SCIENTIFIC REPORTS

OPEN

Sub-picosecond photon-efficient 3D imaging using single-photon sensors

Received: 17 July 2018
Accepted: 30 October 2018
Published online: 07 December 2018

Felix Heide, Steven Diamond, David B. Lindell & Gordon Wetzstein

Active 3D imaging systems have broad applications across disciplines, including biological imaging, remote sensing and robotics. Applications in these domains require fast acquisition times, high timing accuracy, and high detection sensitivity. Single-photon avalanche diodes (SPADs) have emerged as one of the most promising detector technologies to achieve all of these requirements. However, these detectors are plagued by measurement distortions known as pileup, which fundamentally limit their precision. In this work, we develop a probabilistic image formation model that accurately models pileup. We devise inverse methods to efficiently and robustly estimate scene depth and reflectance from recorded photon counts using the proposed model along with statistical priors. With this algorithm, we not only demonstrate improvements to timing accuracy by more than an order of magnitude compared to the state-of-the-art, but our approach is also the first to facilitate sub-picosecond-accurate, photon-efficient 3D imaging in practical scenarios where widely-varying photon counts are observed.

Active 3D imaging has broad applications across disciplines, including remote sensing, non-line-of-sight imaging, autonomous driving, robotics, and microscopic imaging of biological samples[1–11]. Key requirements for most of these applications include high timing accuracy, fast acquisition rates, wide dynamic operating ranges, and high detection sensitivity. In particular, remote sensing and automotive applications[1–4,6] demand acquisition ranges from <1 meter to kilometers, while non-line-of-sight imaging[7–9], for example, relies on the information encoded by the few returning photons of multiply scattered indirect light, in addition to the directly reflected light. To facilitate these applications, ultra-sensitive detectors have been developed that allow for individual photons returning from a pulsed illumination source to be recorded. One of the most sensitive time-resolved detector technologies that can be produced in the CMOS process to date are single-photon avalanche diodes (SPADs), which have been rapidly established as a core detector technology for 3D imaging[2,12–20].

SPADs are reverse-biased photodiodes operated above their breakdown voltage, i.e. in Geiger mode[12]. Photons incident on the active surface of a SPAD can trigger an electron avalanche, which is subsequently time stamped. By repeatedly time stamping photons returning from a synchronously pulsed illumination source, typically operating at MHz rates, a histogram of photon counts over time can be accumulated. The histogram approximates the intensity of the returning light pulse and characterizes the propagation delay, enabling distance, reflectance and 3D geometry to be recovered. While a SPAD can be operated in a free-running mode[21], which allows photon events to be detected at all arrival times simultaneously, some modes of operation enable gated detection wherein only photons within a specified time window between pulses are detected[22]. In this gated mode[23,24], SPADs achieve high accuracy and robustness to ambient light by electronically gating out all but a narrow temporal slice[22]. However, gated acquisition requires a sequential sweep of the temporal slice over the full target range, which restricts it to applications that do not require fast acquisition rates. SPAD detectors operated in free-running mode do not suffer from this limitation, but whether these sensors are used for 3D imaging, microscopy, non-line-of-sight imaging, or any other application, they are subject to a fundamental phenomenon which severely limits accuracy: pileup distortions[25,26].

Pileup is a limitation inherent to the SPAD detector's working principle. After every triggered electron avalanche, the detector needs to be quenched before it is able to detect further photon arrival events. During this dead time, which is typically in the order of tens to hundreds of ns, the detector is inactive. Thus, earlier photons of a single laser pulse are more likely to trigger an avalanche while later ones are more likely to fall into the dead time, thus being ignored. This causes severe measurement skew, which is known as pileup. Pileup cannot be

Stanford University, Department of Electrical Engineering, Stanford, USA. Correspondence and requests for materials should be addressed to F.H. (email: fheide@stanford.edu) or G.W. (email: gordon.wetzstein@stanford.edu)

