Opportunistic Single-Photon Time of Flight

Sotiris Nousias[1*] Mian Wei[1*] Howard Xiao[1] Maxx Wu[1] Shahmeer Athar[1] Kevin J. Wang[1]
Anagh Malik[1] David A. Barmherzig[1] David B. Lindell[1, 2] Kiriakos N. Kutulakos[1, 2]

[1]Dept. of Computer Science, University of Toronto [2]Vector Institute

Abstract

Scattered light from pulsed lasers is increasingly part of our ambient illumination, as many devices rely on them for active 3D sensing. In this work, we ask: can these "ambient" light signals be detected and leveraged for passive 3D vision? We show that pulsed lasers, despite being weak and fluctuating at MHz to GHz frequencies, leave a distinctive sinc comb pattern in the temporal frequency domain of incident flux that is specific to each laser and invariant to the scene. This enables their passive detection and analysis with a free-running SPAD camera, even when they are unknown, asynchronous, out of sight, and emitting concurrently. We show how to synchronize with such lasers computationally, characterize their pulse emissions, separate their contributions, and—if many are present—localize them in 3D and recover a depth map of the camera's field of view. We use our camera prototype to demonstrate (1) a first-of-its-kind visualization of asynchronously propagating light pulses from multiple lasers through the same scene, (2) passive estimation of a laser's MHz-scale pulse repetition frequency with mHz precision, and (3) mm-scale 3D imaging over room-scale distances by passively harvesting photons from two or more out-of-view lasers.

1. Introduction

Our environments are increasingly illuminated by faint but extremely fast light signals. Devices such as lidars on robots, cars, and drones [43, 45, 66], proximity sensors and time-of-flight (ToF) cameras in smartphones [18, 48, 75], along with wireless light-based communication systems [54, 62, 77], all emit light fluctuating at frequencies from MHz to GHz. This light is generally considered useful only to the emitting device, due to the precise synchronization and calibration required to detect it. ToF cameras, for instance, use a dedicated light source that emits periodic signals in sync with their sensor, enabling accurate measurement of the time delay between emission and detection [13, 41]. To nearby cameras, however, these signals blend into the ambient background, contributing to noise or, worse, interfering with their own measurements [30, 58].

In this work we ask: can such fast "ambient" light signals be automatically detected and leveraged for passive 3D vision? As a first step toward answering this question, we show that this is indeed possible to do with a time-resolving single-photon camera, as long as some of the ambient light comes from *pulsed lasers*—a light source technology used in smartphone flash lidars [48] and many advanced beam-scanning lidar systems [44].

Our focus is on scenarios where several independent and unknown pulsed lasers are present in a scene, each acting as a point light source. This occurs, for example, when many flash lidar devices are active, or when beams from different lasers illuminate distinct points in an environment. In such settings, each laser can be thought of as an ultra-fast strobe light, periodically flooding the scene with nanosecond- or sub-nanosecond pulses a few million times a second [60].

To passively harvest this light, we develop an incident flux model suited for cases where an unknown number of lasers emit pulses asynchronously from beyond the camera's line of sight, possibly amid other sources of strong ambient light. We then combine this model with passive ultra-wideband imaging [85] to automatically detect the lasers' presence; characterize their pulse emissions; separate their image contributions; and—if two or more lasers are detected—localize them in 3D while also recovering a depth map of the camera's field of view. The only requirements are that (1) the single-photon camera can passively timestamp photon arrivals concurrently and asynchronously across all its pixels, and (2) both the camera and the lasers remain stationary for the (sub-second) exposure duration. The resulting approach is *entirely passive*: it emits no light of its own, relying instead on detecting and repurposing light from other lasers operating nearby. We call this novel form of 3D imaging *opportunistic single-photon ToF*.

Opportunistic harvesting of various signals has been practiced since at least the 1930s, when reflections from ambient BBC radio signals were used to detect incoming aircraft [9, 40]. Since then, serendipitous ambient signals have enabled passive imaging across diverse domains, including coherent optical [21, 38], acoustic [8, 68, 78, 94], sonar [19, 20, 87], radar [27–29, 53, 79], wifi [1–4] and microwave [50, 92, 93, 95]. Our work is motivated by similar principles, but applied to the domain of incoherent optical ToF. In computer vision, high-speed cameras and motion enhancement techniques have been used to reveal hidden motions in the world [82, 83, 89], sense sound [22], and analyze mechanical vibrations [23, 72]. The visual signals we exploit here are considerably faster and much dimmer.

* Joint first authors: {sotiris,mianwei}@cs.toronto.edu
Project website: https://compimaging.dgp.toronto.edu/opportunistic

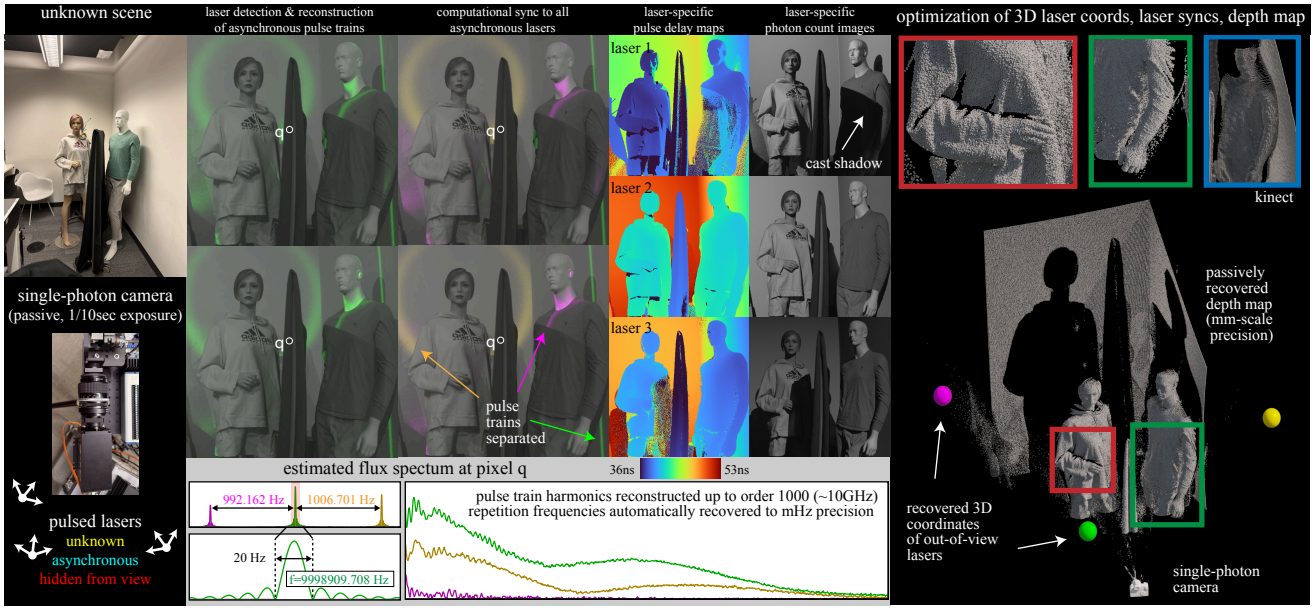

Figure 1. Harvesting scattered photons from unknown ultrafast sources for 3D sensing. **Left:** In this real experiment, three diffused lasers act as independent, unknown flash lidar devices, concurrently and asynchronously flood-illuminating a room with trains of picosecond-scale flashes from different positions. A SPAD camera, with no prior information about these light sources and no direct view of them, passively time-stamps photons with its internal clock. **Middle:** By collecting approximately 10K timestamps per pixel, the individual lasers are automatically detected; their pulse repetition frequencies are resolved to mHz precision; their pulse spectra reconstructed to GHz harmonics; their asynchronous pulse trains separated; and the relative delay of pulses from different lasers to each pixel computed to sub-nanosecond precision. **Right:** Using this information, the lasers are localized in 3D and the scene's depth map computed to mm-scale precision. See supplement Section H for an even more challenging setting with ceiling lights turned on (∼0.01 signal-to-background ratio).

Our approach tackles this challenge by formulating and solving two independent subproblems: a *single-photon imaging* task of converting raw photon detections into laser-specific pulse-delay maps relative to the camera's internal clock, and a purely *geometric* task of computing a depth map, 3D laser coordinates, and camera-to-laser clock offsets that are consistent with those pulse-delay maps.

On the geometric side, our work relates to techniques for radio wave device localization (*e.g.*, GPS [11, 91], wildlife tracking [10, 12, 51], cellphone localization [6, 52, 65, 74], indoor positioning [5, 24, 39, 47, 70]) which extract ToF information from ambient signals emitted by radio beacons. These techniques rely on a known beacon infrastructure, precise synchronization between beacons, and line-of-sight communication with them. In contrast, the laser "beacons" in our work are unknown, asynchronous, out of sight, and can change position independently from one exposure to the next. This results in a geometric problem that is considerably more general, and whose main goal is 3D imaging.

