

Asynchronous Single-Photon 3D Imaging

Anant Gupta* Atul Ingle* Mohit Gupta

{anant,ingle,mohitg}@cs.wisc.edu

University of Wisconsin-Madison

www.SinglePhoton3DImaging.com

Abstract

Single-photon avalanche diodes (SPADs) are becoming popular in time-of-flight depth-ranging due to their unique ability to capture individual photons with picosecond timing resolution. However, ambient light (e.g., sunlight) incident on a SPAD-based 3D camera leads to severe non-linear distortions (pileup) in the measured waveform, resulting in large depth errors. We propose asynchronous single-photon 3D imaging, a family of acquisition schemes to mitigate pileup during data acquisition itself. Asynchronous acquisition temporally misaligns SPAD measurement windows and the laser cycles through deterministically predefined or randomized offsets. Our key insight is that pileup distortions can be “averaged out” by choosing a sequence of offsets that span the entire depth range. We develop a generalized image formation model and perform theoretical analysis to explore the space of asynchronous acquisition schemes and design high-performance schemes. Our simulations and experiments demonstrate an improvement in depth accuracy of up to an order of magnitude as compared to the state-of-the-art, across a wide range of imaging scenarios, including those with high ambient flux.

1. Single-Photon Cameras

Light is fundamentally quantized; any camera records incoming light not continuously, but in discrete packets called photons. A conventional camera typically captures hundreds to thousands of photons per pixel to create an image. What if cameras could record *individual* photons, and, precisely measure their time-of-arrival? Not only would such cameras have extremely high sensitivity, but the captured data will have an additional time-dimension, a rich source of information inaccessible to conventional cameras.

There is an emerging class of sensors, called single-photon avalanche diodes (SPADs) [30] that promise single-photon sensitivity (Fig. 1(a)) and the ability to time-tag photons with picosecond precision. Due to these capabilities, SPADs are driving novel functionalities such as non-line-of-sight (NLOS) imaging [7, 22] and microscopy of bio-

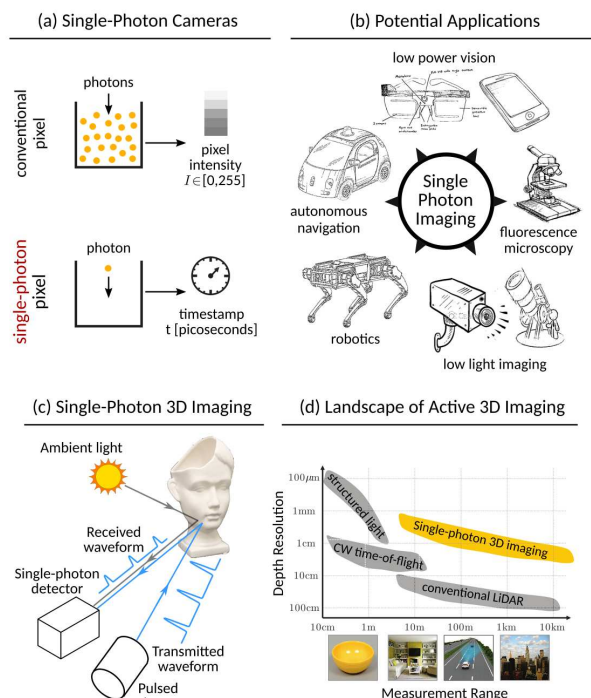


Figure 1. **Single-photon cameras and 3D imaging.** (a) A single-photon camera pixel is sensitive to individual photons and can capture photon arrival times with picosecond resolution. (b) The extreme sensitivity and resolution makes single-photon cameras promising candidates for several applications. (c) A single-photon 3D camera based on time-of-flight consists of a pulsed laser and a single-photon detector that timestamps returning photons. (d) Single-photon 3D cameras have the potential to provide extremely high depth resolution, even at long ranges.

phenomena at nano time-scales [4]. However, so far, SPADs are considered specialized devices suitable only for photon-starved (dark) scenarios, and thus, restricted to a limited set of niche applications. This raises the following questions: Can SPADs operate not just in low-light, but across the entire gamut of imaging conditions, including high-flux scenes [15]? In general, is it possible to leverage the exciting capabilities of SPADs for a broader set of mainstream computer vision applications (Fig. 1(b))?

In this paper, we address the above questions in the context of 3D imaging. Consider a single-photon 3D camera based on time-of-flight (ToF). It consists of a pulsed laser

*Equal contribution

†This research was supported by ONR grant N00014-16-1-2995, DARPA REVEAL program and Wisconsin Alumni Research Foundation.