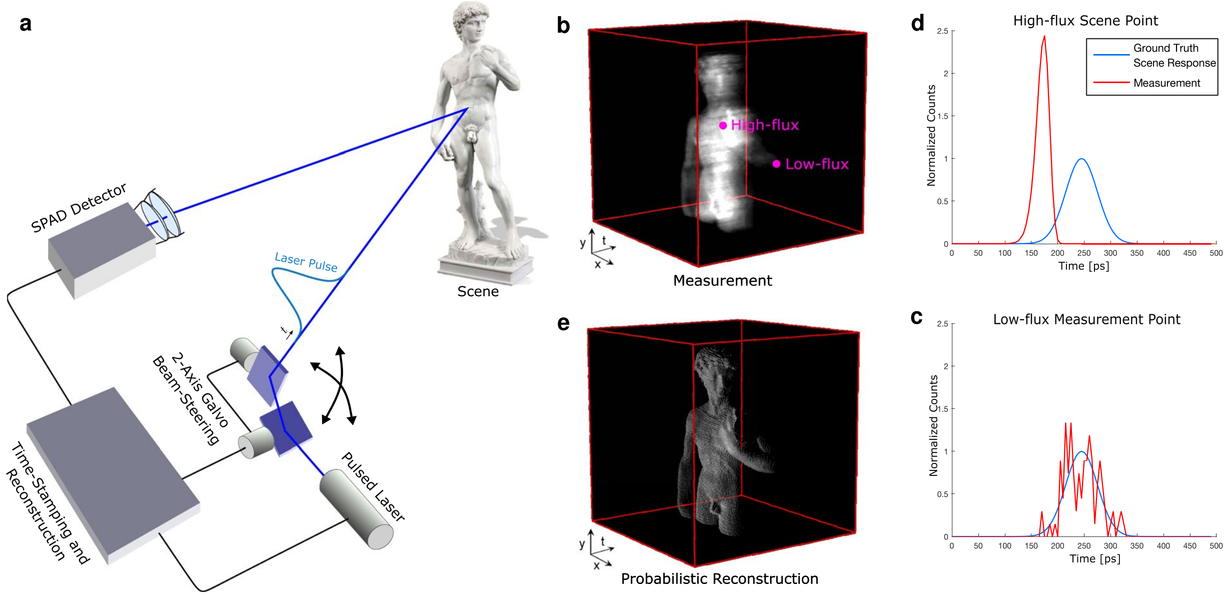

Figure 1. Sub-picosecond 3D Imaging Framework. (**a**) A collimated, pulsed laser illuminates the scene at a single point. The laser is laterally scanned using a 2-axis mirror galvanometer. Timing and control electronics time-stamp each detected photon arrival relative to the last emitted pulse and accumulate these events in a histogram of spatio-temporal photon counts **(b)**. This histogram is processed to estimate both reflectivity and depth information **(c)**. Two points are highlighted, one corresponding to high-flux **(d)** and the other to low-flux **(e)** measurements. Whereas the latter are noisy, high-flux measurements suffer from pileup distortion which introduce a significant bias for the depth estimation of conventional algorithms. The proposed estimation method accurately models both of these scenarios, allowing for reflectance information and travel time to be estimated with sub-picosecond accuracy from severely distorted measurements.

corrected in hardware and results in round-trip time-stamping errors that are orders of magnitudes larger than the timing precision of the detector. Note that the overall instrument dead time is the maximum of the detector dead time and the dead time of the time-to-digital conversion (TDC). However, TDC circuits with dead times of less than a nanosecond are available, such as the PicoQuandt TimeHarp Nano, which is one to two orders of magnitudes lower than typical detector dead times. Although this work is applicable to both TDC or detector-induced pileup, we consider detector-limited systems in the following.

To avoid pileup distortion, active imaging systems can be operated in a low-flux regime, where the average number of incident photons per measurement is extremely low (i.e. $\ll 1$). State-of-the-art depth and reflectivity estimation techniques have demonstrated successful operation in these challenging conditions using probabilistic image formation models and statistical priors in the inverse methods[18,27,28]. However, many 3D imaging applications, for example in robotics, biological imaging, or automotive sensing, operate in environments where both objects returning a low number of signal photons and objects reflecting higher numbers of photons are essential for decision making. The large variance in acquired photon counts of objects at different depths or resulting from varying reflectivity of different objects makes it crucial for a practical 3D imaging technique to operate robustly in both low-flux and high-flux conditions simultaneously (see Fig. 1).

In this work, we introduce a new estimation algorithm that overcomes existing limitations of active 3D imaging systems with free-running SPADs. Specifically, we demonstrate sub-picosecond timing accuracy for pulsed light sources with a full width at half maximum (FWHM) wider than 50 ps. The proposed method improves the accuracy of existing depth and albedo estimation algorithms by more than an order of magnitude across a wide dynamic range, from low-flux to high-flux measurements. These benefits are enabled by an image formation model and corresponding inverse method that lift assumptions of low-flux models[18,27] to broad operating conditions that may be distorted by pileup. To this end, we introduce a probabilistic image formation model that includes pileup and we derive efficient inverse methods for the depth and albedo estimation problem with this model. The inverse methods exploit statistical priors of both depth and albedo. Unlike previous work, our reconstruction framework jointly estimates all unknown parameters, overcoming algorithmic limitations that restrict the timing precision of existing methods. The proposed methods not only facilitate highly accurate and fast 3D imaging but they also open up a new operating regime of photon-efficient 3D imaging in conditions with drastically-varying photon counts. Concurrently with our work, several other methods addressing the pileup problem for SPADs have recently been published[29,30].

Results

Imaging setup. The experimental acquisition setup is illustrated in Fig. 1. The light source is a pulsed, collimated laser that is scanned horizontally and vertically using a 2-axis mirror galvanometer. Light scattered back from the scene is focused on the SPAD (Micro Photon Devices PD-100-CTD) using a microscope objective. A PicoQuant PicoHarp 300 module is used for time stamping, histogram accumulation, and synchronization of laser and detector. We use a pulse rate of 5 MHz, resulting in a maximum path length of 60 m without considering the instrumentation dead time. The dead time of the SPAD detector is 75 ns. Due to the fact that this dead time is multiple orders of magnitude larger than the desired timing accuracy, the pileup distortions of measured histograms can only be ignored in the ultra-low-flux regime with $<10^{-2}$ average photons detected per pulse.

