

Passive Ultra-Wideband Single-Photon Imaging

Mian Wei^{1*} Sotiris Nousias^{1*} Rahul Gulve² David B. Lindell¹ Kiriakos N. Kutulakos¹

¹Dept. of Computer Science, University of Toronto

²Dept. of Electrical and Computer Engineering, University of Toronto

Abstract

We consider the problem of imaging a dynamic scene over an extreme range of timescales simultaneously—seconds to picoseconds—and doing so passively, without much light, and without any timing signals from the light source(s) emitting it. Because existing flux estimation techniques for single-photon cameras break down in this regime, we develop a *flux probing theory* that draws insights from stochastic calculus to enable reconstruction of a pixel’s time-varying flux from a stream of monotonically-increasing photon detection timestamps. We use this theory to (1) show that passive free-running SPAD cameras have an attainable frequency bandwidth that spans the entire DC-to-31 GHz range in low-flux conditions, (2) derive a novel Fourier-domain flux reconstruction algorithm that scans this range for frequencies with statistically-significant support in the timestamp data, and (3) ensure the algorithm’s noise model remains valid even for very low photon counts or non-negligible dead times. We show the potential of this asynchronous imaging regime by experimentally demonstrating several never-seen-before abilities: (1) imaging a scene illuminated simultaneously by sources operating at vastly different speeds without synchronization (bulbs, projectors, multiple pulsed lasers), (2) passive non-line-of-sight video acquisition, and (3) recording *ultra-wideband video*, which can be played back later at 30 Hz to show everyday motions—but can also be played a billion times slower to show the propagation of light itself.

1. Introduction

A basic rule of thumb in high-speed imaging is that *speed needs light*: the faster a scene changes, the more light we need to image it accurately without excessive noise or motion blur. Over the decades, high-speed light sources [1], fast cameras [2–4], and depth sensors [5, 6] have made it possible to image dynamic phenomena occurring in ever-smaller time intervals with the help of actively-controlled light sources and *synchronization*: to collect enough light,

the same picosecond- or nanosecond-scale event may be imaged millions of times by operating a camera and a source in lockstep, at MHz repetition rates or more.

Acquiring videos of ultrafast phenomena this way—from imaging light in flight [7] to fast biological processes [8]—is now quite common. Unfortunately, while these techniques do capture ultrafast events, they cannot simultaneously capture *slower* ones too: time wraps at the sync period, blurring out anything occurring over longer timespans.

But how do we image highly dynamic scenes—both slow and ultrafast—*passively*, without any light sources under our control, no synchronization, and not much light? Very little is known about this problem because existing models for passive low-light imaging [9–13] break down at timescales much shorter than the timespan between photon arrivals. As a result, ultrafast imaging in low light has remained beyond the reach of passive methods.

In this work we seek to bridge these two regimes, active and passive, by revisiting the need for synchronization when imaging ultrafast scenes in low light. Working from first principles, we develop a novel theory of passive single-photon imaging that is specifically designed to eliminate synchronization between a camera and the light sources in a scene. The only requirements are that (1) the camera’s pixels can detect and time-stamp individual photons and (2) their dead-time period does not impair detection of photons from one source significantly more than any other. In this imaging regime, each camera pixel time-stamps the photons it detects using an internal clock that follows the arrow of time, obviating the need for any external timing signals.

Our work is based on the observation that passive (sync-free) imaging is fundamentally *more* powerful than active imaging in such settings. Specifically, without the periodic timing signal from a light source, time never wraps at a sync period; ultrafast scenes can be imaged for arbitrarily long timespans; and flux variations that occur concurrently across 12 orders of magnitude in time (picoseconds to seconds)—and that involve many unknown sources—can be recorded with just one camera.

Because photon timestamps due to all light sources and all

* Joint first authors: {mianwei, sotiris}@cs.toronto.edu
Project website: <https://compimaging.dgp.toronto.edu/ultra-wideband>

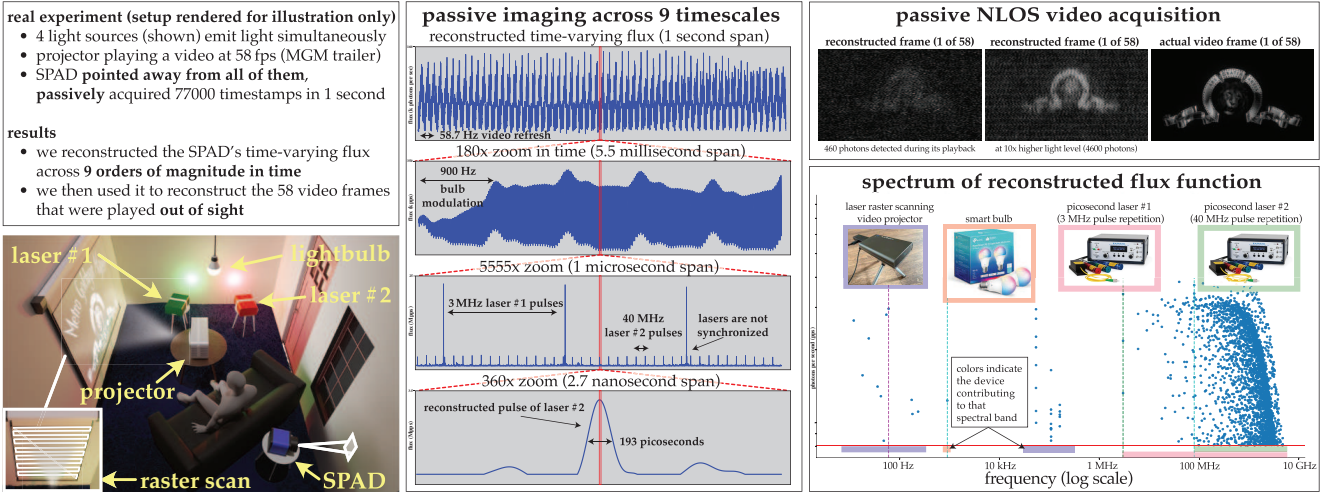


Figure 1: Passive ultra-wideband imaging with a single free-running SPAD pixel. **Left bottom & Center:** In a real, captured experiment, an unsynchronized single-photon avalanche diode (SPAD) passively records indirect light coming from multiple sources operating asynchronously from each other (unsynchronized picosecond lasers, projectors, *etc.*). See Figure 5 (row 2, middle) for actual scene. The incident flux exhibits simultaneous intensity variations with a bandwidth that spans roughly 9 orders of magnitude in frequency. **Bottom right:** Multiple concurrently-occurring phenomena in the flux function can be identified after acquisition: video flicker (58 Hz), the pulse-width modulation of an LED light bulb (900 Hz), a movie projected onto a nearby wall by a raster-scanning laser projector (up to 5 MHz), and two unsynchronized picosecond lasers (40 MHz–10 GHz). **Top right:** By reconstructing the time-varying flux function of the laser projector, video frames are reconstructed at 1280x720 resolution using roughly 450-4500 photons collected during each 1/58 s frame.

timescales are recorded concurrently, the choice of which timescale to show and which light source(s) to use for visual processing can be done *after* acquisition (Figure 1). Thus, just like light field cameras [14] enable post-capture refocusing in space, this new imaging regime enables post-capture refocusing *in time*—from transient to everyday timescales. We demonstrate this one-of-a-kind capability experimentally in Figure 5, where we use photon timestamp data captured by a free-running SPAD camera to play back video of a rapidly-spinning fan at both 1,000 and 250 billion frames per second. We call this novel regime *passive ultra-wideband imaging*.

The key challenge in this regime is how to reconstruct flux functions with an ultra-wide spectrum (DC to over 10 GHz) from a stream of photon timestamps that increase monotonically. To tackle it, we develop a *flux probing theory* that draws on results from stochastic calculus [15, 16] to relate the Fourier series decomposition of a time-varying flux function to the timestamp realizations of an underlying stochastic process [17, 18]. The mathematical underpinnings of our approach are grounded in statistics [19–21] and time series analysis [22–27], and similar methods have explored flux function estimation in optical communications [28–30].