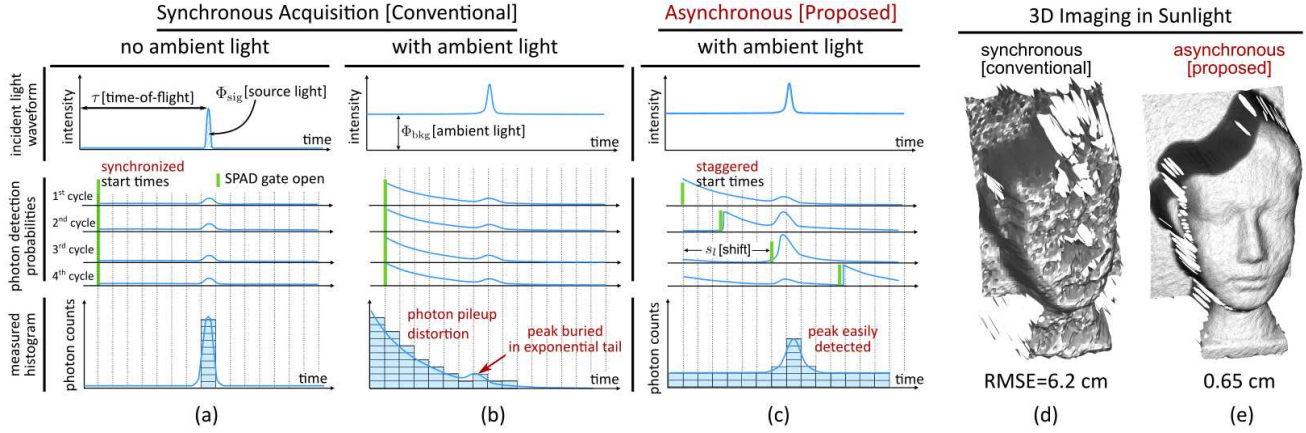


Figure 2. **Imaging model of single-photon 3D cameras.** (a) A single-photon 3D camera records the timestamps of returning photons over many laser cycles and constructs a histogram of photon arrival times. In the absence of ambient light, the peak of this histogram corresponds to the true depth. (b) In the conventional (synchronous) operation, ambient light causes photon pileup which distorts the histogram towards earlier time bins. (c) Asynchronous acquisition prevents pileup by temporally staggering the SPAD cycles with respect to the laser cycles, distributing the effect of pileup uniformly over all histogram bins. (d) 3D shape recovered using synchronous acquisition shows large depth errors due to pileup. (e) Proposed asynchronous method recovers accurate 3D shape even in high ambient light.

emitting periodic pulses of light toward the scene, and a SPAD sensor (Fig. 1(c)). Although several conventional 3D cameras also use the ToF principle, single-photon 3D cameras have a fundamentally different imaging model. The SPAD detects *at most one returning photon* per laser pulse, and records its time-of-arrival. Arrival times over several laser pulses are recorded to create a temporal histogram of photon arrivals, as shown in Fig. 2(a). Under low incident flux, the histogram is approximately a linearly scaled replica of the incident waveform, and thus, can be used to recover scene depths [27, 19]. Due to the high timing resolution of SPADs, single-photon 3D cameras are capable of achieving “laser-scan quality” depth resolution (1–10 mm), at long distances (100–1000 meters) (Fig. 1(d)).

Single-photon 3D imaging in sunlight: Due to the peculiar histogram formation process, single-photon 3D cameras cannot operate reliably under ambient light (e.g., sunlight in outdoor conditions). This is because early arriving ambient photons prevent the SPAD from measuring the signal (laser) photons that may arrive at a later time bin of the histogram. This distorts the histogram measurements towards earlier time bins, as shown in Fig. 2(b). This non-linear distortion, known as *photon pileup* [14, 25, 17] makes it challenging to reliably locate the laser pulse, resulting in large depth errors. Although there has been a lot of research toward correcting these distortions in post-processing [14, 9, 23, 25, 28], strong pileup due to ambient light continues to limit the scope of this otherwise exciting technology.

We propose asynchronous single-photon 3D imaging, a family of computational imaging techniques for SPAD-based 3D cameras with the goal of *preventing* pileup during acquisition itself. In conventional ToF cameras, the laser and sensor are temporally synchronized. In contrast, we *desynchronize* the SPAD acquisition windows with respect

to the laser pulses. This introduces different temporal offsets between laser cycles and SPAD acquisition windows, as shown in Fig. 2(c). The key insight is that cycling through a range of temporal offsets (across different laser cycles) enables detecting photons in later time bins that would otherwise have been masked by early-arriving ambient photons. This distributes the effect of pileup across all histogram bins, thus eliminating the structured distortions caused by the synchronous measurements, as shown in Fig. 2(c).

At first glance, it may appear that such asynchronous measurements may not provide consistent depth information. The main idea lies in *computationally re-synchronizing* the photon timing measurements with the laser cycles. To this end, we develop a generalized image formation model and derive a maximum likelihood estimator (MLE) of the true depth that accounts for arbitrary temporal offsets between measurement and laser cycles. Based on these ideas, we propose two asynchronous acquisition methods: uniform and photon-driven, which shift the SPAD window with respect to laser either deterministically or stochastically. These techniques can be implemented with minimal modifications to existing systems, while achieving up to an order-of-magnitude improvements in depth accuracy. An example is shown in Fig. 2(d–e).

Implications and future outlook: Due to their compatibility with mainstream CMOS sensor fabrication lines, the capabilities of SPAD cameras continue to grow rapidly [11, 32, 12, 20, 18, 1, 3]. As a result, the proposed methods, aided by rapid ongoing advances in SPAD technology, will potentially spur wide-spread adoption of single-photon sensors as *all-purpose* cameras in demanding computer vision and robotics applications, where the ability to perform reliably in both photon-starved and photon-flooded scenarios is critical to success.

2. Related Work

Photon pileup mitigation for SPAD cameras: Perhaps the most widely adopted approach for preventing pileup is attenuation, i.e., optically blocking the total photon flux incident on the SPAD so that only 1-5% of the laser pulses lead to a photon detection [2, 3].¹ Recent work [14, 13] has shown that this rule-of-thumb *extreme attenuation* is too conservative and the optimal operating flux is considerably higher. Various computational [25, 14] and hardware [1, 3, 35] techniques for mitigating pileup have also been proposed. These approaches are complementary to the proposed asynchronous acquisition, and can provide further improvements in performance when used in combination.

Temporally shifted gated acquisition: Fast-gated detectors [6] have been used previously for range-gated LiDAR, confocal microscopy and non-line-of-sight (NLOS) imaging [7] to preselect a specific depth range and suppress undesirable early-arriving photons. A sequence of shifted SPAD gates has been used in FLIM for improving temporal resolution and dynamic range [34, 31, 32] and for extending the unambiguous depth range of pulsed LiDARs [29]. In contrast, we use shifting to mitigate pileup and present a theoretically optimal method for choosing the sequence of shifts and durations of the SPAD measurement gates without any prior knowledge of scene depths.

Photon-driven acquisition: The photon-driven (or free-running) mode of operation has been analyzed for FLIM [16, 9, 1], and recently for LiDAR [28] where a Markov chain model-based iterative optimization algorithm is proposed to recover the incident waveform from the distorted histogram. The focus of these approaches is on designing efficient waveform estimation algorithms. Our goal is different. We explore the space of asynchronous acquisition schemes with the aim of designing acquisition strategies that mitigate depth errors due to pileup in high ambient light under practical constraints such as a fixed time budget. We also propose a generalized closed-form maximum likelihood estimator (MLE) for asynchronous acquisition that can be computed without any iterative optimization routine.