Observation model. We discretize the scene into a grid of $m \times n$ points at distances $\boldsymbol{z} \in \mathbf{R}^{m\times n}$ and denote the reflectance of these points as $\boldsymbol{\alpha} \in [0, 1]^{m\times n}$. Each point is probed with N pulses.

We first describe the observation model for a single point (i, j) by modeling the temporal shape of the laser pulse g as a mixture of Gaussians

$$g(t) = \sum_{k=1}^{K} a_k \exp\left(-\frac{(t-b_k)^2}{c_k^2}\right), \tag{1}$$

where the parameters $a_i, b_i, c_i \in \mathbf{R}$ are calibrated in a pre-processing step (cf. Methods). Given this parametric impulse model, the photon detections on the SPAD detector can be modeled as an inhomogeneous Poisson process with rate function

$$r_{ij}(t, \alpha_{ij}, z_{ij}, s) = \mu\alpha_{ij}\, g(t - 2z_{ij}/c) + s, \tag{2}$$

where c is the speed of light, μ is the SPAD photon detection probability, and s models background detections from ambient light and dark count[18].

The free-running SPAD records photon detections over a time interval $[0, T]$ discretized into uniform bins of size Δ, where T is the inverse of the laser repetition rate. Let h_{ijk} denote the accumulated counts in bin k over the N pulses and $\lambda_{ijk}(\alpha_{ij}, z_{ij}, s) = \int_{k\Delta}^{(k+1)\Delta} r_{ij}(t, \alpha_{ij}, z_{ij}, s)dt$ the Poisson rate for the number of detections in each bin k. The proposed probabilistic reconstruction method infers the latent variables $\boldsymbol{z}, \boldsymbol{\alpha}$ using maximum-a-posteriori (MAP) estimation, which relies on the probability $P(\mathbf{h}_{ij}|\boldsymbol{\lambda}_{ij})$ of observing a histogram $\mathbf{h}_{ij} \in \mathbf{Z}^T$ given means $\boldsymbol{\lambda}_{ij}(\alpha_{ij}, z_{ij}, s) \in \mathbf{R}^T$ and priors on depth and albedo.

Prior work on 3D imaging using SPAD detectors focuses on the low-flux regime[18,27,31,32] in which the expected number of photon detections per pulse is significantly smaller than one. In the low-flux regime we may neglect dead time and approximate the observations h_{ijk} as being conditionally independent across bins k, with corresponding Poisson probability mass function $P(h_{ijk}|\boldsymbol{\lambda}_{ij})$. For medium and high-flux measurements the conditional independence approximation breaks down because dead time ensures that at most one photon detection is recorded per pulse[25]. The probability of an aggregated count in a single histogram bin, taking dead time into account, is given by the multinomial distribution

$$P(\mathbf{h}_{ij}|\boldsymbol{\lambda}_{ij}) = \frac{N!\exp(-\mathbf{1}^T\boldsymbol{\lambda}_{ij})^{N-\mathbf{1}^T\mathbf{h}_{ij}}}{h_{ij1}!\cdots h_{ijT}!(N-\mathbf{1}^T\mathbf{h}_{ij})!}\prod_{\ell=1}^{T}\left(\exp\left(-\sum_{k=1}^{\ell-1}\lambda_{ijk}\right) - \exp\left(-\sum_{k=1}^{\ell}\lambda_{ijk}\right)\right)^{h_{ij\ell}}, \tag{3}$$

where $\mathbf{1}$ is the vector of all ones. We refer to the Supplemental Methods for a detailed derivation of this probabilistic model. The observation model holds across the full measurement range, from low-flux to high-flux regimes.

Reconstruction algorithm. After scanning a scene, temporal histograms $\mathbf{h}_{ij}$ are available for each point (i, j). To reconstruct scene reflectance and depth from these histograms, we find the MAP estimate of $\boldsymbol{\alpha}$ and $\boldsymbol{z}$, along with the ambient term s, using the observation model and the prior assumption that the gradients of reflectivity and depth maps are sparse. Inspired by prior work on depth and reflectivity estimation[18,27,28], we place transverse anisotropic total variation (TV) priors directly on $\boldsymbol{\alpha}$ and $\boldsymbol{z}$. The MAP estimate is given by solving the optimization problem

$$\underset{\boldsymbol{\alpha},\,\mathbf{z},\,s}{\text{minimize}} \quad \sum_{i=1}^{m}\sum_{j=1}^{n} -\log P(\mathbf{h}_{ij}|\boldsymbol{\lambda}_{ij}(\alpha_{ij}, z_{ij}, s)) + \gamma_1\|\nabla\boldsymbol{\alpha}\|_1 + \gamma_2\|\nabla\mathbf{z}\|_1, \tag{4}$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{z}$ are the unknown variables, and ∇ is the gradient operator. To solve this problem, we develop a proximal algorithm for this non-convex optimization problem that decouples the sparsity-promoting prior terms from the likelihood term. Specifically, we minimize the joint objective by introducing slack variables $\mathbf{v}^{\alpha}$, $\mathbf{v}^{z}$, $\mathbf{v}^{s}$ for albedo, depth and ambient terms. We then optimize each unknown term in an alternating fashion (see Supplemental Material for details). As a result of this proximal optimization scheme, the log-likelihood minimization becomes separable in the measurement ij as