In the area of single-photon imaging, systems that combine a pulsed laser with single-photon avalanche diode (SPAD) pixels are widely used for lidar [26, 46, 81], light-in-flight imaging [60], non-line-of-sight imaging [61, 90], and biomedical imaging [35, 42]. These methods fundamentally depend on precise hardware synchronization between laser and detector [31], and on extensive modeling—as well as full control [37]—of the laser device [34, 63, 64]. Our work departs from this paradigm by relinquishing control over the laser, and modeling its operation on the fly from the ambient reflections it produces. We show that this leads to a new form of laser-camera synchronization that is *computational* rather than hardware based, and thus can be done after all photon data has been captured. A key advantage of this approach is that one camera can be (computationally) sync-locked to any number of lasers operating asynchronously, something impossible to achieve with a hardware sync. While we build on the ultra-wideband imaging framework of Wei *et al*. [85] to realize this capability for opportunistic ToF, that framework is agnostic to the characteristics of light sources in a scene and, crucially, is limited to cases where no strong ambient illumination (such as room lighting or sunlight) overpowers other sources.

Central to our approach is the observation that scattered light from pulsed lasers leaves a distinctive pattern in the temporal frequency domain of incident flux. This pattern—a *sinc comb*—is specific to each laser, invariant to the scene, and enables opportunistic, highly precise sync-locking even in the presence of far stronger light sources.

We use this core ability in our experimental system to demonstrate powerful new ultrafast imaging capabilities: (1) a first-of-its-kind visualization of asynchronously propagating light pulses from multiple lasers through the same scene; (2) passive estimation of a laser's MHz-scale pulse repetition frequency with mHz precision from scattered light alone; (3) mm-scale 3D imaging over room-scale distances by passively harvesting photons from out-of-view lasers; and (4) achieving these results under challenging conditions of strong ambient light (signal-to-background ratio ∼1%) and SPAD dead time effects.

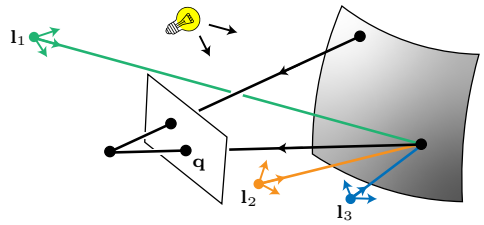

Figure 2. Viewing geometry of opportunistic ToF.

2. Opportunistic Time of Flight

Consider a camera that is imaging an unknown scene in the presence of an unknown number of pulsed lasers, and potentially other (non-ToF) ambient illumination. We model each laser as a point source that periodically emits pulses from a distinct, unknown 3D location (Figure 2). Both the camera and the lasers are assumed to be stationary during an exposure of duration t_{exp}, but can move freely and independently between exposures. Our objective is to automatically detect these lasers and use their scattered light to reconstruct a pixelwise depth map of the camera's 2D field of view.

Passive imaging conditions. We assume that neither the camera nor the individual lasers exchange timing signals of any kind, *i.e.*, they operate independently and asynchronously from each other. In this entirely passive setting, the camera must rely on its own internal clock to measure time. We use t to denote elapsed time according to the camera's clock, with $0 \leq t \leq t_{\text{exp}}$.

Our focus is on the general case where the sources are outside the camera's field of view and their operation—pulse profile, repetition frequency, sync signal, *etc.*—are all unknown. Under such conditions, the time of flight from individual sources to individual camera pixels cannot be measured as in conventional lidar; it must be estimated along with the 3D coordinates, pulse emission properties and sync timing of the lasers themselves.

Paper overview. We begin by introducing the basis of our approach, namely our frequency-domain model for incident flux under pulsed illumination (Section 3) and the framework we build upon to probe it passively with a single-photon camera (Section 4). Section 5 describes the core method of opportunistic ToF—detecting individual pulsed lasers and synching with them computationally—formulated as a sequence of flux-probing operations. Geometric optimization is briefly considered in Section 6.

3. Incident Flux Model for Pulsed Laser Light

The light signal transmitted from a pulsed laser to individual camera pixels carries rich information about the laser itself. This is because the laser's periodic emission creates a distinctive sinc comb pattern in the temporal frequency domain of incident flux. This pattern is invariant to the scene and to multi-path light transport, is highly specific to the laser, and occurs at any pixel that receives non-negligible light from it. Below we focus on the case of negligible indirect light; see supplement Section C for a more general treatment.

More specifically, suppose that laser l emits a train of pulses over the interval $[0, t_{\text{exp}}]$ with repetition frequency f_l and period $T_l = 1/f_l$ (Figure 3, first row). Since the laser is not synchronized with the camera, these pulses are emitted with an unknown clock offset $t = o_l$. The pulses then propagate through the scene at the speed of light, becoming attenuated by various factors as they travel [17]—squared-distance fall-off, surface reflection, radiant intensity of the laser source, *etc*. Consequently, the incident flux at a camera pixel $\mathbf{q}$ due to direct surface reflection will also be an attenuated and time-delayed pulse train.

Radiometrically, this train can be expressed as a time-varying function $\phi_l(\mathbf{q}, t)$ that describes instantaneous flux at time t in units of photons per second. This function depends on the temporal profile of the laser's pulse $\psi_l(t)$; the speed-of-light propagation delay $\tau_l(\mathbf{q})$ from the laser to the pixel; the pulse's attenuation $\alpha_l(\mathbf{q})$; and the laser's repetition frequency and clock offset. Equations (1)-(2) in Figure 3 provide full expressions for the incident flux and its Fourier transform, respectively; we provide further insight about them below.

The laser frequency comb. A basic property of pulse trains is that their spectrum is non-zero only at the harmonics of the repetition frequency, *i.e.*, the integer multiples of f_l (Figure 3, second row). This set of frequencies is known as a *frequency comb* [25] and has three important ramifications in our context. First, with repetition frequencies between 100 kHz and 50 MHz typical for pulsed lasers used in lidar [7, 33, 49, 84], the frequency spectrum of incident flux is extremely sparse, with gaps of potentially millions of Hz between non-zero values. Second, the flux spectrum may have many thousands of non-negligible harmonics because laser pulses must have broad spectral support to allow precise ToF measurements.[1] Third, the frequency comb depends only on the repetition frequency, not the ToF delay at a pixel. Thus, any pixel receiving light from a laser provides essentially equivalent information about the laser's presence and its repetition frequency.

Frequency-blurring impact of exposure. Cameras capture incident light over a finite duration, not indefinitely. Mathematically, restricting the incident flux to the interval $[0, t_{\text{exp}}]$ is equivalent to convolving its frequency-domain representation with an exposure-stretched sinc function. This convolution transforms the laser's frequency comb into a comb of sinc functions, effectively blurring the flux spectrum around each harmonic. For example, an exposure time of $\frac{1}{100}$ seconds will "blur" each harmonic by approximately 60 Hz.

Multiple lasers and the asynchrony advantage. When many lasers emit light in the same scene, the incident flux at a pixel will be a superposition of their individual contributions (Figure 4). Crucially, lasers whose repetition frequencies are not precise multiples of each other will have largely distinct frequency spectra, as their harmonics will rarely fall within each other's narrow sinc neighborhood (except for DC). Thus, asynchronous lasers leave a distinct "signature"

[1] For instance, a one-nanosecond Gaussian pulse—equivalent to roughly 30 cm of light travel—has significant support up to 1 GHz.