Our work ties together several lines of prior research on “extreme” imaging, both passive and active. In passive settings, several techniques have recently been proposed for estimating flux from photon data [13, 31–33]. These rely on a variety of flux constancy assumptions and, as a result, are not applicable to the ultra-wideband regime we consider. In ac-

tive settings, single-photon imaging techniques have relied exclusively on sync-relative timestamps [34–38], where information about sub-MHz flux variations has already been lost. Aside from single-photon imaging, active ultrafast imaging techniques have also used heterodyning to measure flux at one specific modulation frequency [39–43]. These techniques have neither the light efficiency nor the ultra-wide bandwidth we demonstrate in this paper.

2. Passive Ultra-Wideband Imaging

Passive, sync-less imaging. We assume that the imaging system exerts no control over a scene’s appearance: the scene’s light source(s) can be natural, artificial, or both, and their number, operating principle, and time-varying properties are unknown and unconstrained (Figure 1). Importantly, we assume that no electronic timing signals, such as triggers or sync pulses, are received from any of them.

The time-varying flux function. Following standard radiometric conventions [44], we express incident light at a pixel as an unknown time-varying function $\varphi(t)$ that represents the pixel’s instantaneous flux at time $t \geq 0$. Our goal is to acquire a continuous representation of the flux function over a possibly unbounded acquisition interval $[0, t]$ (milliseconds, seconds or much longer). In the following we assume that $\varphi(t)$ is expressed in units of photons per second, is continuous, and has finite spectral support bounded by frequencies f_{\min} and f_{\max} .

Ultra-wideband flux. We seek to reconstruct flux functions that have ultra-wide bandwidth, *i.e.*, whose fre-

quency content spans the entire range from constant flux ($f_{\min} = 0$ Hz) to extreme time-of-flight timescales ($f_{\max} \geq 10$ GHz) [43, 45]. Moreover, we assume that no prior information is available about the spectrum of $\varphi(t)$.

Photon arrival model. Our work applies to the single-photon imaging regime, where the timespan between consecutive photon arrivals is not negligible. In this setting, $\varphi(t)$ is the rate function of an *inhomogeneous Poisson process* governing photon arrivals [29, 46]. The mean value of $\varphi(t)$ represents the average flux received over the observation interval $[0, t]$ in units of photons per second and its inverse, denoted by T_{avg} , is the average timespan between consecutive photon arrivals [33].

Low-flux photon detection model. Modern SPADs can detect and time-stamp the arrival of individual photons with extremely high temporal precision (typically tens of picoseconds [3, 47]). SPADs are not perfect detectors, however, as they exhibit four main non-idealities: quantum efficiency, dead time, timestamp quantization and jitter. *Quantum efficiency* refers to the pixel’s probability of actually detecting a photon when it is in its active state. This probability can be well below 1 depending on wavelength; since it can be thought of as scaling the flux function, we assume it is absorbed in $\varphi(t)$. After a photon detection, SPADs are blind to subsequent photon arrivals for an interval known as the *dead time*. Dead time can skew photon detection statistics quite significantly when photons arrive closely enough in time to fall within a SPAD’s dead-time window with high probability [32, 48, 49]. For simplicity, we focus on low-flux imaging in the main paper, where consecutive arrivals are spaced much farther apart than the SPAD’s dead time.¹ In this case, detections are governed by the same stochastic process that describes photon arrivals [50], with rate function $\varphi(t)$ and average timespan T_{avg} between detections.

Non-negligible dead time. When dead-time intervals become comparable to T_{avg} (or longer), photon detections are not Poisson because the detection of a photon may impact the detection of subsequent ones. We show in supplementary Section C that our stochastic calculus framework covers this case as well, enabling acquisition of $\varphi(t)$ by slightly amending the equations and algorithm in the main paper.

Timestamp model. Photon timestamps are subject to *quantization* from the time-to-digital conversion process and *jitter*, *i.e.*, instabilities in timing electronics. Both can be as low as a few picoseconds for SPADs in the visible range [51]. To simplify our analysis, we assume without loss of generality that timestamp resolution and timestamp accuracy are identical, so that the timestamps’ bin size Q accounts for jitter as well.²

¹For example, T_{avg} was nearly six times our SPAD’s dead time in the experiment of Figure 1 (1.3 microseconds versus 231 nanoseconds).

²When timestamp accuracy is worse than timestamp resolution, the

The stream of absolute detection timestamps. Since no external timing signal is available to serve as a reference, we assume that photon detection timestamps follow the arrow of time, increasing monotonically according to the SPAD’s internal clock. This results in a stream of timestamps $\mathcal{T} = (\tau_1, \dots, \tau_{N(t)})$, where τ_i is the elapsed time from the beginning of acquisition until the i -th photon detection, and $N(t)$ counts the total photons detected up to time t . We refer to timestamps τ_i as *absolute timestamps*. Absolute timestamps can be acquired by operating SPAD pixels in their “passive free-running” mode [32].

Paper roadmap. The remainder of the main paper is aimed at readers with no background in stochastic calculus, as implementation of our flux acquisition algorithms is straightforward and can be done without it. Readers may skip the formal definition of a martingale (supplement Section A) and focus on the comparisons in Section 2.1, the high-level description of our theoretical results in Section 3, and the algorithms in Section 4 (for low-flux settings) and Section C.1 (for non-negligible dead time). For readers familiar with statistical estimation, Section D includes simpler proofs of Section 3’s results, graciously provided by anonymous reviewer R1, which do not invoke stochastic calculus but apply only to low-flux settings where dead time is negligible.

2.1. Imaging with Photon Timestamps

The particular imaging regime outlined above generalizes two broad classes of single-photon imaging research, both of which use photon timestamps as their main input. We distinguish between the two by considering the relation between (a) the rate of photon detections and (b) the maximum reconstructible frequency in each case (Table 1).

Passive inter-photon imaging. Recent work in passive single-photon imaging has proposed treating the timespan between consecutive detections as a noisy sample of the scene’s flux [32, 33]. This implicitly assumes that flux does not vary in that timespan, which makes the rate of photon detections a (loose) upper bound on f_{\max} . As a result, passive low-flux imaging with SPADs has so far been restricted to slow speeds, with f_{\max} on the order of tens of kHz [13].

Active histogram-based imaging. Approaches that employ synchronized light sources [50, 56] occupy the other extreme of the frequency range. Their basic principle is to time-stamp detections relative to a sync signal of a known frequency f_{sync} , so that all timestamps wrap to the same brief interval $[0, 1/f_{\text{sync}}]$ regardless of the actual timespan between them. This forces photons to accumulate in a relatively small number of time bins—typically a few

timestamps’ effective number of bits is reduced [52]. Our system’s timestamps, for example, are quantized to 4 picoseconds but the standard deviation of jitter is 16 picoseconds, so we conservatively use $Q = 16$ for performance modeling purposes. Explicit treatment of jitter is beyond this paper’s scope (*e.g.*, jitter can actually *improve* timing resolution [53–55]).