3. Single-Photon 3D Imaging Model

A SPAD-based 3D camera consists of a pulsed laser that emits short periodic pulses of light toward a scene point, and a co-located SPAD sensor that captures the reflected photons (Fig. 1(c)). Although the incident photon flux is a continuously varying function of time, a SPAD has limited time resolution, resulting in a discrete sampling of the continuous waveform. Let Δ denote the size of each discrete temporal bin (usually on the order of few tens of picoseconds). Assuming an ideal laser pulse modeled as a Dirac-

¹Note that attenuation blocks both ambient and source photons. Attenuation can be achieved through various methods such as spectral filtering, neutral density filtering or using an aperture stop.

delta function $\delta(t)$, the number of photons incident on the SPAD in the i^{th} time bin follows a Poisson distribution with a mean given by:

$$r_i = \Phi_{\text{sig}} \delta_{i,\tau} + \Phi_{\text{bkg}}, \quad (1)$$

where $\delta_{i,j}$ is the Kronecker delta,² $\tau = \lfloor 2z/c\Delta \rfloor$ is the discretized round-trip time delay, z is the distance of the scene point from the camera, and c is the speed of light. Φ_{sig} is the mean number of signal photons (due to the laser pulse) received per bin, and Φ_{bkg} is the (undesirable) background and dark count photon flux per bin. B is the number of time bins in a single laser period. The vector $(r_i)_{i=1}^B$ denotes the *incident photon flux waveform*. A reliable estimate of this waveform is needed to estimate scene depth.

Synchronous acquisition: In order to estimate the incident waveform, SPAD-based 3D cameras employ the principle of time-correlated single-photon counting (TC-SPC) [21, 17, 2, 26, 24, 25]. In conventional *synchronous* acquisition, the SPAD starts acquiring photons immediately after the laser pulse is transmitted, as shown in Fig. 2(a). In each laser cycle (laser repetition period), after detecting the first incident photon, the SPAD enters a *dead time* (~ 100 ns) during which it cannot detect additional photons. The SPAD may remain inactive for longer than the dead time so that the next SPAD acquisition window aligns with the next laser cycle.³ The time of arrival of the first incident photon is recorded with respect to the start of the most recent cycle. A histogram (N_1, \dots, N_B) of the first photon arrival times is constructed over many cycles, where N_i denotes the number of times the first photon arrives in the i^{th} bin. In low ambient light, the histogram, on average, is simply a scaled version of the incident waveform [13], from which, depth can be estimated by locating its peak.

Effect of ambient light in synchronous acquisition: Under ambient light, the incident flux waveform can be modeled as an impulse with a constant d.c. offset, as shown in the top of Fig. 2(b). In high ambient flux, the SPAD detects an ambient photon in the earlier histogram bins with high probability. This skews the measured histogram towards earlier histogram bins, as shown in the bottom of Fig. 2(b). The peak due to the laser source appears only as a small blip in the exponentially decaying tail of the measured histogram. This distortion, called photon pileup [8, 2, 25], significantly lowers the accuracy of depth estimates.

In the next two sections we introduce a generalization of the synchronous TCSPC acquisition scheme, and show how it can be used to mitigate pileup distortion and reliably estimate depths, even in the presence of high ambient light.

² $\delta_{i,j} = 1$ for $i = j$ and 0 otherwise.

³The laser repetition period is set to $2z_{\text{max}}/c$, where z_{max} is the unambiguous depth range. The photon flux is assumed to be 1-5% [33] of the laser repetition rate so that the probability of detecting photons in consecutive laser cycles is negligible. In high ambient light, the dead time from one cycle may extend into the next causing some cycles to be skipped.

4. Theory of Asynchronous Image Formation

In this section we develop a theoretical model for asynchronous single-photon 3D cameras. We derive a histogram formation model and a generalized Coates's estimator [8] for the incident photon flux waveform. In asynchronous acquisition, we decouple the SPAD on/off times from the laser cycles by allowing the SPAD acquisition windows to have arbitrary start times with respect to the laser pulses (Fig. 2(c)). A SPAD cycle is defined as the duration between two consecutive time instants when the SPAD sensor is turned on. The SPAD will remain inactive during some portion of each SPAD cycle due to its dead time.




Figure 3. **Histogram formation for asynchronous acquisition.** (Top) The temporal location of the laser peak in the incident waveform corresponds to the round-trip time-of-flight. A slightly longer SPAD cycle period results in a sequence of increasing shifts with respect to the laser cycles. (Bottom) The histogram formation process involves computational resynchronization of photon arrival times to the laser cycle boundaries, causing a “wrap around.” The measured histogram approaches the true waveform shape when a large number of uniformly spaced shifts is used.

Each laser cycle consists of B time bins which are used to build a photon count histogram. The bin indices are defined with respect to the start of the laser cycle, i.e., the first time bin is aligned with the transmission of each laser pulse. We assume that the laser repetition period is $B\Delta = 2z_{\max}/c$. This ensures that a photon detected by the SPAD always corresponds to an unambiguous depth range $[0, z_{\max}]$. Let s_l ($0 \leq s_l \leq B-1$) denote the bin index (with respect to the most recent laser cycle) at which the SPAD gate is activated during the l^{th} SPAD cycle ($1 \leq l \leq L$). As shown in Fig. 3(top), a SPAD cycle may extend over multiple consecutive laser cycles.

Probability distribution of measured histogram: Due to Poisson statistics, the probability q_i that at least one photon

is incident on the SPAD in the i^{th} bin is:

$$q_i = 1 - e^{-r_i}, \quad (2)$$

where r_i is given by Eq. (1). A photon detection in the i^{th} time bin occurs when no photon is incident in the time bins preceding the i^{th} bin in the current cycle, and at least one photon is incident in the i^{th} bin. The probability $p_{l,i}$ of a photon detection in the i^{th} bin in the l^{th} SPAD cycle depends on the shift s_l , and is given by:

$$p_{l,i} = q_i \prod_{j:j < i} (1 - q_j), \quad (3)$$

where it is understood (see Supplementary Note 1) that $j < i$ denotes the bin indexes preceding the i^{th} bin in a modulo- B sense with a “wrap around” depending on the shift s_l (Fig. 3(bottom)). We introduce an additional $(B+1)^{\text{th}}$ bin in the histogram to record the number of cycles where no photons were detected, with corresponding bin probability $p_{l,B+1} := 1 - \sum_{i=1}^B p_{l,i}$.