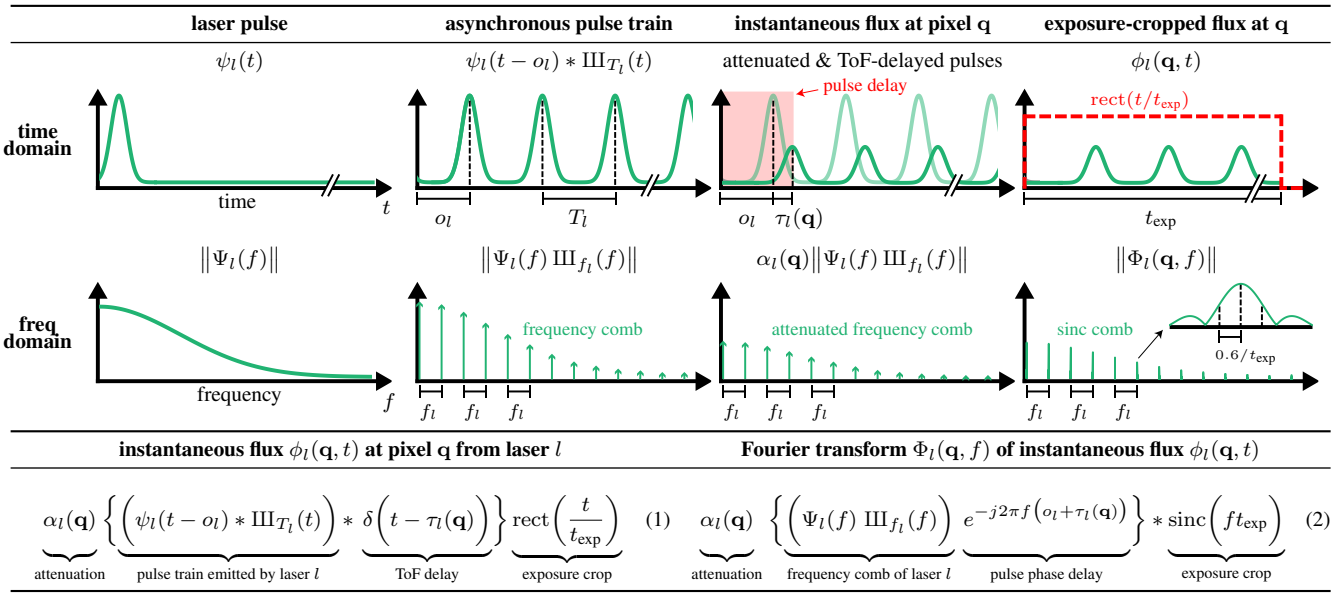

Figure 3. Incident flux model for pulsed laser light. **Top and middle row:** Our model's individual components in the temporal and frequency domain, respectively. Note that the pulse delay $o_l + \tau_l(\mathbf{q})$ in top row (column 3) is the only ToF-related quantity that a camera can measure passively with its internal clock. The full width at half max (FWHM) of the sinc in middle row (column 4) is equal to $1.2/t_{\mathrm{exp}}$. **Bottom row (Eqs. 1 and 2):** We use $\delta(t)$ to denote the impulse function, $\mathrm{III}_{T_l}(t)$ the impulse train with spacing T_l, and $*$ for convolution.

on the spectrum of incident flux that makes their flux contribution potentially separable from each others'. In practice, the degree of asynchrony needed for separability is naturally present in many lasers due to changes in temperature, power supply, and mechanical vibrations, *etc*. [73]. See supplement Section D for a limited experimental investigation of frequency variations of a picosecond laser.

Incident flux from other ambient sources. Unlike the broadband and sparse sinc combs produced by lasers, non-ToF ambient sources contribute negligibly to the incident flux at MHz to GHz frequencies. For instance, sunlight contributes only to DC at the sub-second timescales considered here and indoor lights, which flicker at frequencies ranging from Hz to tens of kHz [32, 71, 86], lack the sub-microsecond precision necessary to produce harmonics in the MHz to GHz range. Consequently, these sources introduce only random noise with near-zero amplitude in the Fourier domain of the incident flux, which we consider small enough to ignore.

4. Flux Probing Framework

Analyzing the flux received at a pixel from pulsed lasers is challenging due to its broad spectral support and weak amplitude—especially when mixed with much stronger contributions from ambient non-ToF sources. To address this problem for opportunistic ToF, we adopt and generalize the flux probing approach of Wei *et al*. [85]. The approach leverages the capabilities of emerging single-photon avalanche diode (SPAD) cameras to passively reconstruct weak flux signals spanning frequencies from Hz to tens of GHz. Here we summarize key aspects relevant to our work.

The basic premise of the approach is that incident flux at a pixel cannot be sampled directly because MHz or GHz frequencies correspond to intensity fluctuations that are orders of magnitude faster than the rate at which individual photons arrive. In this photon-limited regime, the flux $\phi(\mathbf{q},t)$ is the rate function of an inhomogeneous Poisson process that governs photon arrivals [55, 67], and must be inferred by passively detecting photons that arrive at pixel $\mathbf{q}$.

To achieve this, flux probing relies on SPAD cameras whose pixels can time-stamp individual photon detections with picosecond precision relative to an internal real-time clock. Given a frequency f and an asynchronous stream of photon timestamps at a SPAD pixel, the method formulates an unbiased estimator for the Fourier coefficient $\Phi(\mathbf{q},f)$ and a detection criterion to establish if f has significant support in the timestamp data:

Probing flux frequency f [85]: Given a stream of photon timestamps $\mathcal{T}$, the Fourier coefficient estimator

$$\widehat{\Phi}(\mathbf{q},f,\mathcal{T}) \;=\; \frac{1}{t_{\mathrm{exp}}}\sum_{t\in\mathcal{T}}\exp\left(-j2\pi f t\right) \tag{3}$$

approximately follows a complex normal distribution with mean $\Phi(\mathbf{q},f)$ and covariance matrix proportional to $|\mathcal{T}|$, the number of timestamps in $\mathcal{T}$.[2]

CFAR frequency detector: To achieve a constant probability of false alarm p, frequency f is detected if

$$\left\|\widehat{\Phi}(\mathbf{q},f,\mathcal{T})\right\|^2 \;\geq\; \mathrm{CDF}_{\chi^2}^{-1}(1-p)\frac{|\mathcal{T}|}{2t_{\mathrm{exp}}^2} \tag{4}$$

where CDF_{χ^2} is the chi-squared cumulative distribution function.

To reconstruct the time-varying incident flux at a pixel, the SPAD's entire DC-to-GHz bandwidth is densely scanned for frequencies that pass the detection criterion.[3]

[2] See [85] for derivations.

[3] Wei *et al*. [85] propose scanning the range [DC, 15 GHz] at 0.6 Hz based on their system's timestamp precision and exposure time used.

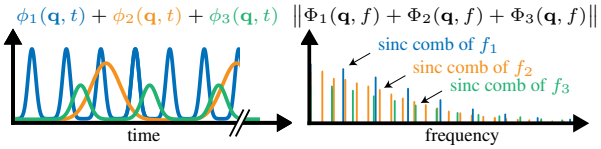

Figure 4. Flux from three asynchronous lasers shown in the temporal and frequency domains.

Function $\phi_l(\mathbf{q}, t)$ is then expressed as a Fourier series using the estimated Fourier coefficients at those frequencies.

Implications for opportunistic ToF. The CFAR detector and the dense frequency-scanning approach of [85] have four limitations in our setting:

Source-agnostic flux estimation—Opportunistic ToF requires detecting and separating the contribution of individual laser sources, not merely recovering their combined contribution to incident flux.

Photon noise—The 3D precision of opportunistic ToF implicitly depends on the GHz-scale harmonics of pulsed lasers because those harmonics allow the arrival of individual pulses to be localized more precisely in time. However, detection and reconstruction of a laser's GHz harmonics is inherently limited by photons coming from other light sources in a scene (including other lasers). Those photons raise the variance of the estimator $\widehat{\Phi}(\mathbf{q}, f, \mathcal{T})$ for every frequency f, contributing noise to the entire spectrum of estimated flux and raising the noise floor of the CFAR detector. This makes GHz harmonics less likely to be detected since they naturally have lower amplitude (Figures 1 and 6).

SPAD dead time—SPADs become inactive for a nanosecond scale interval after each photon detection, known as the dead time. While dead time does not impact the precision of individual photon timestamps, it may cause photons to be missed, particularly as the timespan between consecutive photon arrivals approaches (or falls below) the dead time. Those missed photons reduce the terms included in the sum of Eq. (3), introducing a bias in the estimator. Importantly, when some light sources are much brighter than others (room lights, a strong laser, *etc*.), photon detections from a weaker laser may get so infrequent that the Fourier coefficients of their harmonics no longer pass the CFAR detector. This exacerbates the impact of photon noise.

Computational considerations—Densely scanning the GHz band is inefficient since only a tiny fraction of scanned frequencies will ever fall near a laser's sinc comb (Figure 4).

5. Laser Discovery & Synchronization

We bypass the limitations of flux probing by exploiting the frequency-domain structure of incident flux from pulsed lasers (Section 3). Our method is based on a key observation: the laser's repetition frequency defines the full set of non-zero Fourier frequencies in the flux spectrum, via its sinc comb. Therefore, by precisely identifying the repetition frequency, we can also determine the GHz harmonics of all pixels receiving the laser's light, and vice versa.

Building on this observation, we (1) develop a method that efficiently searches the entire MHz-to-GHz band for candidate sinc combs—rather than isolated frequencies—and (2) formulate a sinc comb detection criterion to evaluate each candidate's support in the timestamp data, enabling the detection of individual lasers. See Figure 5 for an illustration and supplement Section B.1 for pseudocode.