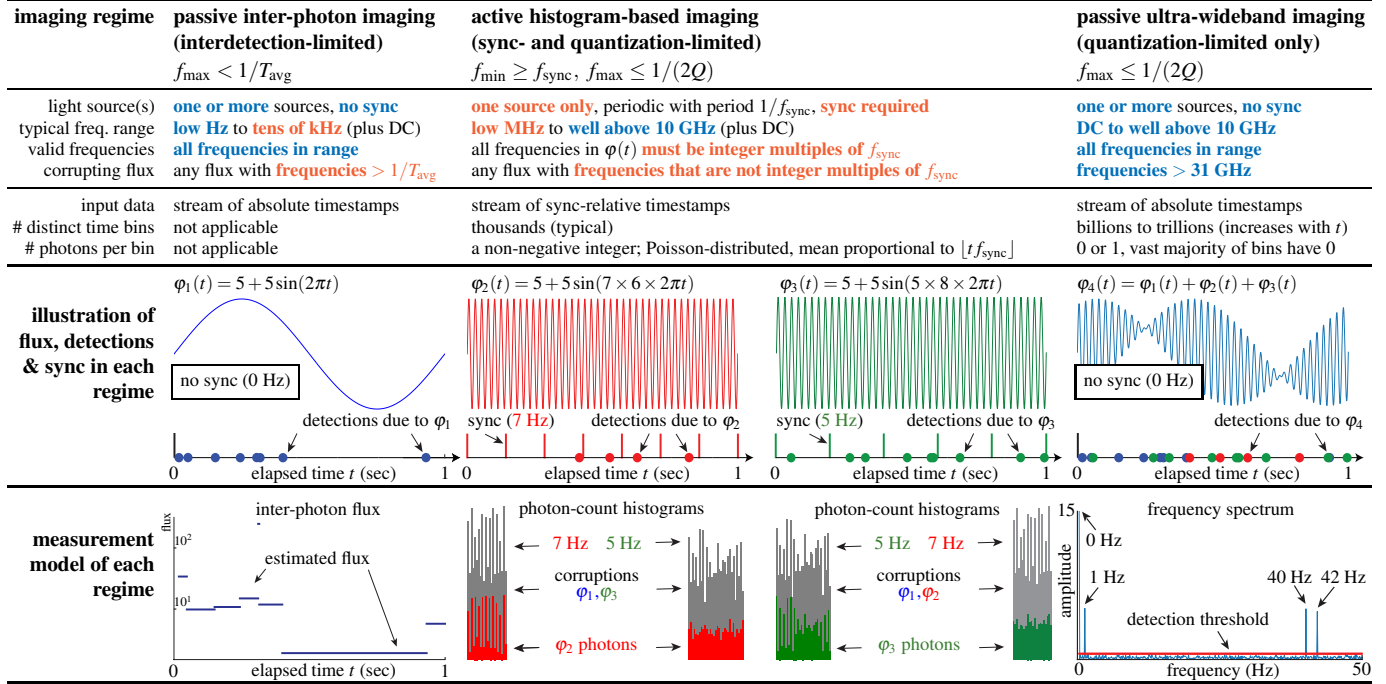


Table 1: Low-flux imaging with photon timestamps. **Left column:** Passive imaging has so far assumed that photons can be detected at a rate (much) greater than the highest flux frequency. **Middle column:** Active techniques are designed to handle the opposite case, where photon detections occur at a rate (much) lower than the flux function’s frequencies. If a light source’s frequency is not a multiple of f_{sync} (e.g., 5 Hz for φ_2 photons and 7 Hz for φ_3 photons in third row above), its photons will land in the wrong time bin. This contributes to noise instead of signal, *i.e.*, the histogram will be “flattened” (compare the two red and two green histograms, respectively, in third row). Moreover, even when f_{sync} is well-matched to one light source, other sources emitting at non-multiples of f_{sync} will corrupt the histogram (third row, gray histograms). Please refer to the supplementary video for another illustration of these effects. **Right column:** Our approach inherits the most important **features** of both regimes, without their **limitations**.

thousand—and yields a photon-count histogram [57] that is a noisy sampling of the scaled and time-wrapped flux function, $\lfloor t f_{\text{sync}} \rfloor \varphi(t - \lfloor t f_{\text{sync}} \rfloor / f_{\text{sync}})$. The maximum reconstructible frequency in this case is governed by the Nyquist theorem, not the photon interdetection time: a bin size of 16 picoseconds, for example, theoretically enables flux acquisition with f_{\max} equal to 31.25 GHz.³

Although this general approach achieves extremely high imaging speeds [12], its reliance on relative timestamps comes with a major constraint: the incident flux must also be a periodic function with a period equal to $1/f_{\text{sync}}$, to ensure that $\varphi(t)$ and its time-wrapped counterpart are identical. This can be trivially satisfied when the only light in the scene comes from a precisely-synchronized source (a pulsed laser, light-emitting diode, *etc.*), but flux variations due to other causes cannot be reconstructed. This includes variations caused by scene motion; light sources that emit at frequencies lower than f_{sync} ; and sources that emit at higher frequencies that are not integer multiples of f_{sync} . Photons from such sources result in histogram artifacts in the form

of additional photon noise [58], beat signals [45], or both.

Role of the sync frequency f_{sync} . Sync frequencies are typically in the low-MHz range in single-photon imaging applications that involve pulsed sources [35, 50, 59, 60]. This choice balances improved signal-to-noise ratio (a faster sync means more laser pulses, more photons detected in each histogram bin, and fewer time bins for them to accumulate in) against the likelihood of photons being missed due to dead time [48, 49, 58], or photons arriving “too late” because time has wrapped already [12, 61]. At such MHz sync frequencies, the memory and compute cost of reconstructing histograms from timestamps can be significant, prompting several recent schemes for just-in-time processing of (sync-relative) photon timestamps [37, 38, 62].

The passive ultra-wideband imaging regime. Intuitively, the regime we tackle in this paper can be understood as the limit case of photon histogramming, where the sync frequency is reduced *all the way to zero*. Specifically, as f_{sync} decreases, the interval $[0, 1/f_{\text{sync}}]$ increases; more histogram bins are needed to span it; fewer photons land into each bin; the time-wrapped flux function is able to represent variations that take place over longer timespans; and the space of reconstructible frequencies (*i.e.*, the integer mul-

³As a reference, 31.25 GHz is the -3 dB cut-off frequency of a Gaussian pulse with a full-width-at-half-max of 36 picoseconds.

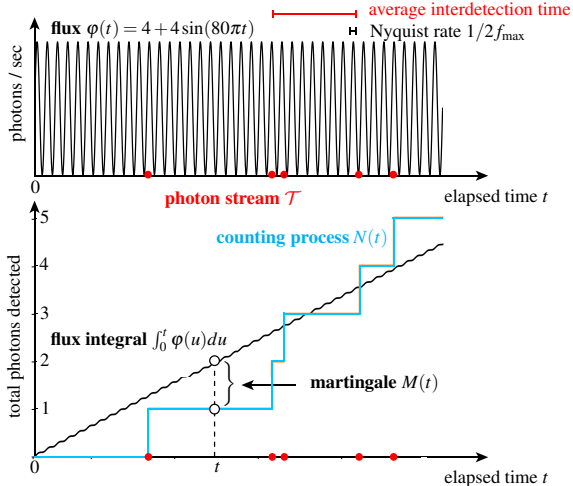


Figure 2: Relation between the stream of absolute timestamps, the counting process, the flux function and its integral. Timestamps are from a computational simulation of the inhomogeneous Poisson process with the flux function $\phi(t)$ indicated above.

triples of f_{sync}) expands. In the limit when f_{sync} is exactly zero, there is no sync at all, timestamps become absolute and every photon lands into its own unique time bin. Crucially, *all frequencies from DC up to the Nyquist limit—and from any light source—become potentially reconstructible.*

Mathematical modeling of this limit case, however, is non-trivial because the concept of a histogram breaks down: photons never accumulate, and the contents—0 or 1—of any given bin provide almost no information about the flux function.⁴ On the computational side, the entire acquisition interval $[0, t]$ is effectively partitioned into time bins at the SPAD’s timing resolution, so acquisitions of a second or more can potentially involve trillions of time bins (most of which are empty). Fortunately, as we show in the next section, both these challenges can be overcome by formulating flux reconstruction in terms of the *photon counting process*, which is not degenerate even when $f_{\text{sync}} = 0$.

3. Probing Flux Functions

Our approach establishes a direct mathematical link between (1) the stream of absolute timestamps detected at a pixel—however few or far apart they may be—and (2) the flux function that produced them. This link allows us to “probe” the Fourier spectrum of an unknown flux function across the entire DC-to-GHz range for frequencies that have statistically-significant support in the timestamp data. We introduce our flux probing theory below and address flux reconstruction in Section 4. For the sake of generality, we consider timestamps to be continuous-valued random variables and model quantization as part of our theory.

