

Passive Ultra-Wideband Single-Photon Imaging: Supplemental Document

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

Contents

A. Preliminaries	3
B. Flux Probing Theory	4
B.1 Photon counting process (proof of Equation 1)	4
B.2 Flux probing operation (proof of Proposition 1)	5
B.3 Maximum frequency (proof of Proposition 2)	6
B.4 Distribution of probing measurements (proof of Proposition 3)	7
B.5 Covariance of probing measurements	9
B.6 Distribution of Fourier probing (proof of Corollary 1)	10
B.7 Distribution of Fourier probing energy (proof of Corollary 2)	11
B.8 Unbiased estimators from probing	12
B.9 Frequency detection	14
B.10 Step size of Δf for probing	14
C. Flux Probing with Non-Negligible Dead Time	16
C.1 Generalized algorithm	16
C.2 Generalized proofs	16
C.3 Dead-time compensation step of Algorithm 1	19
D. Derivations in Special Case of Negligible Dead Time	21
D.1 Alternate proof of Proposition 3	21
D.2 Alternate derivation of the covariance of probing measurements	23
E. Experiments	25
E.1 Experimental hardware	25
E.2 Passive ultra-wideband sensing of 1D signals	25

* Joint first authors: {mianwei, sotiris}@cs.toronto.edu

E.2.1	Additional details on Figure 1 experiment	25
E.2.2	Passive ultra-wideband 1D sensing over room-size distances	26
E.2.3	Passive reconstruction of laser pulse trains	29
E.3	Passive ultra-wideband sensing of 2D signals	31
E.3.1	Fan experiment: freezing time at all timescales	31
E.3.2	Coca-Cola bottle experiment	39
E.4	Recovering passive non-line-of-sight video	41
E.5	Probing with SPAD arrays	45
F.	Supplemental Video	46
G.	Simulations	48
G.1	Simulation of laser pulse trains	48
G.1.1	Model of laser pulse	48
G.1.2	Timestamp stream generation	48
G.1.3	Estimation methods	48
G.2	Simulations with non-negligible dead time	49
G.2.1	Flux function	49
G.2.2	Timestamp stream generation	49
G.2.3	Estimation methods	49
G.2.4	Results	49
G.3	Non-line-of-sight video simulations	51
G.3.1	Model of flux function	51
G.3.2	Timestamp stream generation	51
G.3.3	Estimation methods	51
G.3.4	Reconstructions	51
G.4	Normal distribution of probing measurements	55

A. Preliminaries

We include some clarifications, definitions, and propositions from [1] that will be helpful for some of our core proofs. The equation, proposition, corollary, and figure numbering continue from where the main paper left off.

The exposure time. In the main paper we use t to denote the current instant in time and use $[0, t]$ to denote the time interval since the beginning of acquisition. As a result, t can be thought of as representing two quantities that are usually treated as distinct, *i.e.*, the current instant in time and the exposure time. This distinction is important in conventional settings (*e.g.*, exposure time for a depth measurement in active imaging or a video frame in conventional imaging) because increasing the exposure time eventually leads to temporal blur and information loss. In contrast, expanding the interval $[0, t]$ in our framework leads to a *net gain* of information. This is because flux probing does not introduce temporal blur as t increases and, in fact, frequency scanning and detection can be performed at higher resolution as the acquisition interval expands (see Section B.10). Thus, aside from computational efficiency reasons, it is advantageous to probe flux by taking the entire timestamp stream into account: as time t increases and more photons are detected, the reconstruction of the *entire* flux function from $[0, t]$ can be revised by incorporating timestamps as they arrive in order to achieve higher accuracy. This motivates using t to denote both the current time instant and the total exposure time in the main paper. For the sake of clarity, however, these quantities are treated as distinct in this supplement, with the latter denoted by t_{exp} .

Definition 1 (σ -algebra). Given a space Ω , a σ -algebra is a collection of subsets of Ω that contains the empty set \emptyset , is closed under complement and countable intersections and, by extension, countable unions.

σ -algebras provide a rigorous way to model the stream of timestamps a SPAD outputs. Specifically, a σ -algebra is the collection of all possible sets of timestamps. In other words, each element of the σ -algebra is an observation \mathcal{T} of timestamps.

Definition 2 (Filtration). A **filtration** $\mathbf{F} = (\mathbf{F}_t, t \geq 0)$ is a family of σ -algebras \mathbf{F}_t which is increasing, *i.e.*, $\forall s < t, \mathbf{F}_s \subseteq \mathbf{F}_t$.

Intuitively, a filtration \mathbf{F} models the set of all possible monotonically-increasing timestamp sequences we can observe up to any given time t .

Definition 3 (Adapted Process (Definition 3.2.14 in [1])). A stochastic process X is **F-adapted** if $\forall t \geq 0$, the random variable $X(t)$ is \mathbf{F}_t -measurable.

Intuitively, \mathbf{F}_t -measurable means that the value of $X(t)$ is determined entirely by the information available at time t , *i.e.*, the current value of $X(t)$ does not depend on the future. In our case, since the number of photon timestamps collected up to time t is known, the counting process $N(t)$ in Eq. (1) is \mathbf{F} -adapted with \mathbf{F}_t being the collection of timestamps up to time t .

Definition 4 (\mathbf{F} -Martingale (Definition 5.1.1 in [1])). An \mathbf{F} -adapted process $M = (M(t), t \geq 0)$ is an **F-martingale** if

1. the expectation of the absolute value of $M(t)$, denoted as $E[|M(t)|]$, is finite for all $t \geq 0$ and
2. the conditional expectation of $M(t)$ given the σ -algebra \mathbf{F}_s , $E[M(t)|\mathbf{F}_s]$, equals $M(s)$ almost surely for every pair (s, t) such that $s < t$.

Intuitively, a martingale can be thought of as the stochastic process equivalent of zero-mean additive noise: given the value of the martingale at time s , its expected value at any later time $t > s$, is its value at time s . Now, a key property of martingales is that their square can be expressed as the sum of another martingale $V(t)$ and a non-decreasing function that can be thought of as a stochastic generalization of the concept of variance:

Definition 5 (Quadratic Variation (Remark 11.1.2 in [1])). The **quadratic variation** of a martingale $M(t)$, denoted by $\langle M \rangle$, is the unique increasing process:

$$\langle M \rangle = M^2(t) - V(t), \tag{8}$$

where $V(t)$ is a martingale.

B. Flux Probing Theory

B.1 Photon counting process (proof of Equation 1)

We start with a rigorous re-statement of Equation 1.

Equation 1. Let $N(t)$ be the counting process of an inhomogeneous Poisson process with flux function $\varphi(t)$ and let \mathbf{F} be a filtration such that $N(t)$ is \mathbf{F} -adapted. The process

$$M(t) = N(t) - \int_0^t \varphi(u) du, \quad t \geq 0 \quad (9)$$

is an \mathbf{F} -martingale.

Proof Sketch of Equation 1.

The proof proceeds in two steps:

Step 1: We show $M(t)$ satisfies $E[|M(t)|] < \infty$ which follows from triangle inequality and each term in the triangle inequality being finite.