Initial repetition frequency candidates. We identify a set $\mathcal{F}_{\text{cand}}$ of candidates by densely scanning the expected range of repetition frequencies for lasers in a scene, including in the set only frequencies that pass the CFAR detector.[4]

Frequency localization by high-res scanning. Candidate frequencies may be on the side lobe of a sinc instead of its peak. To ensure each candidate represents a local extremum in the flux spectrum, we perform a high-resolution scan in the neighborhood of each candidate and adjust its frequency accordingly, to obtain a localized set $\mathcal{F}_{\text{loc}}$.

Frequency pruning by second harmonic detection. Noise may produce a large number of candidate frequencies. These frequencies do not define a sinc comb and can be pruned by examining their second harmonics. Specifically, we first localize the second harmonics of all candidates in $\mathcal{F}_{\text{loc}}$ and then prune any candidates whose localized second harmonics fail the CFAR detector.

mHz frequency localization by harmonic hopping. Higher-order harmonics are extremely sensitive to small changes in a laser's repetition frequency. Conversely, by precisely localizing the highest detectable harmonic of a candidate frequency, we can achieve frequency localization at a much finer scale—potentially thousands of times more precise than direct localization. Starting with a candidate in $\mathcal{F}_{\text{loc}}$ we "hop" to its next higher harmonic, localize it, apply the CFAR detector, and repeat recursively until the detector fails. If the order-n harmonic is the highest detected and is localized to frequency f, we replace the candidate frequency in $\mathcal{F}_{\text{loc}}$ with f/n.

Pulse train reconstruction by harmonic probing. Each candidate frequency f still in $\mathcal{F}_{\text{loc}}$ corresponds to a potential laser signal whose harmonics extend to the GHz band, but whose individual harmonics may not be detectable. To leverage the signal those harmonics provide, we reconstruct the pulse train of a putative laser emitting pulses at frequency f by (1) estimating the Fourier coefficients of all harmonics of f up to the SPAD's band limit using Eq. (3), and (2) expressing the train $\widehat{\phi}_f(\mathbf{q}, \mathcal{T})$ as a Fourier series.

CFAR sinc comb detection and synchronization. Including all harmonics of a candidate frequency f that does not form a sinc comb will add only noise. To capture this intuition, we formulate a novel CFAR detection criterion in the time domain to test whether a frequency's reconstructed pulse train is synchronous with any photon detections.[5] Specifically, a candidate frequency f is detected if its pulse train $\widehat{\phi}_f(\mathbf{q}, \mathcal{T})$ exceeds a noise threshold for at least one timestamp in $\mathcal{T}$. The number of detected frequencies is

[4] We scan [DC, 50 MHz] with step size $0.6/t_{\text{exp}}$ in all our experiments.
[5] See supplement Section A for the exact expression and its derivation.

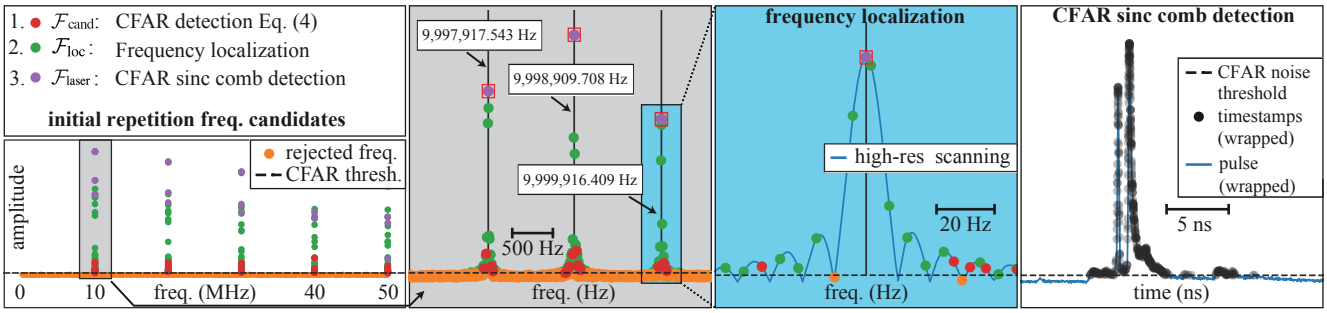

Figure 5. Overview of frequency detection. Our approach consists of four steps: (1) we apply constant false alarm rate (CFAR) detection to identify all candidate frequencies $\mathcal{F}_{\text{cand}}$ from zero to 50 MHz (red, green, purple dots in left and center left panels); (2) we obtain a refined set of frequencies $\mathcal{F}_{\text{loc}}$ (green dots) by performing a high-resolution search around each candidate frequency (center left, center right) and precisely localizing frequencies whose MHz- to GHz-rate harmonics are above the CFAR threshold; (3) we identify all dominant frequencies $\mathcal{F}_{\text{rep}}$ (purple dots) by reconstructing the time-domain flux using each frequency and its harmonics and checking if the flux has statistically significant support at the timestamp locations (right); (4) we output fundamental frequencies (red boxes) by removing all harmonics and side-lobe detections in the dominant frequencies.

the number of pulsed lasers in the scene.

Boosting SNR by local timestamp aggregation. All the above steps involve a sequence of probing operations and thus critically depend on the SNR of individual Fourier coefficients computed by Eq. (3). To boost it, all probing operations described above are applied to the *union of timestamps* detected in small N-sized neighborhoods of pixels. The intuition here is that all pixels in a neighborhood are likely to receive light from the same laser and thus they all contribute to the same Fourier coefficients with correlated phases. This effectively treats each neighborhood as a single "superpixel," and results in an SNR boost whose upper bound is $\sqrt{N}$ from Poisson noise. We find that aggregation is particularly important in low-SBR conditions, where the timestamps at individual pixels are overwhelmed by background photons and blinded by dead-time effects (Figure 6).

Pulse-delay map estimation. Upon detection of a pulsed laser with repetition frequency f, the pulse train arriving from that laser at each pixel $\mathbf{q}$ (Figure 3, third column) is reconstructed independently. Specifically, we apply the harmonic probing procedure to that pixel's timestamps without aggregation to ensure that photons with different ToFs do not affect the phase of the pulse train's Fourier coefficients. Once the train $\widehat{\phi}_f(\mathbf{q}, t, \mathcal{T})$ is reconstructed, we compute the pulse delay $o_f + \tau_f(\mathbf{q})$ by finding the global maximum of its period-wrapped counterpart, $\widehat{\phi}_f(\mathbf{q}, (t \mod 1/f), \mathcal{T})$.

6. Geometric Optimization

Pulse delays from individual lasers to specific pixels impose joint geometric constraints on the 3D positions of the lasers and the scene's depth map. This is because all light paths from lasers in the scene to a given camera pixel share a common 3D segment along the ray through that pixel (Figure 2). More specifically, any pulse delay due to a direct surface reflection from a laser to a pixel must satisfy the relation

$$c \underbrace{\left(\tau_l(\mathbf{q}) + o_l\right)}_{\text{pulse delay at pixel } \mathbf{q}} = \underbrace{\left\|\mathbf{l}_l - d(\mathbf{q})\mathbf{v}(\mathbf{q})\right\|}_{\text{laser } l \text{ to scene point along ray}} + \underbrace{d(\mathbf{q})}_{\text{depth of } \mathbf{q}} \quad (5)$$

where c is the speed of light, $\mathbf{l}_l$ is the laser's 3D position, and $\mathbf{v}(\mathbf{q})$ is the unit vector along the ray through pixel $\mathbf{q}$. It follows that Eq. (5) defines a total of LN algebraic equations for L lasers and N pixels in terms of the N unknown depths and the $4L$ laser-specific unknowns (3D positions and clock offsets). In scenes with complex geometry we expect these equations to result in an overdetermined (but non-convex) problem akin to bundle adjustment [69, 76]. To solve it, we minimize the adaptive loss function of Barron [15] on the geometric error defined by Eq. (5) using stochastic gradient descent [36]. The robust loss provides robustness to occlusions, shadows, multi-bounce specular transport and other global illumination effects. Please refer to supplement Sections B.2 and C.2 for further details.