As in the synchronous case, we construct a histogram of the number of photons detected in each time bin. Let N_i be the number of photons captured in the i^{th} bin over L SPAD cycles. As shown in Supplementary Note 1, the joint distribution of the measured histogram $(N_1, N_2, \dots, N_B, N_{B+1})$ is given by a Poisson-Multinomial Distribution (PMD) [10]. The PMD is a generalization of the multinomial distribution; if $s_l = 0 \forall l$ (conventional synchronous operation), this reduces to a multinomial distribution [14, 25].

Characterizing pileup in asynchronous operation: Similar to the synchronous case, in the low incident flux regime ($r_i \ll 1 \forall i$) the measured histogram is, on average, a linearly scaled version of the incident flux: $\mathbb{E}[N_i] \approx Lr_i$, and the incident flux can be estimated as $\hat{r}_i = N_i/L$. However, in high ambient light, the photon detection probability at a specific histogram bin depends on its position with respect to the beginning of the SPAD cycle. Similar to synchronous acquisition, histogram bins that are farther away from the start of the SPAD cycle record photons with exponentially smaller probabilities compared to those near the start of the cycle. However, unlike the synchronous case, the shape of this pileup distortion wraps around at the B^{th} histogram bin during computational resynchronization. This is shown in Fig. 3(bottom). The segment that is wrapped around depends on s_l and may vary with each SPAD cycle.

Computational pileup correction in asynchronous acquisition: A computational pileup correction algorithm must use the histogram $(N_i)_{i=1}^{B+1}$ to estimate the true waveform r_i via an estimate of q_i and Eq. (2). Recall that a photon detection in a specific histogram bin prevents subsequent bins from recording a photon. Therefore, in the high flux regime, q_i cannot be simply estimated as the ratio of N_i to the number of SPAD cycles (L); the denominator in this

ratio must account for the number of SPAD cycles where the i^{th} histogram bin had an opportunity to record a photon.

Definition 1 (Denominator Sequence). Let $D_{l,i}$ be an indicator random variable which is 1 if, in the l^{th} SPAD cycle, no photon was detected before the i^{th} time bin. The denominator sequence $(D_i)_{i=1}^B$ is defined as $D_i = \sum_{l=1}^L D_{l,i}$.

Note that $D_{l,i} = 1$ indicates that in the l^{th} SPAD cycle, the SPAD had an opportunity to detect a photon in the i^{th} bin. By summing over all SPAD cycles, D_i denotes the total number of photon detection opportunities in the i^{th} histogram bin. Using this corrected denominator, an estimate for q_i is obtained as follows:

$$\hat{q}_i = \frac{N_i}{D_i}.$$

We show in Supplementary Note 1 that \hat{q}_i is in fact the MLE of q_i . The MLE of the incident flux waveform is given by:

$$\hat{r}_i = \ln \left(\frac{1}{1 - \hat{q}_i} \right) \quad (4)$$

which is a generalization of the Coates's estimator [8, 25]. Photon pileup causes later histogram bins to have $D_i \approx 0$ making it difficult to estimate r_i . Intuitively, a larger D_i denotes more “information” in the i^{th} bin, hence a more reliable estimate of the true flux waveform can be obtained.

5. Photon Pileup: Prevention Better than Cure?

In theory, when operating in high ambient light, the generalized Coates's estimator in Eq. (4) can invert pileup distortion for asynchronous acquisition with any given set of shifts s_l . However, if the asynchronous acquisition scheme is not well-designed, this inversion will lead to unreliable waveform estimates. For example, if the shifts s_l are all zero (synchronous acquisition), bins farther from the start of the SPAD cycle will have $D_i \approx 0$ and suffer from extremely noisy flux estimates.

In this section, we design imaging techniques that prevent photon pileup in the acquisition phase itself, even under high ambient light. Our main observation is that delaying the start of the SPAD cycle with respect to the start of a laser cycle increases D_i at later time bins. *The key idea*, as shown in Fig. 3, is to cycle through various shifts s_l for different SPAD cycles. This ensures that each time bin is close to the start in at least a few SPAD cycles. Intuitively, if all possible shifts from 0 to $B-1$ are used, the effect of the exponentially decaying pileup due to ambient photons gets distributed over all histogram bins equally. On the other hand, returning signal photons from the true laser peak add up “coherently” because their bin location remains fixed. As a result, the accumulated histogram has enough photons in all bins (Fig. 3(e)) to enable reliable Coates's estimates.




Figure 4. **Simulated depth RMSE at different ambient and signal flux levels.** Asynchronous acquisition with uniform shifting achieves lower error than synchronous acquisition with no and extreme attenuation [13], over a wide range of flux conditions.

We characterize the space of all shifting strategies by their *shift sequence*, $(s_i)_{i=1}^L$. For now, we only consider *deterministic shift sequences*, which means that the shifts are fixed and known prior to acquisition. Given these definitions, the question that we seek to address is: What is the optimal shifting strategy that minimizes depth estimation error? We now present two key theoretical results towards answering this question for a SPAD-based 3D camera operating in the *high ambient flux regime* where the total number of incident photons is dominated by ambient photons.⁴

Definition 2 (Uniform Shifting). A shifting strategy is said to be uniform if its shift sequence is a uniform partition of the time interval $[0, B\Delta]$, i.e., is a permutation of the sequence $(0, \lfloor B/L \rfloor, \lfloor 2B/L \rfloor, \dots, \lfloor (L-1)B/L \rfloor)$.