Step 2: We show $E[M(t) | \mathbf{F}_s] = M(s)$ which follows from the memoryless property of Poisson processes. \square

Proof of Equation 1.

Step 1: Let $t \geq 0$ and $M(t) = N(t) - \int_0^t \varphi(u) du$. Then, we have:

$$E[|M(t)|] = E\left[\left|N(t) - \int_0^t \varphi(u) du\right|\right] \quad (10a)$$

$$\leq \underbrace{E[|N(t)|]}_{|N(t)|=N(t)} + E\left[\underbrace{\left|\int_0^t \varphi(u) du\right|}_{\varphi(u) \geq 0}\right] \quad (10b)$$

$$\leq \underbrace{E[N(t)]}_{=\int_0^t \varphi(u) du} + \int_0^t \varphi(u) du \quad (10c)$$

$$= 2 \int_0^t \varphi(u) du < \infty. \quad (10d)$$

Step 2: Let $s < t$ and $M(t) = N(t) - \int_0^t \varphi(u) du$. We have to show that

$$E[M(t) | \mathbf{F}_s] = M(s). \quad (11)$$

By substituting Eq. (9) in Eq. (11), we have:

$$E[M(t) | \mathbf{F}_s] = E\left[N(t) - \int_0^t \varphi(u) du \mid \mathbf{F}_s\right] \quad (12a)$$

$$= E[N(t) | \mathbf{F}_s] - \int_0^t \varphi(u) du \quad (12b)$$

$$= E[N(t) - N(s) + N(s) | \mathbf{F}_s] - \int_0^t \varphi(u) du \quad (12c)$$

$$= E[N(t) - N(s) | \mathbf{F}_s] + E[N(s) | \mathbf{F}_s] - \int_0^t \varphi(u) du. \quad (12d)$$

Due to the independent increment property of Poisson process, the photon counts, $N(t) - N(s)$, in the interval $[s, t]$ do not depend on the arrival of photons before s , *i.e.*, its history \mathbf{F}_s . As such, $N(t) - N(s)$ is conditionally independent of \mathbf{F}_s so $E[N(t) - N(s) | \mathbf{F}_s] = E[N(t) - N(s)] = \int_s^t \varphi(u) du$. In a similar vein, the counts at s is determined entirely by its history \mathbf{F}_s

and so $E[N(s) | \mathbf{F}_s] = N(s)$. Substituting these back into Eq. (12) gives us:

$$E[M(t) | \mathbf{F}_s] = \int_s^t \varphi(u) du + E[N(s) | \mathbf{F}_s] - \int_0^t \varphi(u) du \quad (13a)$$

$$= \int_s^t \varphi(u) du + N(s) - \int_0^t \varphi(u) du \quad (13b)$$

$$= N(s) - \int_0^s \varphi(u) du \quad (13c)$$

$$= M(s). \quad (13d)$$

Therefore $M(t)$ is an \mathbf{F} -martingale. \square

Remark. While the decomposition in Eq. (9) is reminiscent of the Doob-Meyer decomposition in stochastic processes [1], the proofs below do not rely on that decomposition. Instead, we derive it directly from first principles because our proofs require an explicit construction.

B.2 Flux probing operation (proof of Proposition 1)

We restate Proposition 1 of the main paper:

Proposition 1 (Flux Probing Equation). The inner product of $p(t)$ and the unknown flux function $\varphi(t)$ over the time interval $[0, t_{\text{exp}}]$ can be expressed as

$$\langle p, \varphi \rangle = p(\mathcal{T}) + M_p(t_{\text{exp}}), \quad (14)$$

where $p(\mathcal{T})$ are ‘‘probing measurements’’ which sum the values of the probing function at the photon timestamps

$$p(\mathcal{T}) \stackrel{\text{def}}{=} \sum_{\tau \in \mathcal{T}} p(\tau), \quad (15)$$

$M_p(t)$ is a martingale, and the inner product is $\int_0^{t_{\text{exp}}} p(u)\varphi(u)du$.

Proof Sketch of Proposition 1. The proof proceeds in three steps:

Step 1: We express $\langle p, \varphi \rangle$ in terms of its Lebesgue-Stieltjes integral $\int_0^{t_{\text{exp}}} p(u)d\Phi(u)$ where $\Phi(t) = \int_0^t \varphi(u)du$. Furthermore, we substitute in the decomposition in Eq. (1) into this Lebesgue-Stieltjes integral.

Step 2: We then simplify the first term of the Lebesgue-Stieltjes integral which is the integral of $p(t)$ with respect to the counting process.

Step 3: We characterize the Lebesgue-Stieltjes integral with respect to a \mathbf{F} -martingale $M(t)$ using the following lemma:

Lemma 1 (Definition 8.3.2.1(i) in [2]). Let $M(t)$ be an \mathbf{F} -martingale. For any bounded deterministic function p , the process $M_p(t)$ given by

$$M_p(t) = - \int_0^t p(u)dM(u) \quad (16)$$

is a \mathbf{F} -martingale.

This lemma effectively says that the integral of a function p with respect to an \mathbf{F} -martingale is still an \mathbf{F} -martingale. \square

Proof of Proposition 1. Step 1: We start by expressing $\langle p, \varphi \rangle$ in terms of the Lebesgue-Stieltjes integral:

$$\langle p, \varphi \rangle = \int_0^{t_{\text{exp}}} p(u)\varphi(u)du \quad (17a)$$

$$= \int_0^{t_{\text{exp}}} p(u)d\Phi(u), \quad (17b)$$

where $\Phi(t) = \int_0^t \varphi(u)du$. Substituting Eq. (1) into the integrator gives us

$$\langle p, \varphi \rangle = \int_0^{t_{\text{exp}}} p(u) dN(u) - \int_0^{t_{\text{exp}}} p(u) dM(u). \quad (18)$$

Step 2: We evaluate the integral of $p(t)$ with respect to the counting process $N(t)$. Since $N(t)$ is a piecewise constant function, the Lebesgue-Stieltjes integral is simply the sum of $p(t)$ evaluated at the location of the jumps in $N(t)$ weighted by the size of the jump. In our case, these jumps occur exactly at our timestamps \mathcal{T} and their magnitude is 1 and therefore we have:

$$\int_0^{t_{\text{exp}}} p(u) dN(u) = \underbrace{\sum_{\tau \in \mathcal{T}} p(\tau)}_{p(\mathcal{T})}. \quad (19)$$

Step 3: Now we consider the integral of $p(t)$ with respect to the \mathbf{F} -martingale $M(t)$. Let us define $M_p(t_{\text{exp}})$ according to Eq. (16). Using Lemma 1, we know that $M_p(t_{\text{exp}})$ is a \mathbf{F} -martingale. Combining Lemma 1 and Eq. (19) gives us the flux probing equation:

$$\langle p, \varphi \rangle = \sum_{\tau \in \mathcal{T}} p(\tau) + M_p(t_{\text{exp}}). \quad (20)$$

□

B.3 Maximum frequency (proof of Proposition 2)

In this section, we prove Proposition 2 of the main paper:

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

An analogous observation can be found in [3] for irregular samples corresponding to conventional (*i.e.*, non-photon based) astronomical measurements of a variable star's flux.