7. Experiments

In practice, 2D single-photon cameras are an emerging technology [57, 59, 80] and we do not have one available to us; nevertheless, we create an experimental setup based on a single-pixel SPAD and one pulsed laser that accurately emulates what a 2D single-photon camera *would capture* in a scene filled with fast ambient light signals caused by asynchronous pulsed laser sources. Using this setup we demonstrate (1) simultaneous discovery and mHz-accurate

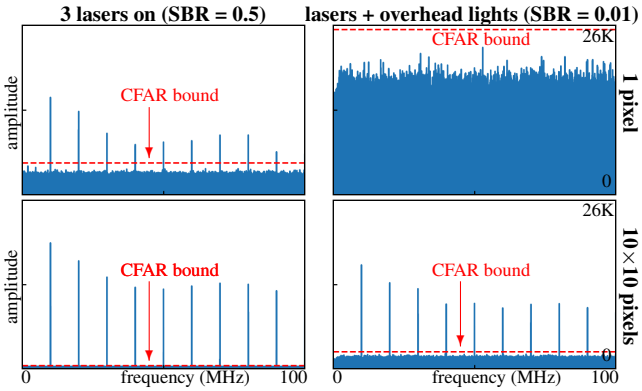

Figure 6. SNR benefit of timestamp aggregation. Estimated flux spectrum of pixel $\mathbf{q}$ in Figure 1 for the range [DC,100 MHz]. With overhead lights on and no notch filter on the SPAD to block them, no laser can be detected by probing a single pixel's timestamps.

computational synchronization to multiple asynchronous laser sources, (2) post-capture separation and visualization of their asynchronously propagating wavefronts, and (3) mm-accurate 3D imaging and laser source localization over room-scale distances, including under strong ambient light with SBR $\ll 1$. In addition to the results below, please refer to supplement Sections E–J for additional experiments and to the supplementary video for transient videos at high and low SBR, flyby visualizations of 3D reconstructions, *etc*.

Capture procedure. We emulate opportunistic single-photon time-of-flight imaging using a picosecond pulsed laser (Katana 05HP) and a single-pixel SPAD (Micro Photon Devices) that is coupled to an objective lens and a 2D scanning galvonometer (Thorlabs GVS012). A pulsed source is created by directing the collimated beam of the laser to a nearby diffuse surface, causing light from that point to illuminate the scene. We scan the scene with the SPAD and asynchronously collect photon timestamps for each pixel. Then, we illuminate a different point with the laser to create a new pulsed source and repeat this process for as many sources as needed. Last, we merge the timestamps at each pixel into a single stream, removing detections within a 231 ns dead time window matched to our SPAD. To model various SBRs, we combine the timestamp streams with ambient timestamps generated by a homogeneous Poisson process with a specific flux level. Here we show results with three lasers operating at 9.998, 9.999, and 10.000 MHz; see supplement Section G.2 for 2–8 lasers.

Opportunistic time-of-flight 3D imaging. We detect three laser sources, separate them, localize their 3D position, and recover geometry and sync offsets in a room-scale scene; this procedure results in the first visualization of light propagation from multiple separated asynchronous laser sources (Figure 1, column 3). We position a single-photon camera in the 4.53 m $\times$ 2.41 m $\times$ 2.84 m room at a position 3 m from the top wall (Figure 1, top left). The laser sources operate at 532 nm with 0.2 W average power.[6] This scene poses unique challenges: the light sources are outside the camera's field of view and have significant estimated sync offsets (6.57 m or 21.9 ns mean offset). Distance falloff, shading, and shadows cast by the divider cause the measurements for each laser source to have different SBRs (0.66, 0.83, 0.29) and photon counts (4,511, 4,367, and 2,161); shadowed regions result in outlier path length estimates.

To assess 3D reconstruction accuracy, we fit planes to the walls by segmenting them with a depth threshold; for the far and right walls, respectively, we calculate the RMSE (3.4 mm and 3.6 mm), inlier ratios (93.2% and 93.2%), and estimate the angle between the walls (90.42°) using $\sim$67000 and $\sim$23000 segmented points. Our results are accurate, recovering the $\sim$1 cm depth variations at the mannequins' nose and mm-scale details in the sleeve's folds.

Laser source discovery & frequency estimation. Figure 7 (top) displays frequency estimation error for the 10 MHz source in the scene of Figure 1; we create the error maps by applying the frequency estimation approach of Section 5 to independent patches of 10$\times$10 pixels. Under high and low SBR conditions, our estimates are typically accurate to within 1 mHz and 10 mHz, respectively (assessed using the laser sync signal). Outlier frequency errors correlate with occlusions and shadowed regions (arrows in Figure 7, top).

Comparison with UWB imaging. We compare to ultra-wideband (UWB) imaging [85] in Figure 7 (row 2). We probe frequencies up to 15 GHz, resulting in 7191 frequencies above the CFAR bound used to reconstruct flux. At high SBR (0.50) our method produces sharper pulses (110 ps FWHM vs 120 ps FWHM), with higher-intensity peaks and fewer ringing artifacts. Our approach is almost unaffected by low SBR (0.01), while UWB fails to recover pulse profiles. Notably, UWB cannot separate the contributions of each source nor computationally sync to them, as our approach does (colored transients in Figure 7, row 2).

Quantitative assessment. We assess depth error in Figure 7 (row 3) using 3D-printed staircase objects with known step sizes of 10 cm, 5 cm, 3 cm, and 1 cm or of 3 cm, 2 cm, 1 cm, 8 mm, 6 mm, 4 mm, and 2 mm, respectively. The mean SBR for the three pulsed lasers is 0.67, 0.36, 0.54, and we align the estimated 3D point clouds for each staircase to the ground-truth CAD model using the ICP algorithm [16]. Our method achieves mm-scale accuracy (2.3 mm average RMSE) and outperforms the Azure Kinect [14] (4.9 mm average RMSE). Reducing the SBR to 0.01 has little effect (2.5 mm average RMSE).

Dynamic acquisition. We consider the case where the camera and the laser sources all move independently from one acquisition to the next (Figure 7, row 4). In each acquisition, our method reliably recovers mm-accurate geometry, and successfully reconstructs concave shapes, such as a bowl (row 4, view 1). The geometry remains consistent despite significant variations in laser source positions (row 4, bottom right)—registering the point clouds sequentially with ICP yields a mean RMSE of 1.1 cm, and we estimate the angle between the walls to be 90.89° by applying plane fitting to the walls of the registered point clouds.

8. Concluding remarks

Our work represents a small, initial step toward leveraging the harmonic structure of pulsed laser signals for 3D vision. While our preliminary results are promising, pulsed lasers are used in a far broader range of settings than the one considered here, including beam-scanning lidar [37], visible-light communication [88], and many more. Developing methods that tackle this broader range of signals is challenging because the structure of those signals is far more complex than the case considered here. Additionally, 2D SPAD cameras are still in their infancy [56, 57], presenting a technological barrier to deployment of our opportunistic 3D sensing approach in the wild. Nevertheless, the specific capabilities we developed in this work—hardware-free computational synchronization, repetition-frequency estimation, low-SBR single-photon sensing, transient imaging with many asynchronous lasers, etc—are directly applicable to existing single-pixel or low-pixel-counts SPADs, as well as active sensing systems that employ pulsed lasers.

[6] This is comparable to commercially-available flash lidars. See supplement Section E.2 for a detailed comparison.