⁴More formally, the Poisson-distributed random variable associated with any given time bin has a mean that goes to zero as $t \rightarrow \infty$ [63].

The photon counting process. Even though a single absolute timestamp provides (almost) no information about the underlying flux function, the stream of absolute timestamps as a whole contains considerable information about it. The specific relation between the two comes from stochastic calculus [15]. Specifically, in the continuous-time domain, a stream \mathcal{T} of real-valued absolute timestamps provides a noisy “reconstruction” of the integral of $\phi(t)$ (Figure 2):

$$\underbrace{N(t)}_{\text{counting process}} = \underbrace{\int_0^t \phi(u) du}_{\text{flux integral up to time } t} + \underbrace{M(t)}_{\text{martingale noise}}. \quad (1)$$

The function $N(t)$ in Eq. (1) counts the photons received up to time t and is completely determined by \mathcal{T} ; formally, it is a *counting process* [15, 19]. Viewed from the perspective of histogram-based single-photon imaging (Table 1, middle), $N(t)$ is the continuous-time analog of the *cumulative photon-count histogram* over the interval $[0, t]$, for $f_{\text{sync}} = 0$. The function $M(t)$ in Eq. (1) is a continuous-time random process called a *martingale* [64] that can be thought of as a form of additive zero-mean noise.⁵

As can be seen from the example of Figure 2, a single random realization of the counting process (cyan curve) is a highly discontinuous function that, on first inspection, bears no resemblance to the flux integral it is supposed to approximate in Eq. (1). These discontinuities introduce dense, spurious frequencies in the Fourier-domain representation of $N(t)$ that do not exist in the actual flux integral.

3.1. Theory of Flux Probing

Our theoretical results use tools from stochastic calculus to address two basic questions. First, what is the highest possible frequency f_{\max} that can be recovered by a passive single-photon imaging system that outputs quantized absolute timestamps? Second, for frequencies within the attainable bandwidth, how can we derive a noise model that allows spurious frequencies to be efficiently detected and discarded, and the accuracy of real frequencies to be quantified as a function of the acquired timestamp stream? We summarize these results below and defer proofs to Section B.

The probing operation. Proposition 1 tells us that we can always probe the flux function to recover a (noisy) measurement of its inner product with practically any other function. Moreover, probing is efficient to compute from the timestamp stream and can be thought of as a continuous-time and sync-free generalization of compressive acquisition schemes for conventional photon-counting histograms [37, 38, 62]. In particular, let \mathcal{T} be the stream of real-valued absolute timestamps up to time t and let $p(t)$ be

⁵One example of a martingale is an unbiased random walk. Like $N(t)$ in Eq. (1), many other increasing stochastic processes can be expressed as the sum of a deterministic increasing function and a martingale [15]. See supplement Section A for the formal definition of a martingale.

an arbitrary known and square-integrable function:

Proposition 1 (Flux Probing Equation). The inner product of the probing function $p(t)$ and the unknown flux function $\varphi(t)$ over the time interval $[0, t]$ satisfies the relation

$$p(\mathcal{T}) = \langle p, \varphi \rangle + M_p(t) \quad (2)$$

where $p(\mathcal{T})$ are “probing measurements” which sum the values of the probing function at the absolute timestamps

$$p(\mathcal{T}) \stackrel{\text{def}}{=} \sum_{\tau \in \mathcal{T}} p(\tau), \quad (3)$$

$M_p(t)$ is a martingale, and the inner product is defined as $\int_0^t p(u)\varphi(u) du$.

Fundamental limit on bandwidth. In the special case of probing with the Fourier basis functions $p_f(t) = e^{-j2\pi ft}$, we prove in supplement Section B that probing with $f > 1/2Q$ yields aliased measurements that “wrap around” the frequency spectrum and are identical to—and indistinguishable from—lower-frequency measurements:

Proposition 2. Given timing resolution Q , the maximum recoverable frequency is $f_{\max} = \frac{1}{2Q}$.

Intuitively, Proposition 2 says that flux frequencies above $1/2Q$ are unrecoverable regardless of whether we detect a few photons or a million.

Noise model. Our model accounts for the inhomogeneous Poisson nature of photon detections and treats the general case of real-valued timestamps. The model is valid for arbitrary flux levels within the low-flux regime and, as we show in supplement Section G.4, it remains valid for low-count acquisitions (e.g., as few as ten photons). More specifically, Proposition 3 tells us that the noise in probing measurements has a distribution that can be estimated from the timestamp stream through another probing operation. Thus, probing gives the means both to observe a flux function and to quantify the uncertainty of that observation:

Proposition 3 (Distribution of Probing Measurements). The probing measurements $p(\mathcal{T})$ are approximately normally distributed with mean $\langle p, \varphi \rangle$ and variance $\langle p^2, \varphi \rangle$.

Fourier probing noise. Finally, Corollaries 1 and 2 allow us to quantify the accuracy by which specific flux frequencies can be estimated from a given timestamp stream:

Corollary 1 (Distribution of Fourier Probing). The Fourier probing measurements $p_f(\mathcal{T})$ approximately follow a complex normal distribution with mean and covariance matrix

$$\mu = \begin{bmatrix} \langle \cos(2\pi ft), \varphi(t) \rangle & \langle -\sin(2\pi ft), \varphi(t) \rangle \\ \langle \cos^2(2\pi ft), \varphi(t) \rangle & 0 \end{bmatrix} \quad (4)$$

$$\Sigma = \begin{bmatrix} \langle \cos^2(2\pi ft), \varphi(t) \rangle & 0 \\ 0 & \langle \sin^2(2\pi ft), \varphi(t) \rangle \end{bmatrix}. \quad (5)$$

Corollary 2 (Distribution of Fourier Probing Energy). The normalized energy of the Fourier basis probing measurements

$$p_f^{\mathcal{E}}(\mathcal{T}) \stackrel{\text{def}}{=} \text{Re} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{1,1}}} \right]^2 + \text{Im} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{2,2}}} \right]^2 \quad (6)$$

follows a non-central χ^2 distribution with 2 degrees of freedom and non-centrality parameter $\mu_1^2/\Sigma_{1,1} + \mu_2^2/\Sigma_{2,2}$.

In supplement Section B we show that unbiased estimators of the parameters of the above distributions can be obtained via probing.

Frequency detection. Given the estimated distribution of the Fourier Probing Energy, we derive the constant false alarm rate (CFAR) detector [65] (see supplement Section B) to identify and remove noisy frequencies based on a desired probability of false alarm α . False alarms occur when we keep f and $E[|p_f(\mathcal{T})|] = 0$; we remove f if $p_f^{\mathcal{E}}(\mathcal{T})$ is lower than $\text{CDF}_{\chi^2}^{-1}(1 - \alpha)$, derived from Corollary 2. Specifically, we detect frequencies for which

$$|p_f(\mathcal{T})|^2 \geq \text{CDF}_{\chi^2}^{-1}(1 - \alpha) \frac{N(t)}{2t^2}. \quad (7)$$

Note that (1) for a fixed α , the probability of detecting a frequency is proportional to the total number of photons detected, and (2) for flux functions dominated by a particular frequency such that $|p_{f_i}(\mathcal{T})|^2$ is large, $N(t)$ also tends to become proportionally larger, reducing the probability of detecting other frequencies.