Result 1 (Denominator Sequence and Probability of Depth Error). *In the high ambient flux regime, among all denominator sequences with a fixed total expected sum $\sum_{i=1}^L \mathbb{E}[D_i]$, an upper bound on the average probability of depth error for the estimator in Eq. (4) is minimized when $\mathbb{E}[D_i] = \mathbb{E}[D_j] \forall i, j$.*

Result 2 (Denominator Sequence for Uniform Shifting). *Uniform shifting achieves a constant expected denominator sequence.*

Interpreting Results 1 and 2: As shown in Supplementary Note 2, for a fixed L , different shift sequences will lead to different denominator sequences but the total expected denominator $\sum_{i=1}^L \mathbb{E}[D_i]$ remains constant. The first result (based on [13]) shows that if a shifting strategy can achieve a constant expected denominator sequence, it will have lower depth error than all other shifting strategies (including synchronous acquisition). The second result shows that there exists a shifting strategy that achieves a constant expected denominator: uniform shifting. As a byproduct, a uniform denominator sequence makes the depth errors invariant to the true bin location, unlike the synchronous case where later time bins suffer from higher depth errors.

⁴We define signal-to-background-ratio $\text{SBR} = \Phi_{\text{sig}}/B\Phi_{\text{bkg}}$. In the high ambient flux regime, $\text{SBR} \ll 1$.

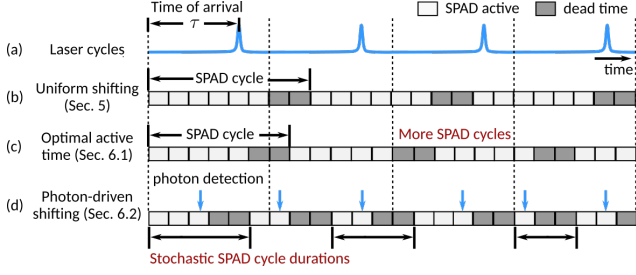


Figure 5. **Different asynchronous acquisition methods.** (a) The incident waveform has a period equal to the laser cycle period. (b) Uniform shifting staggers the laser and SPAD cycles by introducing a mismatch in cycle lengths. (c) Optimizing the SPAD active time enables more SPAD cycles to fit within a fixed total capture time. (d) Photon-driven shifting has random SPAD cycle lengths determined by photon detection events.

Single-pixel simulations: We compare the performance of uniform shifting and conventional synchronous acquisition through Monte Carlo simulations.⁵ We use a histogram with $B = 1000$ and $\Delta = 100$ ps and consider a wide range of background and signal photon flux levels in the discrete delta pulse model of Eq. (1). Uniform shifts are simulated by choosing equally spaced shifts between 0 and 1000 and generating photon counts using Eq. (S3). Depth is estimated using the generalized estimator (Eq. (4)). As seen in Fig. 4, the depth RMSE with uniform shifting is considerably lower than conventional synchronous acquisition. At certain combinations of signal and background flux levels, uniform shifting estimates depths with almost zero RMSE while the conventional methods give a very high error.

6. Practically Optimal Acquisition for Single-Photon 3D Imaging in Bright Sunlight

The theoretical analysis in the previous section shows uniform shifting minimizes an upper bound on the ℓ_0 depth error. It is natural to ask: How can we implement practical uniform shifting approaches that are not just theoretically optimal in the ℓ_0 sense, but also achieve good RMSE (ℓ_2 error) performance under realistic constraints and limited acquisition time? In this section, we design several high-performance shifting schemes based on uniform shifting. These are summarized in Fig. 5.

Uniform shifting can be implemented in practice by making the SPAD cycle period longer than the laser cycle (Fig. 5(b)), and relying on this mismatch to automatically cycle through all possible shifts. Moreover, this can be implemented at a negligible additional cost in terms of total acquisition time as shown in Supplementary Note 3.

6.1. SPAD Active Time Optimization

So far, we have assumed the SPAD active time duration is fixed and equal to $B\Delta$. Programmable fast-gated SPAD

⁵ The maximum possible relative RMSE is 30% because it is defined in a modulo-B sense. See Supplementary Note 8.

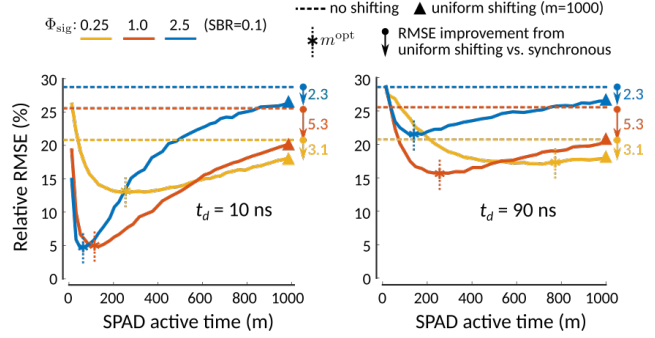


Figure 6. **Effect of SPAD active time on performance.** This plot shows the improvement in RMSE from SPAD active time optimization at different signal strengths and dead times. Note that the RMSE gain (size of vertical arrows) due to uniform shifting over synchronous acquisition remains unchanged across dead times. For low flux levels and long dead times, a longer active time improves RMSE by enabling the SPAD to capture more photons.

detectors [6, 5] allow flexibility in choosing different active time and SPAD cycle durations (Fig. 5(c)). Arbitrary shift sequences can also be implemented by varying the number of active time bins, m , while keeping the inactive duration fixed at t_d . This expands the space of shifting strategies characterized by the active time bins, m , and the shift sequence, $(s_l)_{l=1}^L$. Under a fixed acquisition time constraint:

$$L(m\Delta + t_d) \leq T. \quad (5)$$

Note that L can now vary with m . Can this greater design flexibility be used to improve depth estimates?

Varying m leads to an interesting trade-off. Shortening the active time duration causes a larger proportion of each SPAD cycle to be taken up by dead time. On the other hand, using a very long active time is inefficient because the portion of the active time after the first photon arrival is spent in dead time anyway. This raises the question: What is the optimal active time that minimizes the depth error? In Supplementary Note 4 we show that the optimal active time for uniform shifting is given by:

$$m^{\text{opt}} = \arg \max_m \frac{T}{m\Delta + t_d} \frac{1 - e^{-m\Phi_{\text{bkg}}}}{1 - e^{-\Phi_{\text{bkg}}}}. \quad (6)$$

Simulation results for varying active time: Fig. 6 shows plots of depth RMSE vs. m for a wide range of ambient flux levels and two different values of dead time. Observe that the RMSE curves have local minima which agree with our theory (Eq. (6)). For a wide range of photon flux levels considered here, m^{opt} is *shorter* than the conventionally used active time of $m = B = 1000$ and gives a remarkable improvement in RMSE by up to a factor of 6.