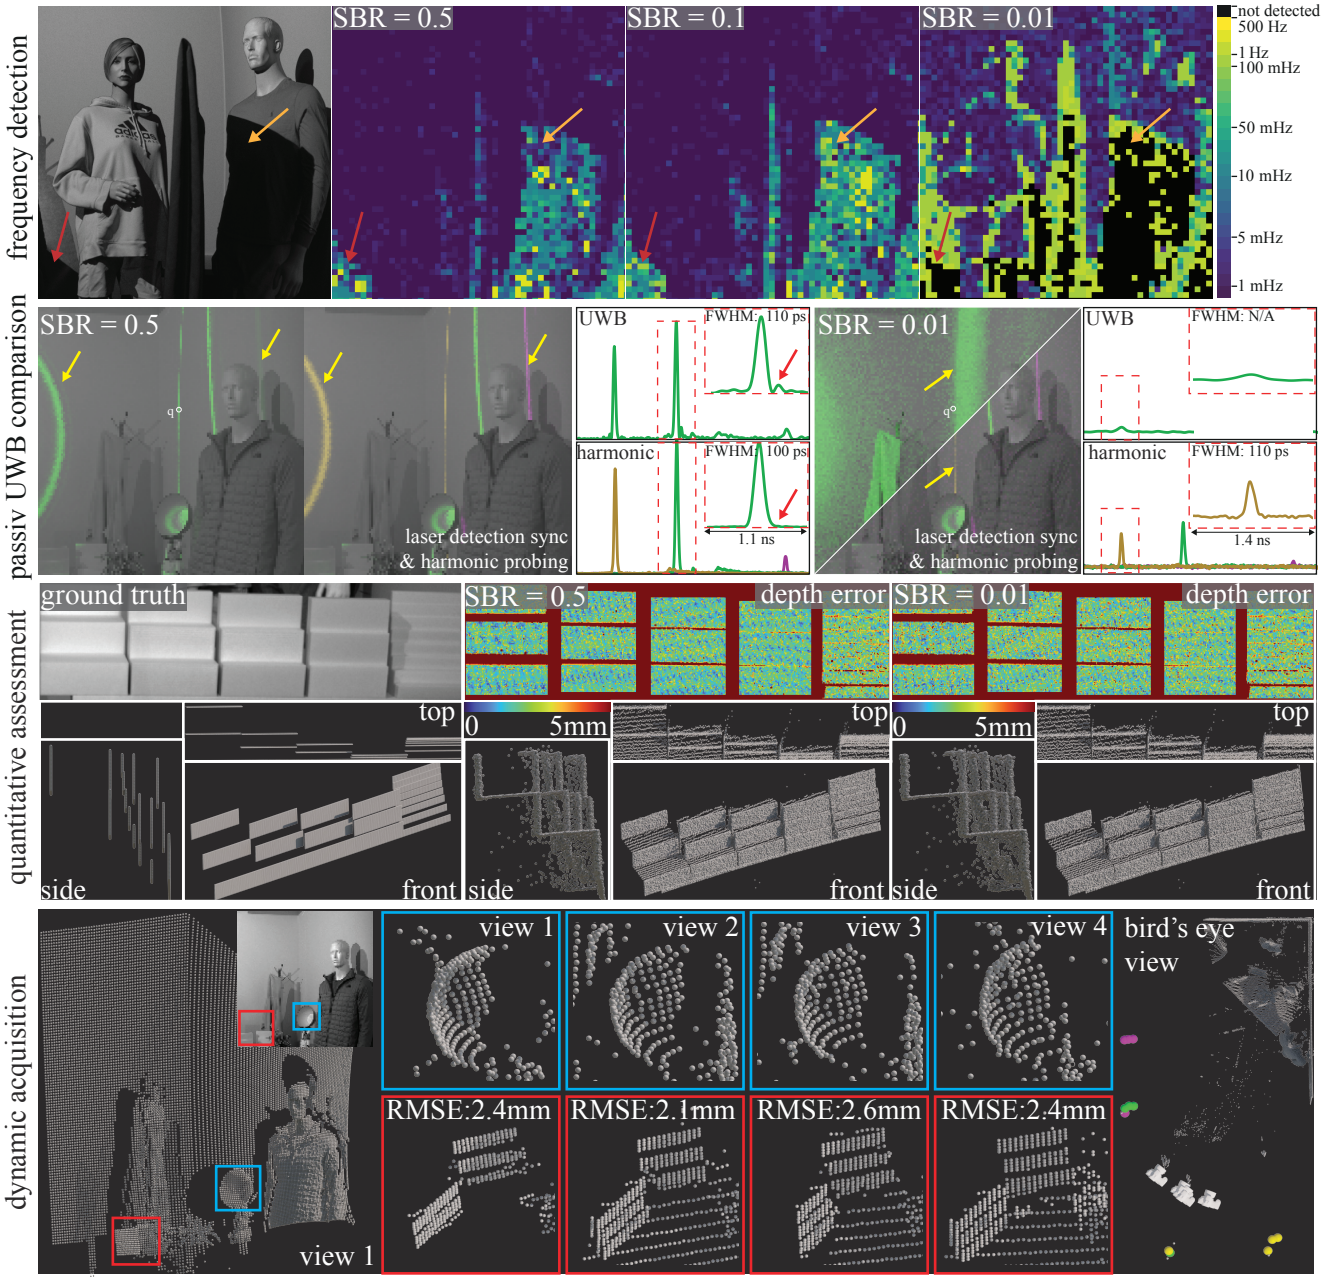

Figure 7. Opportunistic single-photon time of flight. **Row 1:** Per-pixel frequency error across different SBR levels; error in most areas is well below 100 mHz and detection fails only in shadow where the light source is occluded. **Row 2:** Our method improves over UWB [85] in both high SBR and photon-starved regimes, enabling detection and separation of asynchronous laser sources for the first time. **Row 3:** We reconstruct the 3D shape of calibration targets with known geometry; we achieve mm-scale reconstruction accuracy down to SBRs of 0.01. **Row 4:** We turn a single-photon camera into a passive 3D imaging system capable of capturing scenes where the camera and light source move at each frame. We consistently recover mm-accurate geometry despite changes in the imaging configuration.

Acknowledgements

DBL and KNK acknowledge support of NSERC under the RGPIN, RTI, and USRA programs. DBL also acknowledges support from the Canada Foundation for Innovation and the Ontario Research Fund.

References

[1] Fadel Adib and Dina Katabi. See through walls with wifi! In *Proc. ACM SIGCOMM Conf.*, pages 75–86, 2013. 1

[2] Fadel Adib, Zach Kabelac, Dina Katabi, and Robert C Miller. 3D tracking via body radio reflections. In *Proc. USENIX NSDI Symp.*, pages 317–329, 2014.

[3] Fadel Adib, Chen-Yu Hsu, Hongzi Mao, Dina Katabi, and Frédo Durand. Capturing the human figure through a wall. *ACM Trans. Graph.*, 34(6), 2015.

[4] Fadel Adib, Hongzi Mao, Zachary Kabelac, Dina Katabi, and Robert C Miller. Smart homes that monitor breathing and heart rate. In *Proc. ACM CHI Conf.*, pages 837–846, 2015. 1

[5] Faisal Ahmed, Miguel Heredia Conde, Paula López Martínez, Thomas Kerstein, and Bernd Buxbaum. Pseudo-passive time-of-flight imaging: Simultaneous illumination, communication, and 3D sensing. *IEEE Sens. J.*, 22(21): 21218–21231, 2022. 2

[6] Hüseyin Akcan, Vassil Kriakov, Hervé Brönnimann, and Alex Delis. GPS-free node localization in mobile wireless sensor networks. In *Proc. ACM MobiDE*, pages 35–42, 2006. 2

[7] ams OSRAM Group. TMF8820/21/28 Time-of-Flight Sensor Datasheet, 2024. Accessed: 2024-11-14. 3

[8] Anonymous. Sound ranging. *Nature*, 104:278–280, 1919. 1

[9] Edward V Appleton and Miles AF Barnett. On some direct evidence for downward atmospheric reflection of electric rays. *Proc. R. Soc. London, Ser. A*, 109(752):621–641, 1925. 1

[10] Eitam Arnon, Shlomo Cain, Assaf Uzan, Ran Nathan, Orr Spiegel, and Sivan Toledo. Robust time-of-arrival location estimation algorithms for wildlife tracking. *Sensors*, 23(23): 9460, 2023. 2

[11] Rashmi Bajaj, Samantha Lalinda Ranaweera, and Dharma P Agrawal. GPS: location-tracking technology. *IEEE Computer*, 35(4):92–94, 2002. 2

[12] Henrik Baktoft, Karl Ø Gjelland, Finn Økland, and Uffe H Thygesen. Positioning of aquatic animals based on time-of-arrival and random walk models using YAPS (Yet Another Positioning Solver). *Sci. Rep.*, 7(14294), 2017. 2

[13] Cyrus Bamji, John Godbaz, Minseok Oh, Swati Mehta, Andrew Payne, Sergio Ortiz, Satyadev Nagaraja, Travis Perry, and Barry Thompson. A review of indirect time-of-flight technologies. *IEEE Trans. on Electron. Devices*, 69(6): 2779–2793, 2022. 1

[14] Cyrus S Bamji, Swati Mehta, Barry Thompson, Tamer Elkhatib, Stefan Wurster, Onur Akkaya, Andrew Payne, John Godbaz, Mike Fenton, Vijay Rajasekaran, et al. IMpixel 65nm BSI 320MHz demodulated TOF image sensor with 3μm global shutter pixels and analog binning. In *Proc. IEEE ISSCC*, pages 94–96, 2018. 7

[15] Jonathan T Barron. A general and adaptive robust loss function. In *Proc. IEEE/CVF CVPR*, pages 4331–4339, 2019. 6

[16] Paul J Besl and Neil D McKay. A method for registration of 3-D shapes. *IEEE Trans. Pattern Anal. Machine Intell.*, 14 (2):239–256, 1992. 7

[17] Robert W Boyd. *Radiometry and the Detection of Optical Radiation*. Wiley, 1983. 3

[18] Clara Callenberg, Zheng Shi, Felix Heide, and Matthias B. Hullin. Low-cost SPAD sensing for non-line-of-sight tracking, material classification and depth imaging. *ACM Trans. Graph.*, 40(4), 2021. 1

