014 the RSE Kelvin Medal in 2015 the Science of Light Prize from the EPS and in 2017 the Max Born award of the OSA.