6.2. Photon-Driven Shifting

The optimal active time criterion balances the tradeoff between short and long active time windows in an aver-




Figure 7. **Simulation-based evaluation of practically optimal asynchronous acquisition.** (a) Asynchronous acquisition with optimal SPAD active time (Section 6.1) provides an order of magnitude lower depth RMSE as compared to existing methods. (b) Photon-driven shifting (Section 6.2) further lowers RMSE by allowing the active time to vary stochastically on a per-photon basis.

age sense. However, due to the large variance in the arrival time of the first photon in each cycle, a fixed m cannot achieve both these goals on a per-photon basis. It is possible to achieve photon-adaptive active time durations using the free-running mode [28] where the SPAD is always active, except after a photon detection when it enters a dead time.

In the free-running mode the active times and SPAD cycle durations vary randomly due to the stochastic nature of photon arrivals. As shown in Fig. 5(d), this creates different shifts across different SPAD cycles. Over a sufficiently long acquisition time, a uniformly spaced sequence of shifts is achieved with high probability, distributing the effect of pileup over all histogram bins uniformly. We call this randomized shifting phenomenon *photon-driven shifting*.

Depth estimator for photon-driven shifting: Unlike deterministic shifting, the shift sequence in photon-driven shifting is stochastic because it depends on the random photon arrival times. In Supplementary Note 5 we show that the scene depths can still be estimated using the generalized Coates’s estimator (Eq. (4)) as before⁶.

The following result states that photon-driving shifting possesses the desirable property of providing a uniform shift sequence. See Supplementary Note 5 for a proof.

Result 3. As $L \rightarrow \infty$, *photon-driven shifting achieves a uniform shift sequence.*

Result 3 says that photon-driven shifting exhibits a pileup averaging effect, similar to uniform shifting. Although this does not establish a relationship between the shift sequence and depth RMSE, our results show up to an order of magnitude improvement in RMSE compared to conventional synchronous acquisition.⁷

Simulation results: Fig. 7 shows simulated RMSE results for photon-driven shifting over a wide range of signal and

⁶Note that the joint distribution of $(N_i)_{i=1}^L$ derived in Section 4 does not hold for photon-driven shifting because the shift sequence is random.

⁷Proving exact uniformity of the shift sequence requires assuming $L \rightarrow \infty$, although in practice, we found that $L \leq 50$ SPAD cycles is sufficient.

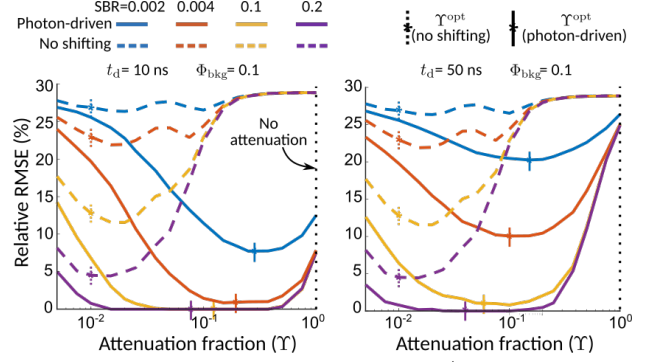


Figure 8. **Optimal attenuation factor Υ^{opt}** for photon-driven shifting is higher than that of synchronous acquisition, leading to efficient flux utilization, while minimizing photon pileup.

ambient flux levels. For some flux levels the proposed shifting methods provide almost zero depth error while the conventional method has the maximum possible error. The RMSE of photon-driven shifting is similar to uniform shifting with m^{opt} , but for some flux levels it can provide a factor of 2 improvement over uniform shifting. Supplementary Note 6 discusses certain regimes where deterministic shifting may be preferable over photon-driven shifting.

6.3. Combination with Flux Attenuation

Recent work [13, 14] has shown that there is an optimal incident flux level at which pileup in a synchronous SPAD-based 3D camera is minimized while maintaining high SNR. This optimal flux can be achieved by optically attenuating the incident photon flux. In the space of acquisition strategies to deal with pileup, attenuation can be considered a complementary approach to asynchronous shifting. In Supplementary Note 7, we show that the optimal attenuation fraction for photon-driven shifting is given by:

$$\Upsilon_{\text{photon-driven}}^{\text{opt}} = \min \left\{ 1.0, \arg \min_{\Upsilon} \frac{1 + t_d(1 - e^{-\Upsilon \Phi_{\text{bkg}}})}{e^{-\Upsilon \Phi_{\text{bkg}}}(1 - e^{-\Upsilon \Phi_{\text{sig}}})} \right\}.$$

Fig. 8 shows simulation results of depth RMSE for the conventional synchronous mode and photon-driven shifting over a range of attenuation factors and two different dead times. The locations of the minima agree with our theory. There are two key observations. First, the optimal attenuation fraction with shifting is much higher than that for conventional synchronous acquisition. Second, combining attenuation with photon-driven shifting can provide a large gain in depth error performance, reducing the RMSE to almost zero under certain conditions.

7. Experiments

Our hardware prototype consists of a 405 nm pulsed laser (Picoquant LDH-P-C-405B), a TCSPC module (Picoquant HydraHarp 400) and a fast-gated SPAD [6] that can be operated in both triggered and free-running modes and




Figure 9. **Experimental demonstration of single-photon 3D imaging under strong ambient light.** A white “PorcelainFace” vase was illuminated with high ambient light of $B\Phi_{\text{bkg}} = 11$ photons and scanned with a low-power laser at an SBR of 0.02. The proposed asynchronous acquisition schemes achieve considerably higher depth quality as compared to conventional synchronous methods.