[19] W Carey and N Yen. The formation of a synthetic aperture with towed hydrophones. *J. Acoust. Soc. Am.*, 75:S62–S63, 1984. 1

[20] G C Carter. Time delay estimation for passive sonar signal processing. *IEEE Trans. Acoust. Speech Signal Process.*, 29 (3):463–470, 1981. 1

[21] Wei-Yu Chen, Aswin C Sankaranarayanan, Anat Levin, and Matthew O'Toole. Coherence as texture–passive textureless 3D reconstruction by self-interference. In *Proc. IEEE/CVF CVPR*, pages 25058–25066, 2024. 1

[22] Abe Davis, Michael Rubinstein, Neal Wadhwa, Gautham Mysore, Fredo Durand, and William T. Freeman. The visual microphone: Passive recovery of sound from video. *ACM Trans. Graph.*, 33(4), 2014. 1

[23] Abe Davis, Katherine L Bouman, Justin G Chen, Michael Rubinstein, Fredo Durand, and William T Freeman. Visual vibrometry: Estimating material properties from small motion in video. In *Proc. IEEE/CVF CVPR*, pages 5335–5343, 2015. 1

[24] Zahid Farid, Rosdiadee Nordin, and Mahamod Ismail. Recent advances in wireless indoor localization techniques and system. *J. Comput. Networks Commun.*, 2013(1):185138, 2013. 2

[25] Tara Fortier and Esther Baumann. 20 years of developments in optical frequency comb technology and applications. *Commun. Phys.*, 2(1):153, 2019. 3

[26] Francesco Gramuglia, Ming-Lo Wu, Claudio Bruschini, Myung-Jae Lee, and Edoardo Charbon. A low-noise CMOS SPAD pixel with 12.1 ps SPTR and 3 ns dead time. *IEEE J. Sel. Top. Quantum Electron.*, 28(2), 2021. 2

[27] Hugh D Griffiths and Christopher J Baker. *An introduction to passive radar*. Artech House, 2022. 1

[28] Hugh D Griffiths and NRW Long. Television-based bistatic radar. In *Proc. IEE F–Commun., Radar Signal Process.*, pages 649–657, 1986.

[29] Hugh D Griffiths and Nicholas Willis. Klein heidelberg–the first modern bistatic radar system. *IEEE Trans. Aerosp. Electron. Syst.*, 46(4):1571–1588, 2010. 1

[30] Sara Grollius, Andre Buchner, Manuel Ligges, and Anton Grabmaier. Probability of unrecognized LiDAR interference for TCSPC LiDAR. *IEEE Sens. J.*, 22(13):12976–12986, 2022. 1

[31] Anant Gupta, Atul Ingle, and Mohit Gupta. Asynchronous single-photon 3d imaging. In *Proc. IEEE/CVF ICCV*, pages 7908–7917, 2019. 2

[32] Edward E Hammer. High frequency characteristics of fluorescent lamps up to 500 kHz. *J. Illum. Eng. Soc.*, 16(1): 52–61, 1987. 4

[33] Daniel Hampf, Ewan Schafer, Fabian Sproll, Toshimichi Otsubo, Paul Wagner, and Wolfgang Riede. Satellite laser ranging at 100 kHz pulse repetition rate. *CEAS Space J.*, 11(4): 363–370, 2019. 3

[34] Felix Heide, Steven Diamond, David B Lindell, and Gordon Wetzstein. Sub-picosecond photon-efficient 3D imaging using single-photon sensors. *Sci. Rep.*, 8(1):17726, 2018. 2

[35] Barmak Heshmat, Genevieve Gariepy, Jonathan Leach, Ramesh Raskar, and Daniele Faccio. SPAD cameras for biomedical imaging: Promise and problems. In *Proc. IEEE CLEO*, 2016. 2

[36] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6

[37] Ahmed Kirmani, Dheera Venkatraman, Dongeek Shin, Andrea Colaço, Franco NC Wong, Jeffrey H Shapiro, and Vivek K Goyal. First-photon imaging. *Science*, 343(6166): 58–61, 2014. 2, 7

[38] Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. Passive micron-scale time-of-flight with sunlight interferometry. In *Proc. IEEE/CVF CVPR*, pages 4139–4149, 2023. 1

[39] Jayakanth Kunhoth, AbdelGhani Karkar, Somaya Al-Maadeed, and Abdulla Al-Ali. Indoor positioning and wayfinding systems: a survey. *Hum. Cent. Comput. Inf. Sci.*, 10(18), 2020. 2

[40] Heiner Kuschel. Approaching 80 years of passive radar. In *Proc. IEEE Int. Conf. Radar*, pages 213–217, 2013. 1

[41] Robert Lange and Peter Seitz. Solid-state time-of-flight range camera. *IEEE J. Quantum Electron.*, 37(3):390–397, 2001. 1

[42] Jongho Lee, Atul Ingle, Jenu V Chacko, Kevin W Eliceiri, and Mohit Gupta. CASPI: collaborative photon processing for active single-photon imaging. *Nat. Commun.*, 14(1): 3158, 2023. 2

[43] You Li and Javier Ibanez-Guzman. Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Process Mag.*, 37(4):50–61, 2020. 1

[44] Zheng-Ping Li, Jun-Tian Ye, Xin Huang, Peng-Yu Jiang, Yuan Cao, Yu Hong, Chao Yu, Jun Zhang, Qiang Zhang, Cheng-Zhi Peng, Feihu Xu, and Jian-Wei Pan. Single-photon imaging over 200 km. *Optica*, 8:344–349, 2021. 1

[45] Yi-Chun Lin, Yi-Ting Cheng, Tian Zhou, Radhika Ravi, Seyyed M Hasheminasab, John E Flatt, Cary Troy, and Ayman Habib. Evaluation of UAV LiDAR for mapping coastal environments. *Remote Sens.*, 11(24):2893, 2019. 1

[46] David B Lindell, Matthew O'Toole, and Gordon Wetzstein. Single-photon 3D imaging with deep sensor fusion. *ACM Trans. Graph.*, 37(4), 2018. 2

[47] Fen Liu, Jing Liu, Yuqing Yin, Wenhan Wang, Donghai Hu, Pengpeng Chen, and Qiang Niu. Survey on WiFi-based indoor positioning techniques. *IET Commun.*, 14:1372–1383, 2020. 2

[48] Gregor Luetzenburg, Aart Kroon, and Anders A Bjørk. Evaluation of the Apple iPhone 12 Pro LiDAR for an application in geosciences. *Sci. Rep.*, 11, 2021. 1

[49] Lumentum. High-power Q-switched diode-pumped UV and green laser. https://www.lumentum.com/en/products/laser-solid-state-q-switched-355-532-q-series, 2024. Accessed: Nov. 10, 2024. 3

[50] Jonathan J Lynch, Harris P Moyer, James H Schaffner, Yakov Royter, Marko Sokolich, Brian Hughes, Yeong J Yoon, and Joel N Schulman. Passive millimeter-wave imaging module with preamplified zero-bias detection. *IEEE Trans. Microwave Theory Tech.*, 56(7):1592–1600, 2008. 1

[51] Robert MacCurdy, Rich Gabrielson, Eric Spaulding, Alejandro Purgue, Kathryn Cortopassi, and Kurt Fristrup. Automatic animal tracking using matched filters and time difference of arrival. *J. Commun.*, 4(7):487–495, 2009. 2

[52] Guoqiang Mao, Barış Fidan, and Brian DO Anderson. Wireless sensor network localization techniques. *Comput. Networks*, 51(10):2529–2553, 2007. 2

[53] Jess Marcum. A statistical theory of target detection by pulsed radar. *IRE Trans. Inf. Theory*, 6(2):59–267, 1960. 1

[54] Luiz EM Matheus, Alex B Vieira, Luiz FM Vieira, Marcos AM Vieira, and Omprakash Gnawali. Visible light communication: concepts, applications and challenges. *IEEE Commun. Surv. Tutor.*, 21(4):3204–3237, 2019. 1

[55] J. E. Mazo and J. Salz. On optical data communication via direct detection of light pulses. *Bell Syst. Tech. J.*, 55(3): 347–369, 1976. 4

[56] Kazuhiro Morimoto, Andrei Ardelean, Ming-Lo Wu, Arin Can Ulku, Ivan Michel Antolovic, Claudio Bruschini, and Edoardo Charbon. A megapixel time-gated SPAD image sensor for 2D and 3D imaging applications. *Optica*, 7 (4):346–354, 2020. 7