Implications for passive single-photon imaging. Proposition 2 implies that rather than being a hindrance, synchless imaging with absolute timestamps confers an extreme bandwidth advantage to SPAD cameras: systems with 16-picosecond resolution, such as our own, can *simultaneously* acquire flux variations that span the entire DC-to-31 GHz range of frequencies, and that are due to *any number of unknown light sources* operating independently. This bandwidth is orders of magnitude broader than intensity cameras—SPADs or otherwise—were thought capable of acquiring directly [41–43, 45], i.e., without resorting to homodyne [43] or heterodyne [41, 42, 45] detection schemes. While Proposition 2 describes *reconstructability* (i.e., frequencies above the limit are unreconstructible), Proposition 3 and Eq. (7) provide insights about the *accuracy* and *detectability* of flux variations at different frequencies. Lastly, although we have not verified the theorized DC-to-31 GHz bandwidth experimentally due to unavailability of lasers that are fast enough, we show several real-world demonstrations of simultaneous DC-to-16.9 GHz imaging under very challenging low-flux conditions (see Figure 1, Section 5, and Section E). These validate our noise models in Eqs. (4)–(6) and (partially) confirm our theoretical bound.

```

procedure FLUXREC( $\mathcal{T}, t, f_{\max}, \alpha$ )
  // Frequency scanning.
   $\Delta f = 0.6/t$  (see supp. Section B)
   $\mathcal{F} = \text{freqs from } 0 \text{ to } f_{\max} \text{ with step } \Delta f$ 
  loop  $f \in \mathcal{F}$ 
     $p_f(\mathcal{T}) = (1/t) \sum_{\tau \in \mathcal{T}} e^{-j2\pi f \tau}$ 
  // Frequency detection.
   $\mathcal{F}_{\text{used}} = \emptyset$ 
  loop  $f \in \mathcal{F}$ 
     $A_f = |p_f(\mathcal{T})|$ ,  $\phi_f = \angle p_f(\mathcal{T})$ 
    reject  $f$  using CFAR (Eq.(7))
     $\mathcal{F}_{\text{used}} = \mathcal{F}_{\text{used}} \cup \{f\}$  if not rejected
  // Flux reconstruction.
   $\hat{\phi}(t) = \sum_{f \in \mathcal{F}_{\text{used}}} A_f \cos(2\pi f t + \phi_f)$ 

```

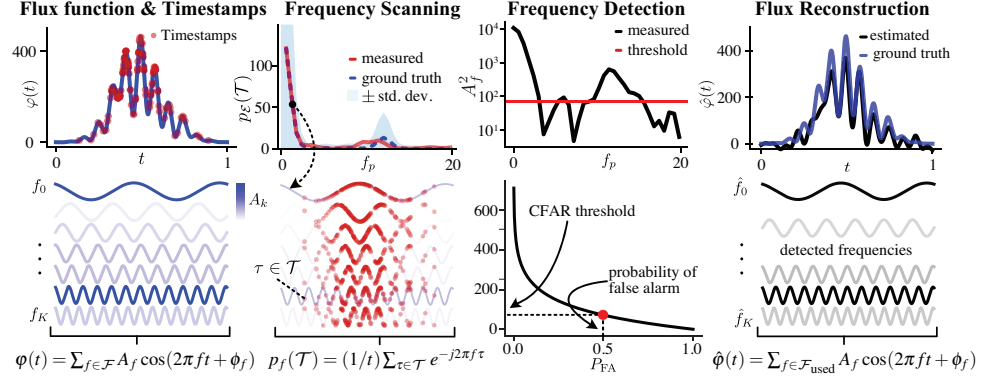


Figure 3: Overview of imaging by flux probing. **Left:** Flux reconstruction algorithm used in our experiments. **Right:** Visual illustration of the algorithm. *Flux function:* A flux function of finite spectral support can be expressed as a sum of sinusoids and produces a stream of absolute timestamps. *Frequency scanning:* We probe the flux function using a Fourier basis and measure the response at each frequency. *Frequency detection:* For each of the probed frequencies, we detect whether it contributes to the flux function if its corresponding amplitude is greater than a threshold (top) which is selected to achieve the desired probability of false alarm. (bottom). *Flux reconstruction:* Finally, we reconstruct a continuous-time flux function from the amplitudes and phases of detected frequencies.

4. Imaging by Flux Probing

Our probing theory leads directly to an algorithm that reconstructs the Fourier transform of the flux function by frequency-scanning the entire DC-to-GHz bandwidth (Figure 3). This algorithm is similar to the Fourier-domain histogramming technique used in conventional active settings [62], but also differs in two key respects: (1) it provides a principled way to estimate frequency uncertainty in an acquired timestamp stream, and (2) it enables tractable operation in a regime involving potentially billions of candidate frequencies—*e.g.*, a 1 Hz-resolution scan of DC-to-20 GHz—by rejecting spurious frequencies and reducing storage requirements. The frequency detection step uses the CFAR detector of Section 3, where we set α so that the expected number of false alarms is less than 1. Figure 3 includes a visual depiction of the flux reconstruction algorithm, along with a complete description in pseudocode.

5. Experiments

We validate our theory experimentally with (1) passive ultra-wideband sensing of both 1D intensity signals and 2D video signals ranging from DC to 16.9 GHz, (2) passive non-line-of-sight (NLOS) video via MHz-rate flux function reconstruction, (3) generalization to 2D SPAD arrays for high speed video, and (4) simulation-based quantitative evaluation of flux probing. We refer readers to the supplemental video and to supplement Sections E, F, and G for more experiments and implementation details.

Passive ultra-wideband sensing. We demonstrate recovering signals with frequencies spanning roughly 9 orders of magnitude from DC to 10 GHz (Figure 1). We place a single-pixel SPAD in the scene to capture flux variations from (1) pulse-width modulation of a light bulb (900 Hz),

(2) backscattered light from a raster-scanning laser projector (60 Hz–10 MHz), and (3) two unsynchronized picosecond lasers (40 MHz–10 GHz). Remarkably, the flux function is reconstructed from only 77,000 photon timestamps.⁶ We recover time-varying flux across billions of frequencies from this minuscule set of photons. Moreover, by employing a brighter and faster laser, we achieve passive DC-to-16.9 GHz sensing over room-size distances (Section E.2.2).

We also demonstrate ultra-wideband video (Figure 5, top) by raster scanning a scene in which a pulsed laser, with 20 MHz repetition and 80 ps FWHM, is diffused to illuminate a fan spinning at 54 Hz. We detect frequencies from DC to 10 GHz and render flux functions at 1 kfps and 250 Gfps, showing both the fan blades rotating and the propagation of the laser pulse. In contrast, conventional approaches reconstruct the scene at only one of the aforementioned frame rates, temporally blurring either slow or fast events. Furthermore, our method can render the flux at whatever timescale, essentially *freezing time at all timescales*. Note that because we only had access to a single pixel SPAD, the timestamps were collected by scanning across the field of view of the SPAD. To temporally align the flux functions *between pixels*, we use synchronization signals from both the laser and the fan. We emphasize that no synchronization signals were used to reconstruct the flux functions themselves—we are demonstrating a new capability of reconstructing the appearance of the scene *as it appeared during each laser pulse*—this is distinct from the use of synchronization and histogramming to estimate the *average appearance* of the flux function over time [7, 12, 66]. We also emphasize that the images are rendered by integrating the flux function over the exposure of each frame. As

⁶Conventional camera pixels collect a few thousand photons to return a single estimate of light intensity.




Figure 4: Simulation-based comparisons. **Left to right:** Laser pulse reconstructed by three methods from 2000 timestamps produced by a 20 MHz pulse train. Reconstruction from a 20 MHz train of much shorter pulses, using just 50 timestamps. Pulse reconstruction error as a function of the number of timestamps given as input. Reconstructed pulses from 100 realizations of a fifty-photon timestamp stream.

such, these images exhibit not only high dynamic range but their intensities are also expressed in physical units of photons, thereby ensuring radiometric calibration by nature.

We show another video example in Figure 5 (row 2), where the same picosecond laser illuminates a Coca-Cola bottle filled with water and a small amount of milk to scatter the light. Within the same scene, a compact fluorescent light-bulb (CFL) flickers at 120 Hz. We render videos at 10 kfps and 200 Gfps to visualize the CFL flicker and light pulses propagating through the bottle (Figure 5, rows 3–4). For the same reasons outlined in the previous paragraph, we use synchronization signals from the laser and the bulb.