$$\underset{\boldsymbol{\alpha},\,\mathrm{z},\,s}{\text{minimize}} \quad -\log P(\mathbf{h}_{ij}\big|\boldsymbol{\lambda}_{ij}(\boldsymbol{\alpha}, \mathrm{z}, s)) + \frac{\xi}{2}(\mathbf{z}-\mathbf{v}^{\mathbf{z}})^2 + \frac{\xi}{2}(\boldsymbol{\alpha}-\mathbf{v}^{\alpha})^2 + \frac{\xi}{2}(s-\mathbf{v}^{s})^2, \tag{5}$$

which is an optimization problem over a non-linear tri-variate loss function and quadratic proximal closeness terms with scalar weight ξ. We solve the minimization problem with per-pixel parallel Newton Methods, as derived in the Supplemental Methods.

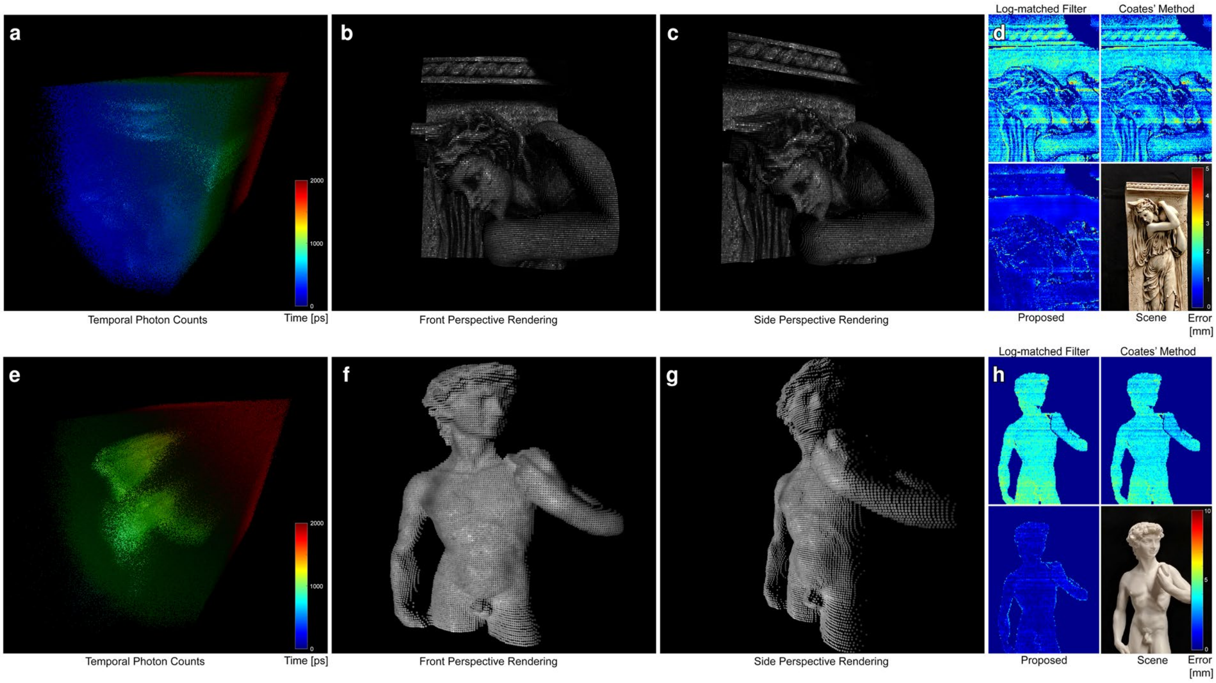

Figure 2. Experimental reconstructions. A recorded spatio-temporal distribution of photon counts **(a,e)** is processed to estimate a 3D point cloud **(b,c,f,g)** that contains both depth and albedo information, here shown for two different scenes (photographs shown in **(d,h)**). The color-coded errors maps **(d,h)** directly compare the results of several depth estimation techniques, including log-matched filtering[32], Coates' method[25] followed by Gaussian fit (on high-flux measurement), and the proposed method.

This reconstruction method improves prior SPAD depth imaging techniques both by accounting for dead time in the observation model, which is crucial for high and medium-flux scenarios, and by jointly estimating $\boldsymbol{\alpha}$ and $\boldsymbol{z}$ directly from the raw histogram data. Previous approaches, on the other hand, apply a sequence of transformations to the data, estimating $\boldsymbol{\alpha}$ and $\boldsymbol{z}$ in separate stages[18,27,31,32], which limits the reconstruction performance.