Proof Sketch of Proposition 2. The main idea here is that every quantized timestamp can be written as an integer multiple of the quantization resolution $\frac{1}{Q}$. Therefore when we probe with frequencies above $\frac{1}{2Q}$, the probing function gets wrapped into a lower frequency. Specifically, we show that for any frequency $f > \frac{1}{2Q}$, there is a frequency $f_0 < f$ such that probing with f_0 is equivalent to probing with f . In other words, there is aliasing for any frequency above $\frac{1}{2Q}$.

The proof proceeds in two steps:

Step 1: We show that $\forall f > \frac{1}{Q}, f_0 = f \bmod Q$ is an alias.

Step 2: We show that $\forall f \in [\frac{1}{2Q}, \frac{1}{Q}], f_0 = \frac{1}{Q} - f$ is an alias. □

Proof of Proposition 2. Let us start by defining the quantization of timestamps with a quantization function $q(t)$ as follows:

$$q(t) = Q \left\lfloor \frac{t}{Q} + \frac{1}{2} \right\rfloor. \quad (21)$$

where Q is the SPAD's timing resolution. Then given timestamps \mathcal{T} detected during $[0, t_{\text{exp}}]$, we have the set of quantized timestamps $q(\mathcal{T}) = \{q(\tau)\}_{\tau \in \mathcal{T}}$.

For simplicity, we consider a flux function $\varphi(t) = N + A \cos(2\pi ft + \phi)$ and probing functions $p_f(t) = \frac{1}{t_{\text{exp}}} e^{-2\pi jft}$.

Step 1 ($f > \frac{1}{Q}$): Since $f > \frac{1}{Q}$, there exist $f_0 \in [0, \frac{1}{Q}]$ and $k \in \mathbb{N}$ such that $f = \frac{k}{Q} + f_0$. It is easy to see that $f_0 < f$ by construction.

We show that probing with f_0 is equivalent to probing with f :

$$p_{f_0}(q(\mathcal{T})) = \sum_{\tau \in q(\mathcal{T})} p_{f_0}(\tau) \quad (22a)$$

$$= \sum_{\tau \in \mathcal{T}} p_{f_0}(q(\tau)) \quad (22b)$$

$$= \frac{1}{t_{\text{exp}}} \sum_{\tau \in \mathcal{T}} e^{-2\pi j f_0 q(\tau)} \quad (22c)$$

$$= \frac{1}{t_{\text{exp}}} \sum_{\tau \in \mathcal{T}} e^{-2\pi j f_0 Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor} \quad (22d)$$

$$= \frac{1}{t_{\text{exp}}} \sum_{\tau \in \mathcal{T}} e^{-2\pi j f_0 Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor} \underbrace{e^{-2\pi j \frac{k}{Q} Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor}}_{=1} \quad (22e)$$

$$= \frac{1}{t_{\text{exp}}} \sum_{\tau \in \mathcal{T}} e^{-2\pi j (\frac{k}{Q} + f_0) Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor} \quad (22f)$$

$$= \frac{1}{t_{\text{exp}}} \sum_{\tau \in \mathcal{T}} e^{-2\pi j f q(\tau)} \quad (22g)$$

$$= p_f(q(\mathcal{T})). \quad (22h)$$

Step 2 ($f \in (\frac{1}{2Q}, \frac{1}{Q}]$): Consider $f_0 = \frac{1}{Q} - f$. It is easy to confirm that $f_0 < f$. We make the following observation:

$$p_{f_0}(q(\tau)) = \frac{1}{t_{\text{exp}}} e^{2\pi j f_0 q(\tau)} \quad (23a)$$

$$= \frac{1}{t_{\text{exp}}} e^{-2\pi j f_0 Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor} \quad (23b)$$

$$= \frac{1}{t_{\text{exp}}} e^{-2\pi j (\frac{1}{Q} - f) Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor} \quad (23c)$$

$$= \frac{1}{t_{\text{exp}}} \underbrace{e^{-2\pi j \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor}}_{=1} \underbrace{e^{2\pi j f Q \lfloor \frac{\tau}{Q} + \frac{1}{2} \rfloor}}_{=p_{-f}(q(\tau))} \quad (23d)$$

$$= p_{-f}(q(\tau)). \quad (23e)$$

From Eq. (23e), we have $|p_{f_0}(\mathcal{T})| = |p_f(\mathcal{T})|$ and $\angle(p_{f_0}(\mathcal{T})) = -\angle(p_f(\mathcal{T}))$. Therefore, the timestamps generated from $\varphi(t)$ are indistinguishable from those generated from $\varphi_0(t) = N + A \cos(2\pi f_0 t - \phi)$. Combining the results of Step 1 and Step 2, we have that for any frequency $f > \frac{1}{2Q}$ in our flux function φ , there exists a frequency $f_0 < f$ such that f_0 is an alias. Therefore, the maximum frequency we can detect with probing is $\frac{1}{2Q}$. \square

B.4 Distribution of probing measurements (proof of Proposition 3)

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

Proof Sketch of Proposition 3. The proof proceeds in three steps:

Step 1: We show that the mean of the probing measurements is the inner product between the probing function and the flux function. We use the following lemma and corollary here:

Lemma 2 (Proposition 8.3.2.1 in [2]). Let $N(t)$ be a counting process with flux function $\varphi(t)$, \mathbf{F} be a filtration such that $N(t)$ is \mathbf{F} -adapted. Given a deterministic function p such that $\forall t \geq 0, \int_0^t |p(u)| \varphi(u) du < \infty$. Then

$$\int_0^t p(u) \varphi(u) du = E \left[\int_0^t p(u) dN \right]. \quad (24)$$

Corollary 3. The \mathbf{F} -martingale $M_p(t) = \int_0^t p(u)dM$ has $E[M_p(t)] = 0$.

Step 2: We show that the variance of the probing measurements is the inner product between the square of the probing function and the flux function. We make use of the following lemmas and corollary:

Lemma 3 (Itô Isometry (Lemma 12.1.4 in [1])). Let $M(t)$ be a martingale and $p(t)$ be a deterministic function such that $E[\int_0^t p^2(u)d\langle M \rangle] < \infty$. Then,

$$E[(\int_0^t p(u)dM)^2] = E[\int_0^t p(u)^2 d\langle M \rangle]. \quad (25)$$

Lemma 4 (Proposition 8.3.2.1(ii) in [2]). Let $M(t)$ be an \mathbf{F} -martingale. For any bounded deterministic function p , the process $V_p(t)$ given by:

$$V_p(t) = \left(\int_0^t p(u)dM \right)^2 - \int_0^t p^2(u)\varphi(u)du, \quad (26)$$

is a \mathbf{F} -martingale.

Corollary 4. The quadratic variation of the martingale $M(t)$ defined in Eq. (9) is $\int_0^t \varphi(u)du$.

Step 3: We show that the probing measurements are approximately normally distributed. \square

Proof of Proposition 3.