Recovering passive NLOS videos. We demonstrate passive NLOS video reconstruction using light measured indirectly from a raster-scanning laser projector (see illustration in Figure 1 and photo in Figure 5, row 2). The SPAD observes a *single* point on a diffuse box during the projector beam’s raster scan, collecting light that bounced twice (*i.e.*, diffuser→diffuse box→SPAD). This configuration is analogous to dual photography [67]. By reconstructing the 1D flux function over a one-second span, we recover the video being played. We show results for the multi-illumination setting of Figure 1 and for a projector-only setting. In the latter case, we recover fine details of each video frame (Figure 5, rows 5–6) even though only 3000 photons were collected on average during a frame’s $1/58$ s raster scan.

Probing with SPAD arrays. Our method can be applied off the shelf to data from 2D SPAD sensors. To demonstrate this, we compare to Seets et al. [13] who recover high-speed video with a 32×32 SPAD array. They assume flux is piecewise constant and identify contiguous sets of timestamps with the same flux. Photon inter-arrivals are then averaged to obtain a *single* flux estimate per set (Figure 5, bottom right). In contrast, we recover a time-varying flux function by probing (Figure 5, bottom left). Because the sensor outputs just one timestamp per 20 microseconds for each pixel—a dead time too long to ignore even at relatively low flux levels—we probe using the generalized algorithm of Section C.1. This yields a periodic flux function truer to the rotating fan’s motion, whereas periodicity and high-frequency variations are lost by the method in [13].

Simulations. Lastly, we consider reconstruction of a flux function corresponding to 20 MHz pulse trains from an ultrafast laser with frequency support of 12.5 GHz and 125 GHz, respectively, *i.e.*, up to the theoretically-attainable limit of a jitter-less SPAD with 4 ps timestamp quantization (Proposition 2). Figure 4 compares the result of three methods: (1) conventional photon-count histogramming which requires synchronization, (2) our sync-less reconstruction algorithm in Figure 3, and (3) “oracle-based” flux probing, which probes the ground-truth set of frequencies instead of relying on frequency scanning. As can be seen, probing can recover pulses up to the theoretical limit from just 50 timestamps and, despite being passive, outperforms histogramming considerably as photon counts increase. Please see Section G for a more detailed quantitative evaluation.

6. Concluding Remarks

The sheer amount of data involved in probing timestamp streams cannot be ignored: even a single pixel can output tens of thousands of timestamps per second in low light, and our ultrawide-bandwidth results require probing billions of frequencies. Sketching [38] and Non-Uniform FFT [68] may offer ways forward but major challenges remain. That said, we believe that passive acquisition and processing of timestamp streams from free-running SPADs opens new directions in dynamic imaging: completely unsynchronized, single-shot observations of ultrafast phenomena with multiple light sources across different timescales; passive depth imaging using uncooperative, environmental light sources; compressive ultrafast video recording from sparse timestamps; temporal “microscopes” that allow monitoring intensity fluctuations across timescales spanning the nine-plus orders of magnitude (*i.e.*, DC to 31 GHz) theoretically captured by SPADs; and more. We are thus looking forward to more advances on these remarkable sensors.

Acknowledgements We thank the anonymous reviewers, and R1 in particular, for their invaluable comments and suggestions. We also thank Louis Zhang for creating an interactive tool for visualizing and sampling flux integrals efficiently using OpenCL, and John Hancock for configuring the servers used in our experiments. Lastly, KNK and DBL gratefully acknowledge the support of NSERC under the RGPIN and RTI programs.