Experimental validation. We evaluate the proposed method on measurements experimentally acquired with the setup illustrated in Fig. 1. In Fig. 2, we assess the performance of the proposed method on two scenes with highly varying reflectance and depth profiles. Both scenes contain objects with complex geometries and varying reflectance properties, including specular behavior for the "Statue of David" scene and Lambertian reflectance with spatially varying albedo in the "Bas-relief" scene. For both scenes, we capture a ground truth reference measurement with a 5% Neutral Density filter in the laser path which eliminates pileup distortions by damping the source intensity. To minimize shot noise fluctuations at the low count rates, we acquire very long sequences of 6 s length per spot at 4 MHz laser repetition rate. We scan every scene at a spatial resolution of 150×150 points. Hence, a full reference measurement is acquired in $150 \cdot 150 \cdot 6\,\mathrm{s} = 37.5\,\mathrm{h}$ per scene. The ground truth depth is extracted from long-exposure measurements using log-matched filtering[32] with the impulse response calibrated using a planar target captured under the same acquisition settings, i.e. unaffected by pileup. The experimental measurements, which serves as input for the proposed method, are acquired without any filters in the optical path, please see the Supplemental Material for additional details.

For each scene in Fig. 2, depth and albedo reconstructions along with the corresponding error maps are shown. We show perspective renderings of the point cloud reconstructions for two different viewpoints. These results verify that the proposed method achieves high-quality depth and albedo reconstructions unaffected by scene-dependent pileup or shot-noise distortions. Specifically, we compare our approach against the conventional log-matched filter estimate[32] as well as Coates' pileup correction method[25] followed by a Gaussian fit. The effect of pileup becomes apparent in the error map for the log-matched filter estimate resulting in more than 5 mm depth error for the "Statue of David" scene. Existing methods do not effectively suppress pileup and therefore suffer from strongly scene-dependent depth precision. In contrast, our method achieves sub-picosecond accuracy independent of scene depth and reflectance, despite the FWHM of the laser pulses being longer than 50 ps. We include additional comparisons in the Supplemental Results where we also evaluate the proposed approach in the low-flux regime and demonstrate that our probabilistic method achieves close to an order of magnitude increased temporal resolution.

Next, we validate the sub-picosecond accuracy of the proposed approach without using spatial priors, i.e. relying solely on the probabilistic pileup model for individual histogram measurements. Specifically, we acquire a sequence of single-point measurements of a planar Lambertian target which is moved along the optical axis using a linear motion stage for 100 measurement points uniformly spaced with 0.5 mm distance. The setups for this measurement scenario and for an additional horizontal scanline measurement are shown in the Supplemental

Scene Target	Pulse Wavelength	Log-matched Filter Depth Error	Coates’ Method	First Photon on Coates	Proposed Method	Log-matched Filter Round-trip Error	Coates’ Method	First Photon on Coates	Proposed Method
Planar	450 nm	4.54 mm	4.40 mm	3.39 mm	**0.14 mm**	15.13 ps	14.68 ps	11.30 ps	**0.46 ps**
Linear Stage	450 nm	2.79 mm	2.65 mm	1.98 mm	**0.24 mm**	9.31 ps	8.85 ps	6.58 ps	**0.91 ps**
Planar	670 nm	3.19 mm	3.00 mm	2.37 mm	**0.16 mm**	10.64 ps	9.99 ps	7.91 ps	**0.52 ps**
Linear Stage	670 nm	8.13 mm	8.09 mm	8.00 mm	**0.23 mm**	27.11 ps	26.97 ps	26.68 ps	**0.76 ps**

Figure 3. Experimental validation of sub-picosecond accuracy on recorded single-pixel data without spatial priors. The average depth and round-trip time error for two scenes are shown, for the 450 nm Alphalas LD-450-50 laser (FWHM of 90 ps) and the 670 nm Alphalas LD-670-50 laser (FWHM of 50 ps), respectively. The background level is 5% for all scenes. We compare reconstructions of the conventional log-matched filter estimate[32], Coates’ method[25] followed by a Gaussian fit, Shin *et al.*[18] on Coates-corrected measurements, and the proposed method.

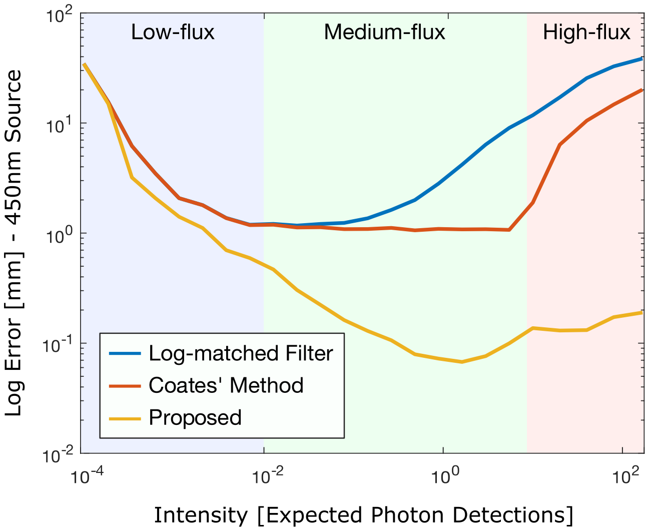

Figure 4. Optimal photon count regime. Depth reconstruction accuracy for varying photon counts for the 450 nm Alphalas LD-450-50 laser (FWHM of 90 ps). The conventional log-matched filter, Coates’ method[25], and the proposed method are compared. The optimal number of photon counts lies around the unconventional region of 1 photon detected per pulse on average, independent of the impulse response and for a broad range of histogram bin widths, see Supplemental Results.