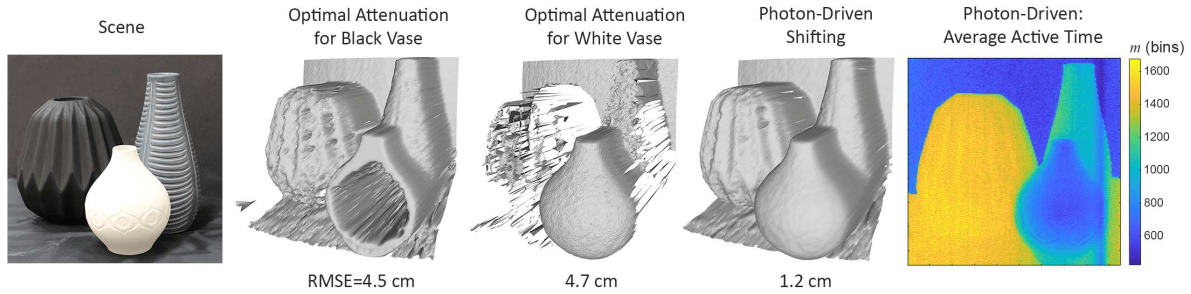


Figure 10. **Adaptivity of photon-driven shifting to different albedos.** The black vase in this “Vases” scene has $1/10^{\text{th}}$ the reflectivity of the white vase. With synchronous acquisition, the attenuation fraction must be adjusted individually for each vase. In contrast, both vases are reliably reconstructed with photon-driven shifting which automatically adapts the active time duration for each pixel.

has a programmable dead time which was set to 50 ns. We operated the laser at a repetition frequency of 10 MHz for an unambiguous depth range of 15 m discretized into 1000 histogram bins. For uniform shifting, we operated the SPAD with its internal clock to obtain shifts between the SPAD measurement windows and the laser cycles.

3D point-scanning results: Fig. 9 shows 3D reconstructions of “PorcelainFace” scene under high ambient illumination. Both uniform and photon-driven shifting ($\Upsilon_{\text{photon-driven}}^{\text{opt}} = 1$) perform better than synchronous acquisition methods. Photon-driven acquisition provides sub-centimeter RMSE, which is an order of magnitude better than the state-of-the-art extreme attenuation method.

The method of [13] uses synchronous acquisition and relies on setting an attenuation factor for different parts of the scene based on the total photon flux and hence requires pixel-wise adaptation. The “Vases” scene in Fig. 10 consists of a black vase with a much lower albedo than the white vase. The attenuation fraction needed for the white vase is too low and causes the black vase to appear noisy, whereas the attenuation fraction for the black vase is too high to avoid pileup distortions at the white vase. The average active time with photon-driven shifting ($\Upsilon_{\text{photon-driven}}^{\text{opt}} = 1$) au-

tomatically adapts to different photon flux levels and reliably captures the depth map for both vases. For darker scene points the average active time is *longer* than the laser cycle period of $B = 1000$.

8. Limitations and Discussion

Incorporating spatial priors: The theoretical analysis and results presented here are limited to a pixel-wise depth estimator which uses the MLE of the photon flux waveform. Further improvements can be obtained by incorporating spatial priors in a regularized optimization framework [14], or data-driven neural network-based approaches [19] that exploit spatial correlations between neighboring pixels and across different training images to improve depth accuracy.

Extension to other active-imaging modalities: The idea of using asynchronous acquisition schemes can be extended to other SPAD-based active-imaging applications that use the principle of TCSPC to recover the true shape of the photon flux waveform. Non-uniform shifting schemes may be required for time-domain FLIM where true waveform shape is an exponential decay and NLOS imaging where the photon flux waveform can have arbitrary shapes.