[57] Paul Mos, Scott Lindner, Chao Zhang, Michael A Wayne, Tommaso Milanese, Claudio Bruschini, and Edoardo Charbon. Piccolo gated: a CMOS 32x32 SPAD camera with all-solid-state nanosecond time gating and PCIe readout for single-photon time-domain DCS and near-infrared optical tomography. In *Proc. SPIE 12895, Quantum Sensing and Nano Electronics and Photonics XX*, number 1289507, 2024. 6, 7

[58] Nikhil Naik, Achuta Kadambi, Christoph Rhemann, Shahram Izadi, Ramesh Raskar, and Sing Bing Kang. A light transport model for mitigating multipath interference in time-of-flight sensors. In *Proc. IEEE/CVF CVPR*, pages 73–81, 2015. 1

[59] Bakhrom G Oripov, Dana S Rampini, Jason Allmaras, Matthew D Shaw, Sae Woo Nam, Boris Korzh, and Adam N McCaughan. A superconducting nanowire single-photon camera with 400,000 pixels. *Nat.*, 622(7984):730–734, 2023. 6

[60] Matthew O'Toole, Felix Heide, David B Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing transient images from single-photon sensors. In *Proc. IEEE/CVF CVPR*, pages 1539–1547, 2017. 1, 2

[61] Matthew O'Toole, David B Lindell, and Gordon Wetzstein. Confocal non-line-of-sight imaging based on the light-cone transform. *Nat.*, 555(7696):338–341, 2018. 2

[62] Parth H Pathak, Xiaotao Feng, Pengfei Hu, and Prasant Mohapatra. Visible light communication, networking, and sensing: A survey, potential and challenges. *IEEE Commun. Surv. Tutor.*, 17(4):2047–2077, 2015. 1

[63] Joshua Rapp and Vivek K Goyal. A few photons among many: Unmixing signal and noise for photon-efficient active imaging. *IEEE Trans. Comp. Imaging*, 3(3):445–459, 2017. 2

[64] Joshua Rapp, Robin MA Dawson, and Vivek K Goyal. Improving lidar depth resolution with dither. In *Proc. IEEE ICIP*, pages 1553–1557, 2018. 2

[65] Theodore S Rappaport, Jeffrey H Reed, and Brian D Woerner. Position location using wireless communications on highways of the future. *IEEE Commun. Mag.*, 34(10):33–41, 1996. 2

[66] Ricardo Roriz, Jorge Cabral, and Tiago Gomes. Automotive lidar technology: A survey. *IEEE Trans. Intell. Transp. Syst.*, 23(7):6282–6297, 2021. 1

[67] Sheldon M. Ross. *Stochastic Processes*. Wiley, 1983. 4

[68] Richard N Scarth. *Mirrors by the Sea: An Account of the Hythe Sound Mirror System Based on Contemporary Letters and Reports*. Hythe Civic Society, 1995. 1

[69] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proc. IEEE/CVF CVPR*, pages 4104–4113, 2016. 6

[70] Saad Mehmood Sheikh, Hafiz M Asif, Kaamran Raahemifar, and Fadi Al-Turjman. Time difference of arrival based indoor positioning system using visible light communication. *IEEE Access*, 9:52113–52124, 2021. 2

[71] Mark Sheinin, Yoav Y Schechner, and Kiriakos N Kutulakos. Computational imaging on the electric grid. In *Proc. IEEE/CVF CVPR*, pages 6437–6446, 2017. 4

[72] Mark Sheinin, Dorian Chan, Matthew O'Toole, and Srinivasa G Narasimhan. Dual-shutter optical vibration sensing. In *Proc. IEEE/CVF CVPR*, pages 16324–16333, 2022. 1

[73] Anthony E Siegman. *Lasers*. University Science Books, Mill Valley, Calif, 1986. 4

[74] Santosh Subedi and Jae-Young Pyun. A survey of smartphone-based indoor positioning system using RF-based wireless technologies. *Sensors*, 20(24), 2020. 2

[75] Stefano Tavani, Andrea Billi, Amerigo Corradetti, Marco Mercuri, Alessandro Bosman, Marco Cuffaro, Thomas Seers, and Eugenio Carminati. Smartphone assisted fieldwork: Towards the digital transition of geoscience fieldwork using lidar-equipped iphones. *Earth Sci. Rev.*, 227:103969, 2022. 1

[76] Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment — a modern synthesis. In *Vision Algorithms: Theory and Practice*, pages 298–372. Springer Berlin Heidelberg, 2000. 6

[77] Dobroslav Tsonev, Stefan Videv, and Harald Haas. Light fidelity (Li-Fi): towards all-optical networking. In *Proc. SPIE 9007, Broadband Access Communication Technologies VIII*, 2014. 1

[78] W S Tucker and E T Paris. A selective hot-wire microphone. *Philos. Trans. R. Soc. London, Ser. A*, 221:389–430, 1921. 1

[79] Harry Urkowitz. Energy detection of unknown deterministic signals. *Proc. IEEE*, 55(4):523–531, 1967. 1

[80] Federica Villa, Rudi Lussana, Danilo Bronzi, Simone Tisa, Alberto Tosi, Franco Zappa, Alberto Dalla Mora, Davide Contini, Daniel Durini, Sasha Weyers, and Werner Brockherde. CMOS imager with 1024 SPADs and TDCs for single-photon timing and 3-D time-of-flight. *IEEE J. Sel. Top. Quantum Electron.*, 20(6):364–373, 2014. 6

[81] Federica Villa, Fabio Severini, Francesca Madonini, and Franco Zappa. SPADs and SiPMs arrays for long-range high-speed light detection and ranging (LiDAR). *Sensors*, 21(11): 3839, 2021. 2

[82] Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Phase-based video motion processing. *ACM Trans. Graph.*, 32(4), 2013. 1

[83] Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Riesz pyramids for fast phase-based video magnification. In *Proc. IEEE ICCP*, 2014. 1

[84] Peiyuan Wang, Michael A Steindorfer, Franz Koidl, Georg Kirchner, and Erich Leitgeb. Megahertz repetition rate satellite laser ranging demonstration at Graz observatory. *Opt. Lett.*, 46(5):937–940, 2021. 3

[85] Mian Wei, Sotiris Nousias, Rahul Gulve, David B Lindell, and Kiriakos N Kutulakos. Passive ultra-wideband single-photon imaging. In *Proc. IEEE/CVF ICCV*, pages 8135–8146, 2023. 1, 2, 4, 5, 7, 8

[86] Arnold J Wilkins and C Clark. Modulation of light from fluorescent lamps. *Light. Res. Technol.*, 22(2):103–109, 1990. 4

[87] Ross E Williams. Creating an acoustic synthetic aperture in the ocean. *J. Acoust. Soc. Am.*, 60(1):60–73, 1976. 1

[88] Emma E Wollman, Jason P Allmaras, Andrew D Beyer, Boris Korzh, Marcus C Runyan, Lautaro Narváez, William H Farr, Francesco Marsili, Ryan M Briggs, Gregory J Miles, and Matthew D Shaw. An SNSPD-based detector system for NASA's Deep Space Optical Communications project. *arXiv preprint arXiv:2409.02356*, 2024. 7

[89] Hao-Yu Wu, Michael Rubinstein, Eugene Shih, John Guttag, Frédo Durand, and William Freeman. Eulerian video magnification for revealing subtle changes in the world. *ACM Trans. Graph.*, 31(4), 2012. 1

[90] Shumian Xin, Sotiris Nousias, Kiriakos N. Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. A theory of fermat paths for non-line-of-sight shape reconstruction. In *Proc. IEEE/CVF CVPR*, pages 6793–6802, 2019. 2

[91] Guochang Xu and Yan Xu. *GPS — Theory, Algorithms and Applications*. Springer, 2007. 2

[92] Larry Yujiri, Hiroshi H Agravante, Mike Biedenbender, G Samuel Dow, Martin R Flannery, Steven W Fornaca, Bruce I Hauss, Ronald L Johnson, Roger T Kuroda, Karen Jordan, Paul S Lee, Dennis Lo, Bill H Quon, Arlen W Rowe, Thomas K Samec, Merit Shoucri, Karen E Yokoyama, and John Yun. Passive millimeter-wave camera. In *Proc. SPIE Passive Millimeter-Wave Imaging Technology*, pages 15–22, 1997. 1

[93] Larry Yujiri, Merit Shoucri, and Philip Moffa. Passive millimeter wave imaging. *IEEE Trans. Microwave Theory Tech.*, 4(3):39–50, 2003. 1

[94] David Zimmerman. Tucker's acoustical mirrors: Aircraft detection before radar. *War & Society*, 15(1):73–99, 1997. 1

[95] H. Jay Zwally and Per Gloersen. Passive microwave images of the polar regions and research applications. *Polar Record*, 18(116):431–450, 1977. 1