Information. In addition to the conventional log-matched filter estimate[32] and Coates’ method[25], we compare the proposed method against and Shin *et al.*’s method[18] applied on the Coates-corrected histogram data, which adds censoring and background signal suppression[18] to Coates’ method. Figure 3 shows the average absolute error in depth and round-trip time for both of these measurements. The results demonstrate that, even without spatial priors, the proposed probabilistic method outperforms existing state-of-the-art approaches by more than an order of magnitude. Adding the proposed spatial prior reduces temporal error by a factor of 2× on average. We refer to the Supplemental Results for additional experimental results and extensive evaluation in simulation.

Optimal regime of photon counts. Next, we analyze the performance of the proposed method with a fixed dwell or exposure time for a varying incident photon flux. The proposed probabilistic method achieves optimal accuracy in an unconventional, pileup-affected regime with around 1 photon detection per pulse (see Fig. 4). This plot is generated in simulation without the use of spatial priors for the 450 nm Alphalas LD-450-50 laser with $N = 10^4$ shots. As expected, the log-matched filtering approach performs best when the measurements are neither too noisy nor too much affected by pileup. Existing sequential pileup correction methods, such as Coates’ method[25], alleviate pileup and effectively extend the range where optimal performance can be reached beyond 1 expected photon detection per pulse. The accuracy of the proposed method is identical to previous methods for the low-flux regime, because we only consider a single pixel and no spatial priors are used. As the photon count increases, our approach substantially improves upon all existing methods by up to two orders of magnitude. Optimal precision is achieved for expected photon counts around 1 per pulse, which is significantly higher than the low-flux regime in which existing methods operate. These results motivate an optimal photon flux regime for 3D imaging that is far outside the conventional low-flux regime, see Supplemental Results for additional discussion. In the Supplemental Results, we also analyze the effect of the histogram bin width on timing accuracy in this optimal flux regime. The accuracy is consistent across a broad range of histogram bin-widths from sub-2 ps up to 80 ps, demonstrating that the proposed approach not only improves on the state-of-the-art by more than an order of magnitude in accuracy, but also reduces the timing resolution requirements on the detector side.

Discussion

Despite laser pulse widths being larger than 50 ps FWHM, the proposed probabilistic image formation model and corresponding inverse methods achieve sub-picosecond accuracy for active 3D imaging. Moreover, the proposed method achieve this precision across a wide dynamic range, from low-flux to high-flux measurements. These capabilities are achieved by accurately modeling pileup distortions in the image formation model and by solving the inverse problem jointly for all latent variables using statistical priors. In the future, the proposed method could also facilitate long-range acquisition without distance-dependent reduction in repetition rate, by multiplexing multiple pileup-affected responses into the same histogram and unwrapping the pileup-corrected peaks. The proposed method paves the way for fast and precise photon-efficient 3D imaging systems in practical scenarios where widely-varying photon counts are observed in a scene. Applications across disciplines may significantly benefit from these capabilities, including 3D mapping and navigation, autonomous driving, robotic and machine vision, geographic information science, industrial imaging, and microscopic imaging.

Methods

Equipment details. The hardware setup consists of a time-resolved sensor, pulsed laser, illumination and collection optics, and a set of scanning mirrors to achieve a raster scan illumination pattern. The sensor is a PDM SPAD from Micro Photon Devices with a $100 \times 100\,\mu$m sensor area, 27 ps timing jitter (measured with a 100 kHz laser at 675 nm), and 40.9 dark counts per second. Timing of photon arrivals is captured with a PicoHarp 300 Time-Correlated Single Photon Counting (TCSPC) module. The illumination source is a 450 nm or 670 nm picosecond laser (ALPHALAS PICOPOWER-LD-450-50, PICOPOWER-LD-670-50). The 450 nm and 670 nm versions have pulse widths of 90 ps and 50 ps and average power of 0.406 mW and 0.11 mW respectively at a 10 MHz pulse repetition rate. The collection optics are designed to extend the field of view of the SPAD across the area scanned by the illumination source and consist of a 75 mm objective lens (Thorlabs AC508-075-A-ML), a 30 mm relay lens (Thorlabs AC254-030-A-ML) and a microscope objective (Olympus UPLFLN 20× objective). The laser spot is minified using a 50 mm (Thorlabs AC254-050-A-ML) and 250 mm (Thorlabs AC254-250-A-ML) lens relay and scanned with mirrors driven by a two-axis galvanometer (Thorlabs GVS012).