Step 1: Taking the expectation of the flux probing equation (Eq. (14)) and using the facts that (a) the inner product is deterministic and (b) $E[M(t)] = 0$ (Corollary 3), it follows that $E[p(\mathcal{T})] = \langle p, \varphi \rangle$.

Step 2: We now turn our attention to the variance of the probing measurements. For brevity, we denote $M(t)$ as M . Starting from Eq. (14) we have:

$$\mathbf{var}[\sum_{\tau \in \mathcal{T}} p(\tau)] = \mathbf{var}[\underbrace{\int_0^{t_{\text{exp}}} p(u)\varphi(u)du}_{=0 \text{ (deterministic)}}] + \mathbf{var}[M_p(t_{\text{exp}})] \quad (27a)$$

$$= E[M_p(t_{\text{exp}})^2] - \underbrace{E[M_p(t_{\text{exp}})]^2}_{=0 \text{ (Cor. 3)}} \quad (27b)$$

$$= E[(\int_0^{t_{\text{exp}}} p(u)dM)^2] \quad (27c)$$

$$= E[\underbrace{\int_0^{t_{\text{exp}}} p(u)^2 d\langle M \rangle}_{\text{Itô isometry}}] \quad (27d)$$

$$= E[\int_0^{t_{\text{exp}}} p(u)^2 d \int_0^u \varphi(s)dsdu] \quad (27e)$$

$$= E[\underbrace{\int_0^{t_{\text{exp}}} p(u)^2 \varphi(u)du}_{\text{deterministic}}] \quad (27f)$$

$$= \langle p^2, \varphi \rangle. \quad (27g)$$

The second term in Eq. (27b) is equal to 0 from Corollary 3. Eq. (27d) uses Itô isometry (Lemma 3) and Eq. (27e) uses the quadratic variation $\langle M \rangle$ from Corollary 4.

Step 3: Finally, we focus on the distribution of the probing measurements. To this end, we observe that the probing measurement is the sum of independent identically-distributed random variables because each timestamp is a random variable. As the number of terms in the sum increases, the density approaches a normal density under the central limit theorem; in practice, a sum of approximately ten such random variables results in a density that closely approximates a normal distribution [4]. \square

We also verify this property empirically, via simulation, in Section G.4.

Proof of Corollary 3. We know from Eq. (1) that $M_p(t) = \int_0^t p(u)\varphi(u)du - \int_0^t p(u)dN$. Taking the expectation of both sides and using Lemma 2, we have that $E[M_p(t)] = 0$. \square

Proof of Corollary 4. By setting $p(t) = 1$ in Eq. (26) we have that $M^2(t) - \int_0^t \varphi(u)du$ is a martingale. Thus from the definition of the quadratic variation in Definition 5 we obtain that $\langle M \rangle = \int_0^t \varphi(u)du$. It is straightforward to show that $\langle M \rangle(0) = 0$. \square

B.5 Covariance of probing measurements

In this section, we prove the following proposition which we then use to prove Corollary 1 in the main paper:

Proposition 4 (Covariance of Probing Measurements). For arbitrary probing functions $p_1(t), p_2(t)$ we have that $\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du$.

Proof Sketch of Proposition 4. The proof proceeds in two steps:

Step 1: We use the definition $\mathbf{Cov}(X, Y) = E[XY] - E[X]E[Y]$ to expand $\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T}))$ and then simplify $E[p_1(\mathcal{T})p_2(\mathcal{T})]$.

Step 2: We simplify the expansion to derive $\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du$. \square

Proof of Proposition 4.

Step 1: From the definition of covariance, we have:

$$\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = E[p_1(\mathcal{T})p_2(\mathcal{T})] - E[p_1(\mathcal{T})]E[p_2(\mathcal{T})]. \quad (28)$$

First, we focus on the first term of the right hand side of Eq. (28):

$$E[p_1(\mathcal{T})p_2(\mathcal{T})] = E\left[\left(\int_0^{t_{\text{exp}}} p_1 dM + \int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du\right)\left(\int_0^{t_{\text{exp}}} p_2 dM + \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du\right)\right] \quad (29a)$$

$$= E\left[\int_0^{t_{\text{exp}}} p_1 dM \int_0^{t_{\text{exp}}} p_2 dM + \int_0^{t_{\text{exp}}} p_1 dM \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du\right] \quad (29b)$$

$$+ \int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du \int_0^{t_{\text{exp}}} p_2 dM + \int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du \quad (29c)$$

$$= \underbrace{E\left[\int_0^{t_{\text{exp}}} p_1 dM \int_0^{t_{\text{exp}}} p_2 dM\right]}_{\int_0^{t_{\text{exp}}} p_1(u)p_2(u)d\langle M \rangle \text{ (It\^o Isometry 3)}} + \underbrace{E\left[\int_0^{t_{\text{exp}}} p_1 dM\right] \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du}_{=0 \text{ (Corollary 3)}} \quad (29d)$$

$$+ \int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du \underbrace{E\left[\int_0^{t_{\text{exp}}} p_2 dM\right]}_{=0 \text{ (Corollary 3)}} + \underbrace{E\left[\int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du\right]}_{\text{deterministic}} \quad (29e)$$

$$= E\left[\int_0^{t_{\text{exp}}} p_1(u)p_2(u)d\langle M \rangle\right] + \underbrace{\int_0^{t_{\text{exp}}} p_1(u)\varphi(u)du \int_0^{t_{\text{exp}}} p_2(u)\varphi(u)du}_{=E[p_1(\mathcal{T})]E[p_2(\mathcal{T})] \text{ (Proposition 1)}} \quad (29f)$$

$$= \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du + E[p_1(\mathcal{T})]E[p_2(\mathcal{T})] \quad (29g)$$

where Eq. (29g) was obtained from Eq. (29f) by substituting the quadratic variation $\langle M \rangle$ from Corollary 4.

Step 2: Inserting Eq. (29g) in Eq. (28) we obtain:

$$\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = E[p_1(\mathcal{T})p_2(\mathcal{T})] - E[p_1(\mathcal{T})]E[p_2(\mathcal{T})] \quad (30a)$$

$$= \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du + E[p_1(\mathcal{T})]E[p_2(\mathcal{T})] - E[p_1(\mathcal{T})]E[p_2(\mathcal{T})] \quad (30b)$$

$$= \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du. \quad (30c)$$

□

B.6 Distribution of Fourier probing (proof of Corollary 1)

In this section, we prove the following corollary from the main paper:

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

$$\mu = [\langle \cos(2\pi ft), \varphi(t) \rangle \quad \langle -\sin(2\pi ft), \varphi(t) \rangle]$$

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

Proof Sketch of Corollary 1. The proof proceeds in three steps:

Step 1: We derive the mean of the real and imaginary parts of the Fourier probing measurements. This follows as a consequence of Euler's formula.

Step 2: We derive the covariance between the real and imaginary parts of the Fourier probing measurements and show that it is always 0.

Step 3: We state the conditions under which the Fourier probing measurements approximately follow a complex normal distribution. □

Proof of Corollary 1.