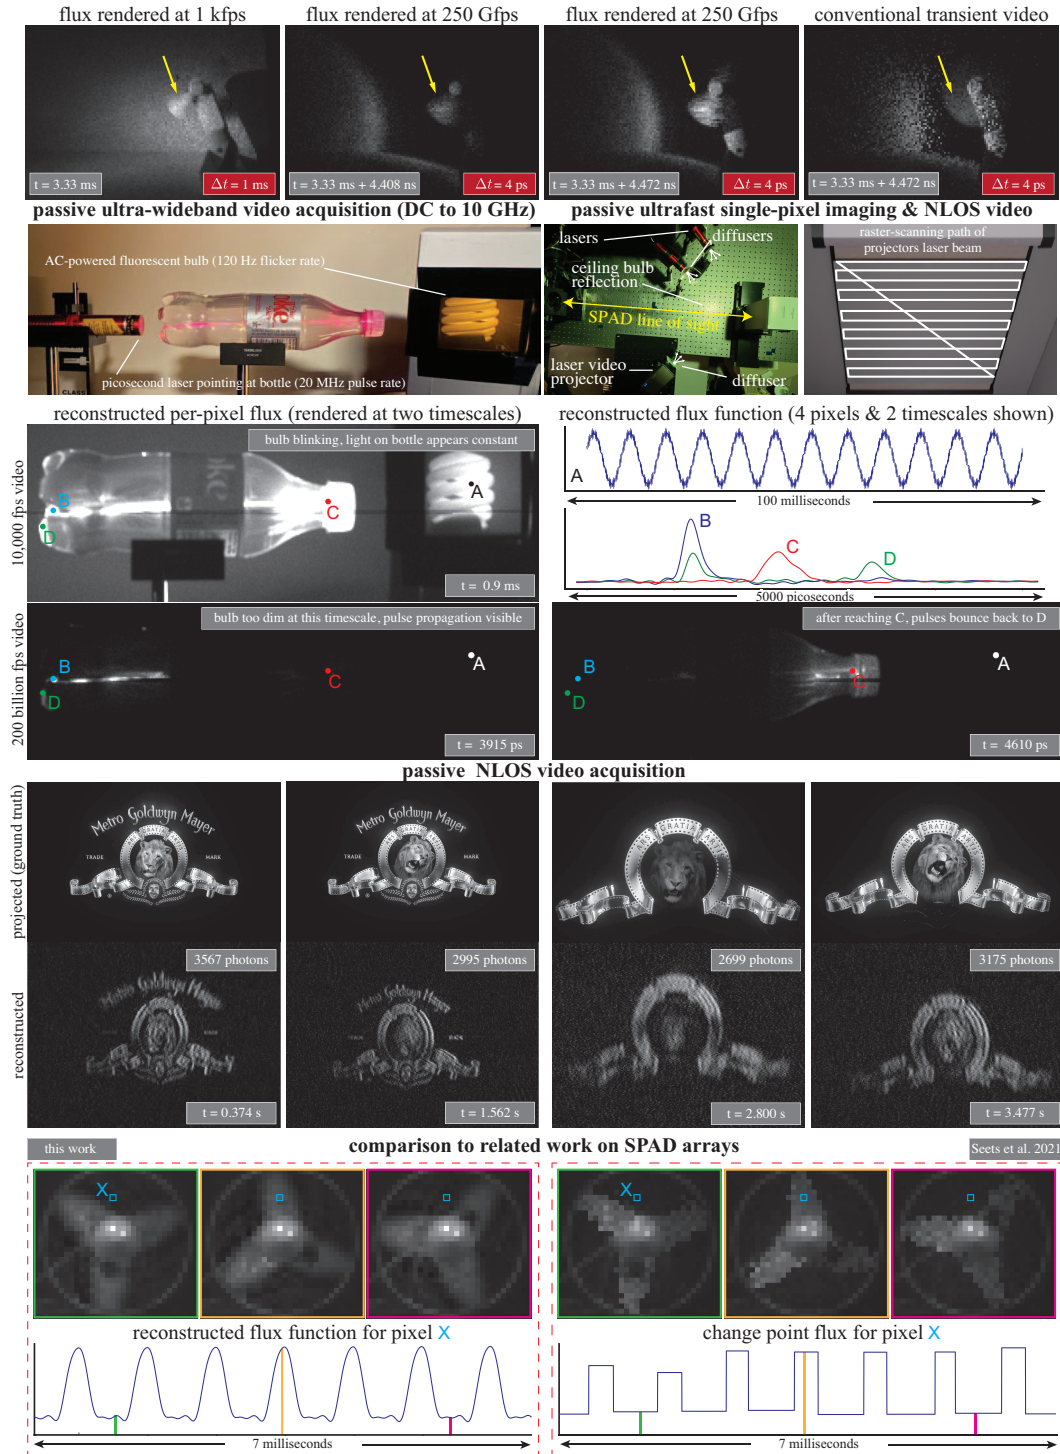


Figure 5: Passive ultra-wideband imaging experiments. Row 1: Simultaneous recovery of spinning fan and light propagation (see text for experiment setup). Rendering the flux function at 1 kfps reveals motion of the fan (yellow arrow), but light propagation is invisible; at 250 Gfps light propagation is visible, and the fan freezes. Conventional histogramming (right) synced to the laser fails to recover the (unsynced) fan rotation. **Row 2:** Experimental setup for ultra-wideband video acquisition (left) and NLOS video imaging in a scene with multiple light sources (right). **Rows 3–4:** Flux function images at two timescales (left, bottom right) and for different points in the scene illuminated by a pulsed laser and flickering light bulb (top right). The three peaks at B, C, and D correspond to a light pulse entering the bottle (B) propagating to the cap (C), and reflecting back (D). **Rows 5–6:** Passive NLOS acquisition using a raster-scanning laser projector with ground truth and reconstructed frames. **Rows 7–8:** We use SPAD array data from Seets et al. [13] to reconstruct per-pixel flux (left). The common assumption of piecewise constant flux (right) does not hold even for this simple scene of a rotating fan.

References

- [1] C. Rullière, *Femtosecond Laser Pulses*. Springer New York, 2005.
- [2] D. J. Bradley, B. Liddy, and W. E. Sleat, “Direct linear measurement of ultrashort light pulses with a picosecond streak camera,” *Opt. Commun.*, vol. 2, no. 8, pp. 391–395, 1971.
- [3] F. Zappa, S. Tisa, A. Tosi, and S. Cova, “Principles and features of single-photon avalanche diode arrays,” *Sens. Actuators A: Phys.*, vol. 140, no. 1, pp. 103–112, 2007.
- [4] L. Gao, J. Liang, C. Li, and L. V. Wang, “Single-shot compressed ultrafast photography at one hundred billion frames per second,” *Nature*, vol. 516, no. 7529, pp. 74–77, 2014.
- [5] C. Niclass, A. Rochas, P.-A. Besse, and E. Charbon, “Design and characterization of a CMOS 3-D image sensor based on single photon avalanche diodes,” *IEEE J. Solid-State Circuits*, vol. 40, no. 9, pp. 1847–1854, 2005.
- [6] R. Schwarte, Z. Xu, H.-G. Heinol, J. Olk, R. Klein, B. Buxbaum, H. Fischer, and J. Schulte, “New electro-optical mixing and correlating sensor: Facilities and applications of the photonic mixer device (PMD),” in *Proc. SPIE*, vol. 3100, pp. 245–253, 1997.
- [7] G. Gariepy, N. Krstajić, R. Henderson, C. Li, R. R. Thomson, G. S. Buller, B. Heshmat, R. Raskar, J. Leach, and D. Faccio, “Single-photon sensitive light-in-flight imaging,” *Nat. Commun.*, vol. 6, no. 6021, 2015.
- [8] W. Becker, “Fluorescence lifetime imaging—techniques and applications,” *J. Microsc.*, vol. 247, no. 2, pp. 119–136, 2012.
- [9] E. R. Fossum, J. Ma, S. Masoodian, L. Anzagira, and R. Zizza, “The quanta image sensor: Every photon counts,” *Sensors*, vol. 16, no. 8, 2016.
- [10] K. Morimoto, A. Ardelean, M.-L. Wu, A. C. Ulku, I. M. Antolovic, C. Bruschini, and E. Charbon, “Megapixel time-gated SPAD image sensor for 2D and 3D imaging applications,” *Optica*, vol. 7, no. 4, pp. 346–354, 2020.
- [11] Y. Chi, A. Gnanasambandam, V. Koltun, and S. H. Chan, “Dynamic low-light imaging with quanta image sensors,” in *Proc. Eur. Conf. on Computer Vision (ECCV)*, pp. 122–138, Springer, 2020.
- [12] M. O’Toole, F. Heide, D. B. Lindell, K. Zang, S. Diamond, and G. Wetzstein, “Reconstructing transient images from single-photon sensors,” in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 1539–1547, 2017.
- [13] T. Seets, A. Ingle, M. Laurenzis, and A. Velten, “Motion adaptive deblurring with single-photon cameras,” in *Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV)*, pp. 1944–1953, 2021.
- [14] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, “Light field photography with a hand-held plenoptic camera,” *Technical Report*, 2005.
- [15] J. L. Doob, *Stochastic Processes*. Wiley New York, 1953.
- [16] S. N. Cohen and R. J. Elliott, *Stochastic Calculus and Applications*, vol. 2. Birkhäuser New York, 2015.
- [17] N. Chen, D. K. K. Lee, and S. N. Negahban, “Super-resolution estimation of cyclic arrival rates,” *Ann. Stat.*, vol. 47, no. 3, pp. 1754–1775, 2019.
- [18] K.-S. Lii and M. Rosenblatt, “Estimation for almost periodic processes,” *Ann. Stat.*, vol. 34, no. 3, pp. 1115–1139, 2006.
- [19] M. S. Bartlett, “The spectral analysis of point processes,” *J. R. Stat. Soc., B: Stat. Methodol.*, vol. 25, no. 2, pp. 264–296, 1963.
- [20] D. Vere-Jones, “On the estimation of frequency in point-process data,” *J. Appl. Probab.*, vol. 19, pp. 383–394, 1982.
- [21] Á. Gajardo and H.-G. Müller, “Cox point process regression,” *IEEE Trans. Inf. Theory*, vol. 68, no. 2, pp. 1133–1156, 2022.
- [22] M. S. Bartlett, “Periodogram analysis and continuous spectra,” *Biometrika*, vol. 37, no. 1/2, pp. 1–16, 1950.
- [23] G. U. Yule, “On a method of investigating periodicities in disturbed series, with special reference to Wolfer’s sunspot numbers,” *Philos. Trans. R. Soc. Lond., A*, vol. 226, pp. 267–298, 1927.
- [24] M. G. Kendall, “On the analysis of oscillatory time-series,” *J. R. Stat. Soc.*, vol. 108, no. 1/2, pp. 93–141, 1945.
- [25] N. R. Lomb, “Least-squares frequency analysis of unequally spaced data,” *Astrophys. Space Sci.*, vol. 39, no. 2, pp. 447–462, 1976.
- [26] S. M. Ransom, S. S. Eikenberry, and J. Middleditch, “Fourier techniques for very long astrophysical time-series analysis,” *Astron. J.*, vol. 124, no. 3, pp. 1788–1809, 2002.
- [27] J. T. VanderPlas, “Understanding the Lomb–Scargle periodogram,” *Astrophys. J. Suppl. Ser.*, vol. 236, no. 1, pp. 1–28, 2018.