Calibration details. All parameters of the image formation model are described in the observation model section. The calibration of these parameters is performed once for a given laser, i.e., pulse shape, and detector configuration. The detector's photon detection probability $\mu = 0.34$ at 450 nm ($\mu = 0.33$ at 670 nm) and dark count rate $d = 40.9$ c/s were calibrated using the method of Polyakov[33]. The remaining unknowns are $\{a_k, b_k, c_k | k \in 1, \ldots, K\}$ for the Gaussian mixture model $\tilde{g}$ we use $K = 8$ and calibrate for the 450 nm Alphalas LD-450-50 laser the values {(0.57, 220.4, 32.0), (0.0059, 203.8, 10.2), (0.003, 214.6, 18.7), (−0.57, 220.3, 32.0), (0.003, 255.6, 47.5), (0.002, 297.9, 67.4), (0.009, 199, 7.1), (0.0003, 357.4, 143.2)}, and for 670 nm Alphalas LD-670-50 laser the values {(0.04, 199.4, 6.4), (0.004, 206.5, 7.0), (−0.001, 187.9, 12.5), (0.005, 204.8, 17.8), (0.003, 225.1, 27.2), (0.002, 254.6, 42.3), (0.0008, 301.0, 69.3), (0.0003, 388.7, 136.0)}. Measurements of a scene consisting of a single diffuse reflector placed at 1 m distance are acquired to estimate these parameters. Specifically, an ND filter with 1% transmission is placed in the illumination path, and a low-noise histogram of a single point reflector is accumulated with $N = 10^9$ pulses, requiring approximately 17 minutes at a 1 MHz laser repetition rate. The high-absorption ND filter ensures a low-flux regime where pileup can be ignored and where the histogram counts are Poisson distributed. For the high number of 10^9 shots, these histogram measurements can be well approximated by a Gaussian distribution. Finally, the mixture parameters are estimated using efficient, conventional expectation-maximization algorithms[34]. Please see the Supplemental Methods for calibration results for the for the 450 nm Alphalas LD-450-50 laser and the 670 nm Alphalas LD-670-50 laser.

Algorithm parameters. The algorithm hyper-parameters are $\gamma_1 = \gamma_2 = 10^{-1}$. These values were determined using simulations. Initial values of $\boldsymbol{\alpha} = 1$, $\boldsymbol{z} = 0$ are used for the reconstruction algorithm in the algorithm section. In practice, warm-starting with the log-matched filter or Coates' estimate reduces the iterations needed to converge. The termination criteria for the proposed non-convex optimization algorithm are detailed in the Supplemental Methods. The algorithm run time with unoptimized Matlab code averages around 100 s on a Intel i7 2.6 GHz notebook computer. Note that the likelihood proximal operator from Eq. (5), which dominates the algorithm runtime, is embarrassingly parallel across all measurement points, and hence is suited for implementation on modern GPU hardware. Specifically, the operator is parallelizable into $150 \cdot 150 = 22500$ independent threads for the proposed setup configuration, promising run times on the order of a few milliseconds in future implementations, see Supplemental Material.

Code and Data Availability

The code and data used to generate the findings of this study are available on GitHub: https://github.com/computational-imaging/spad_pileup.

References

1. Schwarz, B. Lidar: Mapping the world in 3d. *Nature Photonics* **4**, 429 (2010).
2. Bronzi, D. *et al.* Automotive three-dimensional vision through a single-photon counting spad camera. *IEEE Transactions on Intelligent Transportation Systems* **17**, 782–795 (2016).
3. McCarthy, A. *et al.* Kilometer-range, high resolution depth imaging via 1560 nm wavelength single-photon detection. *Opt. Express* **21**, 8904–8915 (2013).
4. Pawlikowska, A. M., Halimi, A., Lamb, R. A. & Buller, G. S. Single-photon three-dimensional imaging at up to 10 kilometers range. *Opt. Express* **25**, 11919–11931 (2017).