Step 1: Applying Euler's formula,

$$\text{Re}[p_f(\mathcal{T})] = \sum_{\tau \in \mathcal{T}} \cos(2\pi f\tau) \iff \quad (31a)$$

$$E[\text{Re}[p_f(\mathcal{T})]] = E\left[\sum_{\tau \in \mathcal{T}} \cos(2\pi f\tau)\right] \iff \quad (31b)$$

$$E[\text{Re}[p_f(\mathcal{T})]] \stackrel{\text{Proposition 3}}{=} \langle \cos(2\pi ft), \varphi(t) \rangle. \quad (31c)$$

Similarly, we have $E[\text{Im}[p_f(\mathcal{T})]] = \langle -\sin(2\pi ft), \varphi(t) \rangle$.

Step 2: Using Proposition 4 and setting $p_1(\mathcal{T}) = \text{Re}[p_f(\mathcal{T})]$ and $p_2(\mathcal{T}) = \text{Im}[p_f(\mathcal{T})]$ in Eq. (30), we have

$$\mathbf{Cov}(\text{Re}[p_f(\mathcal{T})], \text{Im}[p_f(\mathcal{T})]) = \int_0^{t_{\text{exp}}} \cos(2\pi fu)\sin(2\pi fu)\varphi(u)du \quad (32a)$$

$$= \frac{1}{2} \langle \sin(4\pi ft), \varphi(t) \rangle. \quad (32b)$$

We integrate over full periods. Letting A_{2f} be the amplitude of frequency $2f$ in φ , Eq. (32b) tells us that the covariance is

bounded by $\frac{A_{2f}}{2}$. We note the variances:

$$\langle \cos^2(2\pi ft), \varphi(t) \rangle = \frac{1}{2} \underbrace{\langle 1, \varphi(t) \rangle}_{=N(t)} + \frac{1}{2} \langle \cos(4\pi ft), \varphi(t) \rangle, \quad (33a)$$

$$\langle \sin^2(2\pi ft), \varphi(t) \rangle = \frac{1}{2} \underbrace{\langle 1, \varphi(t) \rangle}_{=N(t)} - \frac{1}{2} \langle \cos(4\pi ft), \varphi(t) \rangle. \quad (33b)$$

In general, we expect $N(t)$ to be greater than A_{2f} , so the covariance matrix Σ is approximately diagonal. In the case where $A_{2f} = 0$, this is exact.

Step 3: We know from Proposition 3 that the real and imaginary parts of the Fourier probing measurements are normally distributed. As a consequence, their linear combinations are also normally distributed and thus their joint distribution is a bivariate normal. We also know from Step 2 that their covariance is 0. Finally, we know that random variables whose joint probability density function is a normal distribution with covariance 0 are independent.

To conclude, the real and imaginary parts are independent normal distributions and by definition the complex random variable $p_f(\mathcal{T}) = \text{Re}[p_f(\mathcal{T})] + j\text{Im}[p_f(\mathcal{T})]$ is a complex normal random variable with mean

$$\mu = [\langle \cos(2\pi ft), \varphi(t) \rangle \quad \langle -\sin(2\pi ft), \varphi(t) \rangle], \quad (34)$$

and covariance matrix

$$\Sigma = \begin{bmatrix} \text{var}[\text{Re}[p_f(\mathcal{T})]] & 0 \\ 0 & \text{var}[\text{Im}[p_f(\mathcal{T})]] \end{bmatrix} \quad (35a)$$

$$\stackrel{\text{Proposition 3}}{=} \begin{bmatrix} \langle \cos^2(2\pi ft), \varphi(t) \rangle & 0 \\ 0 & \langle \sin^2(2\pi ft), \varphi(t) \rangle \end{bmatrix}. \quad (35b)$$

□

B.7 Distribution of Fourier probing energy (proof of Corollary 2)

In this section, we prove the following:

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

$$p_f^{\mathcal{E}}(\mathcal{T}) \stackrel{\text{def}}{=} \text{Re} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{1,1}}} \right]^2 + \text{Im} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{2,2}}} \right]^2 \quad (36)$$

follows a non-central χ^2 distribution with 2 degrees of freedom and non-centrality parameter:

$$\frac{\mu_1^2}{\Sigma_{1,1}} + \frac{\mu_2^2}{\Sigma_{2,2}}. \quad (37)$$

Proof Sketch of Corollary 2. The proof proceeds in three steps:

Step 1: We show that $\text{Re} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{1,1}}} \right] \sim \mathcal{N}(\mu_1, 1)$.

Step 2: We show that $\text{Im} \left[\frac{p_f(\mathcal{T})}{\sqrt{\Sigma_{2,2}}} \right] \sim \mathcal{N}(\mu_2, 1)$.

Step 3: We derive the statistics of the non-central χ^2 distribution. □

Proof of Corollary 2.

Step 1: Let X be the random variable $\text{Re}[p_f(\mathcal{T})]$. From Corollary 1, we have $X \sim \mathcal{N}(\mu_1, \Sigma_{1,1})$. Then we know that

$$\frac{1}{\sqrt{\Sigma_{1,1}}}X \sim \mathcal{N}\left(\mu_1, \frac{\Sigma_{1,1}}{\sqrt{\Sigma_{1,1}}^2}\right) = \mathcal{N}(\mu_1, 1). \quad (38)$$

Step 2: Let Y be the random variable $\text{Im}[p_f(\mathcal{T})]$. From Corollary 1, we have $Y \sim \mathcal{N}(\mu_2, \Sigma_{2,2})$. Then we know that

$$\frac{1}{\sqrt{\Sigma_{2,2}}}Y \sim \mathcal{N}\left(\mu_2, \frac{\Sigma_{2,2}}{\sqrt{\Sigma_{2,2}}^2}\right) = \mathcal{N}(\mu_2, 1). \quad (39)$$

Step 3: From Corollary 1, we know that X and Y are independent. Therefore, $X/\sqrt{\Sigma_{1,1}}$ and $Y/\sqrt{\Sigma_{2,2}}$ are also independent. We define Z as follows:

$$Z = \frac{X^2}{\Sigma_{1,1}} + \frac{Y^2}{\Sigma_{2,2}}. \quad (40)$$

Then Z is a non-central χ^2 distribution with 2 degrees of freedom and non-centrality parameter

$$E\left[\frac{X}{\sqrt{\Sigma_{1,1}}}\right]^2 + E\left[\frac{Y}{\sqrt{\Sigma_{2,2}}}\right]^2 = \frac{\mu_1^2}{\Sigma_{1,1}} + \frac{\mu_2^2}{\Sigma_{2,2}}. \quad (41)$$

□

B.8 Unbiased estimators from probing

In this section, we show that our estimators from probing are unbiased. Specifically, we show that

$$\forall \varepsilon > 0, \lim_{t_{\text{exp}} \rightarrow \infty} \Pr(|p_f(\mathcal{T}) - \langle p_f, \varphi \rangle| > \varepsilon) = 0 \quad (42)$$

and

$$\forall \varepsilon > 0, \lim_{t_{\text{exp}} \rightarrow \infty} \Pr(|p_f^2(\mathcal{T}) - \langle p_f^2, \varphi \rangle| > \varepsilon) = 0. \quad (43)$$

Therefore these estimators converge to the true mean and variance in probability. We make a mild assumption here that the flux is always finite, *i.e.*, $\|\varphi\|_\infty < \infty$.