- [28] O. Macchi and B. Picinbono, "Estimation and detection of weak optical signals," *IEEE Trans. Inf. Theory*, vol. 18, no. 5, pp. 562–573, 1972.
- [29] J. E. Mazo and J. Salz, "On optical data communication via direct detection of light pulses," *Bell Syst. Tech. J.*, vol. 55, no. 3, pp. 347–369, 1976.
- [30] A. Komae, "Maximum likelihood and minimum mean squared error estimations for measurement of light intensity," in *Proc. Conf. on Information Sciences and Systems (CISS)*, pp. 1–6, 2010.
- [31] S. Ma, S. Gupta, A. C. Ulku, C. Bruschini, E. Carbon, and M. Gupta, "Quanta burst photography," *ACM Trans. Graph. (SIGGRAPH)*, vol. 39, no. 4, pp. 1–16, 2020.
- [32] A. Ingle, A. Velten, and M. Gupta, "High flux passive imaging with single-photon sensors," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 6753–6762, 2019.
- [33] A. Ingle, T. Seets, M. Buttafava, S. Gupta, A. Tosi, M. Gupta, and A. Velten, "Passive inter-photon imaging," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 8585–8595, 2021.
- [34] D. B. Lindell, M. O'Toole, and G. Wetzstein, "Towards transient imaging at interactive rates with single-photon detectors," in *Proc. IEEE Int. Conf. on Computational Photography (ICCP)*, pp. 1–8, 2018.
- [35] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. Wong, and J. H. Shapiro, "Photon-efficient imaging with a single-photon camera," *Nat. Commun.*, vol. 7, no. 12046, 2016.
- [36] Q. Sun, X. Dun, Y. Peng, and W. Heidrich, "Depth and transient imaging with compressive SPAD array cameras," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 273–282, 2018.
- [37] F. Gutierrez-Barragan, A. Ingle, T. Seets, M. Gupta, and A. Velten, "Compressive single-photon 3D cameras," in *Proc. IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 17854–17864, 2022.
- [38] M. P. Sheehan, J. Tachella, and M. E. Davies, "A sketching framework for reduced data transfer in photon counting lidar," *IEEE Trans. Comput. Imaging*, vol. 7, pp. 989–1004, 2021.
- [39] F. Heide, M. B. Hullin, J. Gregson, and W. Heidrich, "Low-budget transient imaging using photonic mixer devices," *ACM Trans. Graph. (SIGGRAPH)*, vol. 32, no. 4, pp. 1–10, 2013.
- [40] F. Li, F. Willomitzer, M. M. Balaji, P. Rangarajan, and O. Cossairt, "Exploiting wavelength diversity for high resolution time-of-flight 3D imaging," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 43, no. 7, pp. 2193–2205, 2021.
- [41] F. Li, F. Willomitzer, P. Rangarajan, M. Gupta, A. Velten, and O. Cossairt, "SH-ToF: Micro resolution time-of-flight imaging with superheterodyne interferometry," in *Proc. IEEE Int. Conf. on Computational Photography (ICCP)*, pp. 1–10, 2018.
- [42] S.-H. Baek, N. Walsh, I. Chugunov, Z. Shi, and F. Heide, "Centimeter-wave free-space neural time-of-flight imaging," *ACM Trans. Graph. (SIGGRAPH)*, vol. 42, no. 1, pp. 1–18, 2022.
- [43] M. Gupta, S. K. Nayar, M. B. Hullin, and J. Martin, "Phasor imaging: A generalization of correlation-based time-of-flight imaging," *ACM Trans. Graph. (SIGGRAPH)*, vol. 34, no. 5, pp. 1–18, 2015.
- [44] R. W. Boyd, *Radiometry and the Detection of Optical Radiation*. Wiley, 1983.
- [45] A. Kadambi and R. Raskar, "Rethinking machine vision time of flight with GHz heterodyning," *IEEE Access*, vol. 5, pp. 26211–26223, 2017.
- [46] S. M. Ross, *Stochastic Processes*. Wiley, 1983.
- [47] M. A. Itzler, R. Ben-Michael, C.-F. Hsu, K. Slomkowski, A. Tosi, S. Cova, F. Zappa, and R. Ispasoiu, "Single photon avalanche diodes (SPADs) for 1.5 μm photon counting applications," *J. Mod. Opt.*, vol. 54, no. 2-3, pp. 283–304, 2007.
- [48] A. K. Pediredla, A. C. Sankaranarayanan, M. Buttafava, A. Tosi, and A. Veeraraghavan, "Signal processing based pile-up compensation for gated single-photon avalanche diodes," *arXiv preprint arXiv:1806.07437*, 2018.
- [49] J. Rapp, Y. Ma, R. M. A. Dawson, and V. K. Goyal, "Dead time compensation for high-flux ranging," *IEEE Trans. Signal Process.*, vol. 67, no. 13, pp. 3471–3486, 2019.
- [50] G. A. Kirmani, *Computational Time-resolved Imaging*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [51] S. Riccardo, E. Conca, V. Sesta, A. Velten, and A. Tosi, "Fast-gated 16 \times 16 SPAD array with 16 on-chip 6 ps time-to-digital converters for non-line-of-sight imaging," *IEEE Sens. J.*, vol. 22, no. 17, pp. 16874–16885, 2022.
- [52] T. C. Carusone, D. Johns, and K. Martin, *Analog Integrated Circuit Design*. Wiley, 2011.

- [53] J. Rapp, R. M. A. Dawson, and V. K. Goyal, "Dithered depth imaging," *Opt. Express*, vol. 28, no. 23, pp. 35143–35157, 2020.
- [54] J. Rapp, R. M. A. Dawson, and V. K. Goyal, "Estimation from quantized Gaussian measurements: When and how to use dither," *IEEE Trans. Signal Process.*, vol. 67, no. 13, pp. 3424–3438, 2019.
- [55] A. Raghuram, A. Pediredla, S. G. Narasimhan, I. Gkioulekas, and A. Veeraraghavan, "STORM: Super-resolving transients by oversampled measurements," in *Proc. IEEE Int. Conf. on Computational Photography (ICCP)*, pp. 1–11, 2019.
- [56] X. Liu, I. Guillén, M. L. Manna, J. H. Nam, S. A. Reza, T. H. Le, A. Jarabo, D. Gutierrez, and A. Velten, "Non-line-of-sight imaging using phasor-field virtual wave optics," *Nature*, vol. 572, no. 7771, pp. 620 – 623, 2019.
- [57] Y. Chen, J. D. Müller, P. T. C. So, and E. Gratton, "The photon counting histogram in fluorescence fluctuation spectroscopy," *Biophys. J.*, vol. 77, no. 1, pp. 553–567, 1999.
- [58] A. Gupta, A. Ingle, and M. Gupta, "Asynchronous single-photon 3D imaging," in *Proc. IEEE/CVF Int. Conf. on Computer Vision (ICCV)*, pp. 7909–7918, 2019.
- [59] F. Heide, S. Diamond, D. B. Lindell, and G. Wetzstein, "Sub-picosecond photon-efficient 3D imaging using single-photon sensors," *Sci. Rep.*, vol. 8, no. 17726, 2018.
- [60] V. Zickus, M.-L. Wu, K. Morimoto, V. Kapitanov, A. Fatima, A. Turpin, R. Insall, J. Whitelaw, L. Machesky, C. Bruschini, D. Faccio, and E. Charbon, "Fluorescence lifetime imaging with a megapixel SPAD camera and neural network lifetime estimation," *Sci. Rep.*, vol. 10, no. 20986, 2020.
- [61] A. M. Pawlikowska, A. Halimi, R. A. Lamb, and G. S. Buller, "Single-photon three-dimensional imaging at up to 10 kilometers range," *Opt. Express*, vol. 25, no. 10, pp. 11919–11931, 2017.
- [62] X. Liu, S. Bauer, and A. Velten, "Phasor field diffraction based reconstruction for fast non-line-of-sight imaging systems," *Nat. Commun.*, vol. 11, no. 1645, 2020.
- [63] D. L. Snyder and M. I. Miller, *Random Point Processes in Time and Space*. Springer New York, 2012.
- [64] J. Ville, "Etude critique de la notion de collectif," *Bull. Am. Math. Soc.*, vol. 45, no. 11, p. 824, 1939.
- [65] L. L. Scharf and C. Demeure, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis*. Prentice Hall, 1991.
- [66] A. Velten, D. Wu, A. Jarabo, B. Masia, C. Barsi, C. Joshi, E. Lawson, M. Bawendi, D. Gutierrez, and R. Raskar, "Femto-photography: Capturing and visualizing the propagation of light," *ACM Trans. Graph. (SIGGRAPH)*, vol. 32, no. 4, pp. 1–8, 2013.
- [67] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch, "Dual photography," *ACM Trans. Graph. (SIGGRAPH)*, vol. 24, no. 3, p. 745–755, 2005.
- [68] A. H. Barnett, J. Magland, and L. A. Klinteberg, "A parallel nonuniform fast Fourier transform library based on an "exponential of semicircle" kernel," *SIAM J. Sci. Comput.*, vol. 41, no. 5, pp. C479–C504, 2019.