5. Raibert, M., Blankespoor, K., Nelson, G. & Playter, R. Bigdog, the rough-terrain quadruped robot. *IFAC Proceedings Volumes* **41**, 10822–10825 (2008).
6. Côté, J.-F., Widlowski, J.-L., Fournier, R. A. & Verstraete, M. M. The structural and radiative consistency of three-dimensional tree reconstructions from terrestrial lidar. *Remote Sensing of Environment* **113**, 1067–1081 (2009).
7. Velten, A. *et al*. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature communications* **3**, 745 (2012).
8. O'Toole, M., Wetzstein, G. & Lindell, D. B. Confocal Non-Line-of-Sight Imaging Based on the Light Cone Transform. *Nature* 338–341 (2018).
9. Gariepy, G., Tonolini, F., Henderson, R., Leach, J. & Faccio, D. Detection and tracking of moving objects hidden from view. *Nature Photonics* **10**, 23–26 (2016).
10. Morris, P. A., Aspden, R. S., Bell, J. E., Boyd, R. W. & Padgett, M. J. Imaging with a small number of photons. *Nature communications* **6**, 5913 (2015).
11. Chen, Y., Müller, J. D., So, P. T. & Gratton, E. The photon counting histogram in fluorescence fluctuation spectroscopy. *Biophysical journal* **77**, 553–567 (1999).
12. Aull, B. F. *et al*. Geiger-mode avalanche photodiodes for three-dimensional imaging. *Lincoln Laboratory Journal* **13**, 335–349 (2002).
13. Rochas, A. *et al*. First fully integrated 2-d array of single-photon detectors in standard cmos technology. *IEEE Photonics Technology Letters* **15**, 963–965 (2003).
14. Niclass, C., Rochas, A., Besse, P.-A. & Charbon, E. Design and characterization of a cmos 3-d image sensor based on single photon avalanche diodes. *IEEE Journal of Solid-State Circuits* **40**, 1847–1854 (2005).
15. Richardson, J. A., Grant, L. A. & Henderson, R. K. Low dark count single-photon avalanche diode structure compatible with standard nanometer scale cmos technology. *IEEE Photonics Technology Letters* **21**, 1020–1022 (2009).
16. Villa, F. *et al*. Cmos imager with 1024 spads and tdcs for single-photon timing and 3-d time-of-flight. *IEEE journal of selected topics in quantum electronics* **20**, 364–373 (2014).
17. Bronzi, D. *et al*. 100 000 frames/s 64 × 32 single-photon detector array for 2-d imaging and 3-d ranging. *IEEE journal of selected topics in quantum electronics* **20**, 354–363 (2014).
18. Shin, D. *et al*. Photon-efficient imaging with a single-photon camera. *Nature Communications* **7** (2016).
19. Lindell, D. B., O'Toole, M. & Wetzstein, G. Single-photon 3d imaging with deep sensor fusion. *ACM Transactions on Graphics (TOG)* **37**, 113 (2018).
20. Rapp, J. & Goyal, V. K. A few photons among many: Unmixing signal and noise for photon-efficient active imaging. *IEEE Transactions on Computational Imaging* **3**, 445–459 (2017).
21. Migdall, A., Polyakov, S. V., Fan, J. & Bienfang, J. C. *Single-photon generation and detection: physics and applications*, vol. 45 (Academic Press, 2013).
22. Tosi, A. *et al*. Fast-gated single-photon avalanche diode for extremely wide dynamic-range applications. In *Design and Quality for Biomedical Technologies II*, vol. 7170, 71700K (International Society for Optics and Photonics, 2009).
23. Cova, S., Longoni, A. & Ripamonti, G. Active-quenching and gating circuits for single-photon avalanche diodes (spads). *IEEE Transactions on nuclear science* **29**, 599–601 (1982).
24. Busck, J. & Heiselberg, H. Gated viewing and high-accuracy three-dimensional laser radar. *Applied optics* **43**, 4705–4710 (2004).
25. Coates, P. The correction for photon 'pile-up' in the measurement of radiative lifetimes. *Journal of Physics E: Scientific Instruments* **1**, 878 (1968).
26. O'Connor, D. *Time-correlated single photon counting* (Academic Press, 2012).
27. Kirmani, A. *et al*. First-photon imaging. *Science* **343**, 58–61 (2014).
28. Shin, D., Xu, F., Wong, F. N. C., Shapiro, J. H. & Goyal, V. K. Computational multi-depth single-photon imaging. *Opt. Express* **24**, 1873–1888 (2016).
29. Pediredla, A.K. *et al.* Signal Processing Based Pile-up Compensation for Gated Single-Photon Avalanche Diodes. arXiv:1806.07437, 1-17 (2018).
30. Rapp, J. *et al.* Dead Time Compensation for High-Flux Ranging. arXiv:1810.11145, 1-13 (2018).
31. Gariepy, G., Tonolini, F., Henderson, R., Leach, J. & Faccio, D. Detection and tracking of moving objects hidden from view. *Nature Photonics* (2015).
32. Shin, D., Kirmani, A., Goyal, V. K. & Shapiro, J. H. Photon-efficient computational 3-d and reflectivity imaging with single-photon detectors. *IEEE Transactions on Computational Imaging* **1**, 112–125 (2015).
33. Polyakov, S. V. *Single-Photon Detector Calibration* (Elsevier, 2015).
34. Murphy, K. P. *Machine learning: a probabilistic perspective* (The MIT Press, Cambridge, MA, 2012).

Acknowledgements

This work was in part supported by a National Science Foundation CAREER award (IIS 1553333), by a Sloan Fellowship, by the DARPA REVEAL program, and by the KAUST Office of Sponsored Research through the Visual Computing Center CCF grant. The authors would like to thank Rafael Setra, Kai Zang, Matthew O'Toole, Amy Fritz, and Mark Horowitz for fruitful discussions in early stages of this project. S.D. was supported by a National Science Foundation GraduateResearch Fellowship and D.B.L. was supported by a Stanford Graduate Fellowship in Science and Engineering.

Author Contributions

F.H. and G.W. conceived the idea. F.H. derived the algorithm. F.H. and S.D. implemented the algorithm. F.H. processed all data. F.H., D.B.L. and G.W. designed the experiments. D.B.L. captured experimental data. F.H., S.D. and G.W. wrote the manuscript. G.W. supervised all aspects of the project.

Additional Information

Supplementary information accompanies this paper at https://doi.org/10.1038/s41598-018-35212-x.

Competing Interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit http://creativecommons.org/licenses/by/4.0/.

© The Author(s) 2018