References

- [1] Giulia Acconcia, Alessandro Cominelli, Massimo Ghioni, and Ivan Rech. Fast fully-integrated front-end circuit to overcome pile-up limits in time-correlated single photon counting with single photon avalanche diodes. *Opt. Express*, 26(12):15398–15410, Jun 2018. [2](#), [3](#)
- [2] Wolfgang Becker. *Advanced time-correlated single photon counting applications*, volume 111. Springer, 2015. [3](#)
- [3] Maik Beer, Jan Haase, Jennifer Ruskowski, and Rainer Kokozinski. Background light rejection in spad-based lidar sensors by adaptive photon coincidence detection. *Sensors*, 18(12):4338, Dec 2018. [2](#), [3](#)
- [4] Claudio Bruschini, Harald Homulle, Ivan Michel Antolovic, Samuel Burri, and Edoardo Charbon. Single-photon spad imagers in biophotonics: Review and outlook. *arXiv preprint*, 2019. [1](#)
- [5] Samuel Burri, Yuki Maruyama, Xavier Michalet, Francesco Regazzoni, Claudio Bruschini, and Edoardo Charbon. Architecture and applications of a high resolution gated SPAD image sensor. *Optics Express*, 22(14):17573, Jul 2014. [6](#)
- [6] Mauro Buttafava, Gianluca Boso, Alessandro Ruggeri, Alberto Dalla Mora, and Alberto Tosi. Time-gated single-photon detection module with 110 ps transition time and up to 80 mhz repetition rate. *Review of Scientific Instruments*, 85(8):083114, 2014. [3](#), [6](#), [7](#)
- [7] Mauro Buttafava, Jessica Zeman, Alberto Tosi, Kevin Elieiri, and Andreas Velten. Non-line-of-sight imaging using a time-gated single photon avalanche diode. *Optics express*, 23(16):20997–21011, 2015. [1](#), [3](#)
- [8] P. B. Coates. The correction for photon ‘pile-up’ in the measurement of radiative lifetimes. *Journal of Physics E: Scientific Instruments*, 1(8):878, 1968. [3](#), [4](#), [5](#)
- [9] Alessandro Cominelli, Giulia Acconcia, Pietro Peronio, Massimo Ghioni, and Ivan Rech. High-speed and low-distortion solution for time-correlated single photon counting measurements: A theoretical analysis. *Review of Scientific Instruments*, 88(12):123701, Dec 2017. [2](#), [3](#)
- [10] Constantinos Daskalakis, Gautam Kamath, and Christos Tzamos. On the structure, covering, and learning of poisson multinomial distributions. *arXiv preprint*, 2015. [4](#)
- [11] Neale A. W. Dutton, Istvan Gyongy, Luca Parmesan, Salvatore Gnechchi, Neil Calder, Bruce R. Rae, Sara Pellegrini, Lindsay A. Grant, and Robert K. Henderson. A SPAD-based QVGA image sensor for single-photon counting and quanta imaging. *IEEE Transactions on Electron Devices*, 63(1):189–196, Jan 2016. [2](#)
- [12] Abhiram Gnanasambandam, Omar Elgendy, Jiaju Ma, and Stanley H Chan. Megapixel photon-counting color imaging using quanta image sensor. *Optics Express*, 27(12):17298–17310, 2019. [2](#)
- [13] Anant Gupta, Atul Ingle, Andreas Velten, and Mohit Gupta. Photon-flooded single-photon 3d cameras. In *Proceedings of the IEEE CVPR*, pages 6770–6779, 2019. [3](#), [5](#), [7](#), [8](#)
- [14] Felix Heide, Steven Diamond, David Lindell, and Gordon Wetzstein. Sub-picosecond photon-efficient 3d imaging using single-photon sensors. *Scientific Reports*, 8(1), Dec 2018. [2](#), [3](#), [4](#), [7](#), [8](#)
- [15] Atul Ingle, Andreas Velten, and Mohit Gupta. High flux passive imaging with single photon sensors. In *Proc. CVPR*, June 2019. [1](#)
- [16] Sebastian Isbaner, Narain Karedla, Daja Ruhlandt, Simon Christoph Stein, Anna Chizhik, Ingo Gregor, and Jörg Enderlein. Dead-time correction of fluorescence lifetime measurements and fluorescence lifetime imaging. *Optics express*, 24(9):9429–9445, 2016. [3](#)
- [17] Peter Kapusta, Michael Wahl, and Rainer Erdmann. Advanced Photon Counting Applications, Methods, Instrumentation. *Springer Series on Fluorescence*, 15, 2015. [2](#), [3](#)
- [18] Myung-Jae Lee and Edoardo Charbon. Progress in single-photon avalanche diode image sensors in standard cmos: From two-dimensional monolithic to three-dimensional-stacked technology. *Japanese Journal of Applied Physics*, 57(10):1002A3, 2018. [2](#)
- [19] David Lindell, Matthew O’Toole, and Gordon Wetzstein. Single-Photon 3D Imaging with Deep Sensor Fusion. *ACM Trans. Graph. (SIGGRAPH)*, 37(4), 2018. [2](#), [8](#)
- [20] Jiaju Ma, Saleh Masoodian, Dakota A Starkey, and Eric R Fossum. Photon-number-resolving megapixel image sensor at room temperature without avalanche gain. *Optica*, 4(12):1474–1481, 2017. [2](#)
- [21] Matthew O’Toole, Felix Heide, David Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing transient images from single-photon sensors. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2289–2297, July 2017. [3](#)
- [22] Matthew O’Toole, David B. Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nature*, 555:338–341, Mar 2018. [1](#)
- [23] Matthias Patting, Paja Reisch, Marcus Sackrow, Rhys Dowler, Marcelle Koenig, and Michael Wahl. Fluorescence decay data analysis correcting for detector pulse pile-up at very high count rates. *Optical engineering*, 57(3):031305, 2018. [2](#)
- [24] Agata M. Pawlikowska, Abderrahim Halimi, Robert A. Lamb, and Gerald S. Buller. Single-photon three-dimensional imaging at up to 10 kilometers range. *Optics Express*, 25(10):11919, May 2017. [3](#)
- [25] Adithya K Pediredla, Aswin C Sankaranarayanan, Mauro Buttafava, Alberto Tosi, and Ashok Veeraraghavan. Signal processing based pile-up compensation for gated single-photon avalanche diodes. *arXiv preprint arXiv:1806.07437*, 2018. [2](#), [3](#), [4](#), [5](#)
- [26] Sara Pellegrini, Gerald S Buller, Jason M Smith, Andrew M Wallace, and Sergio Cova. Laser-based distance measurement using picosecond resolution time-correlated single-photon counting. *Measurement Science and Technology*, 11(6):712, 2000. [3](#)
- [27] Joshua Rapp and Vivek Goyal. A few photons among many: Unmixing signal and noise for photon-efficient active imaging. *IEEE Transactions on Computational Imaging*, 3(3):445–459, Sept 2017. [2](#)
- [28] Joshua Rapp, Yanting Ma, Robin Dawson, and Vivek K Goyal. Dead time compensation for high-flux ranging. *arXiv preprint arXiv:1810.11145*, 2018. [2](#), [3](#), [7](#)

- [29] Daniel Reilly and Gregory Kanter. High speed lidar via GHz gated photon detector and locked but unequal optical pulse rates. *Optics Express*, 22(13):15718, jun 2014. [3](#)
- [30] Alexis Rochas. *Single Photon Avalanche Diodes in CMOS Technology*. PhD thesis, EPFL, 2003. [1](#)
- [31] Alberto Tosi, Alberto Dalla Mora, Franco Zappa, Angelo Gulinatti, Davide Contini, Antonio Pifferi, Lorenzo Spinelli, Alessandro Torricelli, and Rinaldo Cubeddu. Fast-gated single-photon counting technique widens dynamic range and speeds up acquisition time in time-resolved measurements. *Optics Express*, 19(11):10735, May 2011. [3](#)
- [32] Arin Can Ulku, Claudio Bruschini, Ivan Michel Antolovic, Yung Kuo, Rinat Ankri, Shimon Weiss, Xavier Michalet, and Edoardo Charbon. A 512x512 spad image sensor with integrated gating for widefield flim. *IEEE Journal of Selected Topics in Quantum Electronics*, 25(1):1–12, Jan 2019. [2](#), [3](#)
- [33] Michael Wahl. Time-Correlated Single Photon Counting, 2016. [3](#)
- [34] Xue Feng Wang, Teruo Uchida, David M. Coleman, and Shigeo Minami. A two-dimensional fluorescence lifetime imaging system using a gated image intensifier. *Applied Spectroscopy*, 45(3):360–366, mar 1991. [3](#)
- [35] Chao Zhang, Scott Lindner, Ivan Antolovic, Martin Wolf, and Edoardo Charbon. A CMOS SPAD Imager with Collision Detection and 128 Dynamically Reallocating TDCs for Single-Photon Counting and 3D Time-of-Flight Imaging. *Sensors*, 18(11):4016, Nov 2018. [3](#)