Proof Sketch. The proof proceeds in two steps:

Step 1: We show that the estimated mean converges to the true mean in probability. We leverage Chebyshev's inequality for this.

Step 2: Similarly, we show that the estimated variance converges to the true variance in probability. □

Proof.

Step 1: Consider a set of N timestamps \mathcal{T} . To recover the Fourier coefficient of frequency f for φ in $[0, t_{\text{exp}}]$, we probe with $p_f(t) = \frac{1}{t_{\text{exp}}} e^{-2\pi jft}$. From Proposition 3, we have:

$$E[p_f(\mathcal{T})] = \langle p_f, \varphi \rangle. \quad (44)$$

Using Chebyshev's inequality we have:

$$\Pr(|p_f(\mathcal{T}) - \langle p_f, \varphi \rangle| > \varepsilon) \leq \frac{\text{var}(p_f(\mathcal{T}))}{\varepsilon^2}. \quad (45)$$

We now show that the real component of the estimated mean converges to the real component of the true mean in probability:

$$Pr(|\text{Re}[p_f(\mathcal{T})] - \text{Re}[\langle p_f, \varphi \rangle]| > \varepsilon) \leq \frac{\mathbf{var}(\text{Re}[p_f(\mathcal{T})])}{\varepsilon^2} \quad (46a)$$

$$\leq \frac{\langle \text{Re}[p_f]^2, \varphi \rangle}{\varepsilon^2} \quad (46b)$$

$$\leq \frac{1}{\varepsilon^2 t_{\text{exp}}^2} \langle \cos^2(2\pi ft), \varphi \rangle \quad (46c)$$

$$\leq \frac{1}{\varepsilon^2 t_{\text{exp}}^2} \underbrace{\|\cos^2(2\pi ft)\|_1 \|\varphi\|_\infty}_{\text{Hölder's inequality}} \quad (46d)$$

$$\leq \frac{\|\varphi\|_\infty}{\varepsilon^2 t_{\text{exp}}^2} \underbrace{4\pi f t_{\text{exp}} + \sin(4\pi f t_{\text{exp}})}_{\in \mathcal{O}(t_{\text{exp}})} \quad (46e)$$

$$\leq \frac{\|\varphi\|_\infty \mathcal{O}(t_{\text{exp}})}{\varepsilon^2 t_{\text{exp}}^2} \rightarrow 0. \quad (46f)$$

We use Hölder's inequality with $p = 1$ and $q = \infty$. Since $\|\sin^2(2\pi ft)\|_1 \in \mathcal{O}(t_{\text{exp}})$, the same inequalities in Eq. (46a) hold for the imaginary component of the estimated mean. Therefore, the means of both the real and imaginary component converge in probability to their true mean.

Step 2: Let $p_f^2(\mathcal{T})$ denote the estimated variance of $p_f(\mathcal{T})$. We show that $p_f^2(\mathcal{T})$ converges to $\langle p_f^2, \varphi \rangle$ in probability. From Proposition 3,

$$\mathbf{var}(p_f(\mathcal{T})) = \langle p_f^2, \varphi \rangle. \quad (47)$$

Using Proposition (1), we have:

$$\langle p_f^2, \varphi \rangle = p_f^2(\mathcal{T}) + M(t_{\text{exp}}). \quad (48)$$

It follows from Proposition (3) that $\mathbf{var}(M(t_{\text{exp}})) = \langle p_f^4, \varphi \rangle$.

From Eq. (48) and Chebyshev's inequality,

$$Pr(|p_f^2(\mathcal{T}) - \langle p_f^2, \varphi \rangle| > \varepsilon) = Pr(|M(t_{\text{exp}})| > \varepsilon) \quad (49)$$

$$\leq \frac{\mathbf{var}(M(t_{\text{exp}}))}{\varepsilon^2}. \quad (50)$$

To recover the Fourier coefficient for φ in $[0, t_{\text{exp}}]$, we probe with $p_f(t) = \frac{1}{t_{\text{exp}}} e^{-2\pi jft}$. We now show that the variance of the real and imaginary components of $M(t_{\text{exp}})$ converges in probability to the true variance.

We start with the real component:

$$Pr(|\text{Re}[M(t_{\text{exp}})]| > \varepsilon) \leq \frac{\mathbf{var}(\text{Re}[M(t_{\text{exp}})])}{\varepsilon^2} \quad (51a)$$

$$\leq \frac{1}{\varepsilon^2 t_{\text{exp}}^4} \langle \cos^4(2\pi ft), \varphi \rangle \quad (51b)$$

$$\leq \frac{1}{\varepsilon^2 t_{\text{exp}}^4} \underbrace{\|\cos^4(2\pi ft)\|_1 \|\varphi\|_\infty}_{\text{Hölder's inequality}} \quad (51c)$$

$$\leq \frac{\|\varphi\|_\infty}{\varepsilon^2 t_{\text{exp}}^4} \underbrace{24\pi f t_{\text{exp}} + 8\sin(4\pi f t_{\text{exp}}) + \sin(8\pi f t_{\text{exp}})}_{\in \mathcal{O}(t_{\text{exp}})} \quad (51d)$$

$$\leq \frac{\|\varphi\|_\infty \mathcal{O}(t_{\text{exp}})}{\varepsilon^2 t_{\text{exp}}^4} \rightarrow 0. \quad (51e)$$

We use Hölder's inequality with $p = 1$ and $q = \infty$.

Since $\|\sin^4(2\pi ft)\|_1 \in \mathcal{O}(t_{\text{exp}})$, the same result holds for the imaginary component. Therefore, the variance of both the real and imaginary component converges in probability to the true variance. \square

B.9 Frequency detection

In this section, we prove Equation 7 of the main paper:

$$|p_f(\mathcal{T})|^2 \geq \text{CDF}_{\chi^2}^{-1}(1 - \alpha) \frac{N(t_{\text{exp}})}{2t_{\text{exp}}^2}. \quad (52)$$

Proof Sketch. The proof proceeds in two steps:

Step 1: We show that the real and imaginary parts of the Fourier Probing Energy $p_f(\mathcal{T})$ have roughly equal variance.

Step 2: We derive the threshold for the amplitude of the Fourier probing measurements. \square

Proof.

Step 1: We probe with $p_f(t) = \frac{1}{t_{\text{exp}}} e^{-2\pi jft}$. We know from Corollary 1 that the variance of the random variables $\text{Re}[p_f(\mathcal{T})]$ and $\text{Im}[p_f(\mathcal{T})]$ is, respectively,

$$\text{var}[\text{Re}[p_f(\mathcal{T})]] = \left\langle \frac{\cos^2(2\pi ft)}{t_{\text{exp}}^2}, \boldsymbol{\varphi}(t) \right\rangle = \frac{N(t)}{2t_{\text{exp}}^2} + \frac{\overbrace{\langle \cos(4\pi ft), \boldsymbol{\varphi}(t) \rangle}^{A_{2f}}}{2t_{\text{exp}}^2}, \quad (53a)$$

$$\text{var}[\text{Im}[p_f(\mathcal{T})]] = \left\langle \frac{\sin^2(2\pi ft)}{t_{\text{exp}}^2}, \boldsymbol{\varphi}(t) \right\rangle = \frac{N(t)}{2t_{\text{exp}}^2} - \frac{\langle \cos(4\pi ft), \boldsymbol{\varphi}(t) \rangle}{2t_{\text{exp}}^2}, \quad (53b)$$

where A_{2f} is the amplitude of frequency $2f$. In general, we expect $N(t)$ to be greater than A_{2f} . Therefore,

$$\underbrace{\text{var}[\text{Re}[p_f(\mathcal{T})]]}_{\Sigma_{1,1}} = \underbrace{\text{var}[\text{Im}[p_f(\mathcal{T})]]}_{\Sigma_{2,2}} \approx \frac{N(t_{\text{exp}})}{2t_{\text{exp}}^2}. \quad (54)$$

Step 2: Combining Step 1 with the definition of the Fourier Probing Energy from Corollary 2 we obtain

$$p_f^{\mathcal{E}}(\mathcal{T}) = \frac{2t_{\text{exp}}^2}{N(t_{\text{exp}})} |p_f(\mathcal{T})|^2. \quad (55a)$$

From the CFAR detection test, we detect frequency f if

$$p_f^{\mathcal{E}}(\mathcal{T}) \geq \text{CDF}_{\chi^2}^{-1}(1 - \alpha) \quad (56a)$$

$$|p_f(\mathcal{T})|^2 \frac{2t_{\text{exp}}^2}{N(t_{\text{exp}})} \geq \text{CDF}_{\chi^2}^{-1}(1 - \alpha) \quad (56b)$$

$$|p_f(\mathcal{T})|^2 \geq \text{CDF}_{\chi^2}^{-1}(1 - \alpha) \frac{N(t_{\text{exp}})}{2t_{\text{exp}}^2}, \quad (56c)$$

which is Equation (7) in the main paper. \square

B.10 Step size of Δf for probing

We provide additional details on how to set Δf for frequency scanning in procedure FLUXREC (Figure 3 of the main paper).

We assume that the flux function has a known bandwidth $f \in [0, f_{\text{max}}]$ and was observed for an interval $[0, t_{\text{exp}}]$. Since a finite exposure time corresponds to windowing the flux function, the measured Fourier spectrum is the result of convolving $\boldsymbol{\varphi}(t)$ with a sinc function whose full width at half maximum (FWHM) is approximately $1.2/t_{\text{exp}}$ Hz [5]. We therefore scan

the desired range of frequencies with $\Delta f = 0.6/t_{\text{exp}}$ Hz in order to ensure that we “hit” the main lobe of every sinc function contributed by each frequency.

We use this scheme for the entire frequency range except near DC, where the Fourier-domain sinc lobe corresponding to the interval $[0, t_{\text{exp}}]$ can interfere with detection of low frequencies when the DC component of flux is strong. In such cases, we use $\Delta f = 1/t_{\text{exp}}$ instead. Please refer to Section E for the precise scheme followed in individual experiments. Note that our scheme is designed to balance resolution and computational efficiency, and that finer-grained frequency scanning is also possible (*e.g.*, as employed in astronomical imaging [3]).

C. Flux Probing with Non-Negligible Dead Time

For the sake of simplicity, we did not emphasize non-negligible dead time in the main paper since most high-level ideas, illustrations and algorithms can be communicated via the easier-to-grasp setting of negligible dead time and independent detections. As we show in this section, however, our stochastic calculus framework does allow handling dead time and the probabilistic dependencies it induces in the timestamp stream.

More specifically, photon detections are no longer Poisson [6] but Eqs. (1) and (2) still apply because they do not rely on the independence of photon detections. In that case, the timestamp stream is described by a more general point process that we consider in Section C.2; the random process $N(t)$ in Eq. (1) is still a counting process and the noise is still a martingale; Eq. (2) still holds but the inner product is over the known timespan that the SPAD is active, not over the full interval $[0, t_{\text{exp}}]$; and the algorithm in Figure 3 can be readily modified to convert inner products over the SPAD's active time to inner products over $[0, t_{\text{exp}}]$. This algorithm is described below.

C.1 Generalized algorithm

In the case when T_{avg} is comparable to dead time, the probing operation requires solving a complex least squares problem. Use of complex least squares is necessary because when dead time is not negligible, the union of time intervals when the SPAD is active is not equal to $[0, t_{\text{exp}}]$, and thus the Fourier basis is not orthogonal over that union. We refer to this method as *Flux probing with non-negligible dead time* (Algorithm 1). Note that this algorithm employs a frequency-scanning step size equal to $1/t_{\text{exp}}$ instead of $0.6/t_{\text{exp}}$ used in the algorithm of Figure 3. This ensures that the resulting system of linear equations in the dead-time compensation step is invertible.

The main idea of Algorithm 1 is to identify frequencies through CFAR detection, followed by the solution of a complex linear system of equations to obtain the amplitudes and phases compensated for the SPAD's active period. The formulation of this linear system is discussed in Section C.3.

Algorithm 1 Flux probing with non-negligible dead time

```

procedure FLUXRECDEADTIME( $\mathcal{T}$ ,  $t_{\text{exp}}$ ,  $f_{\text{max}}$ ,  $\alpha$ ,  $d$ )
  // Frequency scanning.
   $\Delta f = 1/t_{\text{exp}}$ 
   $\mathcal{F} =$  freqs from 0 to  $f_{\text{max}}$  with step  $\Delta f$ 
  loop  $f \in \mathcal{F}$ 
     $p_f(\mathcal{T}) = (1/t_{\text{exp}}) \sum_{\tau \in \mathcal{T}} e^{-j2\pi f \tau}$ 
  // 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
  // Dead-time compensation
  Solve Eq. (85) to obtain  $C_f$ 
   $A_f = |C_f|$ ,  $\phi_f = \angle C_f$ 
  // Flux reconstruction.
   $\hat{\phi}(t) = \sum_{f \in \mathcal{F}_{\text{used}}} A_f \cos(2\pi f t + \phi_f)$ 

```

C.2 Generalized proofs

In this section we extend our theory to the case where dead time is not negligible. To this end, we incorporate dead time d into our probing equation. Our key result is as follows:

Proposition 5 (Generalized Flux Probing Equation). The inner product of $p(t)$ and the unknown flux function $\phi(t)$ over the time interval $[0, t_{\text{exp}}]$ can be expressed as:

$$\langle p, \phi \rangle_{\text{active}} = p(\mathcal{T}) + M(t_{\text{exp}}), \quad (57)$$

where $\langle p, \phi \rangle_{\text{active}} = \int_{\Omega_{\text{active}}} p(u)\phi(u)du$ and Ω_{active} is the subset of $[0, t_{\text{exp}}]$ where the SPAD is active.

Proof Sketch of Proposition 5. The proof proceeds in three steps:

Step 1: We introduce **marked point processes** as a model for timestamp detections and introduce relevant conditional probabilities.

Step 2: We define the **flux process** $\Phi(t)$ in terms of these conditional probabilities and show that it is equivalent to $\varphi(t)$ while the SPAD is active and 0 otherwise¹. This gives us a decomposition similar to Eq. (1), but takes dead time into account.

Step 3: We derive the dead-time flux probing equation. □

Proof of Proposition 5. Note that for clarity purposes, statements of Lemmas 5-9 used in the proof below can be found after the proof of Proposition 5.

Step 1: Suppose that in addition to the set of photon arrival timestamps \mathcal{T} themselves, each timestamp τ_n is assigned a binary mark m_n , representing the two possible outcomes, *i.e.*, detected (1) and not-detected (0). The sequence (τ_n, m_n) is called a *marked point process* [7] and specifically a 2-variate point process.

We begin by making no assumption about how marks are assigned. Let $N^{(1)}(t)$ be the counting process of all detected timestamps, *i.e.*, all timestamps with $m_n = 1$. We define the history $\mathbf{F}_{\tau_{j-1}}$ to be the sequence $(\tau_n, m_n)_{n < j}$. Now consider the conditional probability of the j^{th} timestamp arriving at time t given the history $\mathbf{F}_{\tau_{j-1}}$:

$$P_j(\tau_j = t | \mathbf{F}_{\tau_{j-1}}) = \varphi(t) e^{-\int_{\tau_{j-1}}^t \varphi(u) du}. \quad (58)$$

Equation (58) follows from the Poisson arrival of τ_j [8].

Now, from Lemma 9 it follows that the conditional probability P_j of the j^{th} timestamp arriving at time t **and** being assigned mark $m_j = 1$ is

$$P_j(\tau_j = t, m_j = 1 | \mathbf{F}_{\tau_{j-1}}) = \mathbf{1}_{t > \tau_i + d} \varphi(t) e^{-\int_{\tau_{j-1}}^t \varphi(u) du}, \quad (59)$$

where τ_i is the last timestamp before τ_j for which $m_i = 1$, and $\mathbf{1}_{t > \tau_i + d}$ is the indicator function for t being outside the dead time (which occurs when $t > \tau_i + d$).

Step 2: Let us define the following stochastic process $\Phi(t)$:

$$\Phi(t) = \sum_{n \geq 0} \frac{P_{n+1}(\tau_{n+1} = t, m_{n+1} = 1 | \mathbf{F}_{\tau_n})}{1 - \int_0^{t - \tau_n} P_{n+1}(\tau_{n+1} = u + \tau_n | \mathbf{F}_{\tau_n}) du} \mathbf{1}_{\tau_n \leq t < \tau_{n+1}}. \quad (60)$$

We call $\Phi(t)$ the *flux process*. The reason for this will become evident shortly. Now consider n such that $\tau_n \leq t < \tau_{n+1}$. Then we have:

$$\Phi(t) = \frac{P_{n+1}(\tau_{n+1} = t, m_{n+1} = 1 | \mathbf{F}_{\tau_n})}{1 - \int_0^{t - \tau_n} P_{n+1}(\tau_{n+1} = u + \tau_n | \mathbf{F}_{\tau_n}) du} \quad (61a)$$

$$= \frac{\mathbf{1}_{t > \tau_i + d} \varphi(t) e^{-\int_{\tau_n}^t \varphi(u) du}}{1 - \int_0^{t - \tau_n} \varphi(u + \tau_n) e^{-\int_{\tau_n}^{u + \tau_n} \varphi(v) dv} du} \quad (61b)$$

$$= \frac{\mathbf{1}_{t > \tau_i + d} \varphi(t) e^{-\int_{\tau_n}^t \varphi(u) du}}{1 - \int_{\tau_n}^t \varphi(u) e^{-\int_{\tau_n}^u \varphi(v) dv} du}. \quad (61c)$$

Using Lemma 5 with interval $[\tau_n, \tau]$ and Eq. (61c) gives us

$$\Phi(t) = \mathbf{1}_{t > \tau_i + d} \varphi(t). \quad (62)$$

Let Ω_{active} be the set of active SPAD times in the interval $[0, t_{\text{exp}}]$. Then

$$\int_0^{t_{\text{exp}}} \Phi(u) du = \int_0^{t_{\text{exp}}} \mathbf{1}_{t > \tau_i + d} \varphi(t) du = \int_{\Omega_{\text{active}}} \varphi(t) du. \quad (63)$$

¹The approach taken here can be thought of as a generalization of [6] to the case of flux probing.

Now we use Lemma 7 to show that the stochastic process

$$M(t_{\text{exp}}) = N^{(1)}(t_{\text{exp}}) - \int_0^{t_{\text{exp}}} \Phi(u) du = N^{(1)}(t_{\text{exp}}) - \int_{\Omega_{\text{active}}} \varphi(t) du, \quad (64)$$

is an \mathbf{F} -martingale on $[0, t_{\text{exp}}]$.

Step 3: Finally, from Lemma 8,

$$E\left[\int_0^{t_{\text{exp}}} p(u) dN^{(1)}(u)\right] = E\left[\int_0^{t_{\text{exp}}} p(u) \Phi(u) du\right]. \quad (65)$$

Therefore we have

$$M(t_{\text{exp}}) = \int_0^{t_{\text{exp}}} p(u) dN^{(1)}(u) - \int_0^{t_{\text{exp}}} p(u) \Phi(u) du \quad (66a)$$

$$= \sum_{\tau \in \mathcal{T}^{(1)}} p(\tau) - \int_{\Omega_{\text{active}}} p(u) \varphi(t) du, \quad (66b)$$

where $\mathcal{T}^{(1)}$ is the set of detected timestamps. □

Lemma 5. Given an integrable function $\varphi(t)$ and an interval $[a, b]$, the following holds:

$$\frac{e^{-\int_a^b \varphi(u) du}}{1 - \int_a^b \varphi(u) e^{-\int_a^u \varphi(v) dv} du} = 1. \quad (67)$$

Proof. It suffices to show that

$$1 - \int_a^b \varphi(u) e^{-\int_a^u \varphi(v) dv} du = e^{-\int_a^b \varphi(u) du}. \quad (68)$$

Let us define the definite integral $\Phi(t) = \int_a^t \varphi(u) du$. It is easy to verify that

$$\frac{d\Phi}{dt}(t) = \varphi(t), \quad \Phi(a) = 0. \quad (69)$$

By substituting Eq. (69) into the integral of the left-hand side of Eq. (68) and using a change of variables, we have

$$- \int_a^b \varphi(u) e^{-\int_a^u \varphi(v) dv} du = \int_a^b -\frac{d\Phi}{dt}(u) e^{-\Phi(u)} du \quad (70a)$$

$$= \int_{-\Phi(a)}^{-\Phi(b)} e^u du \quad (70b)$$

$$= e^{-\Phi(b)} - e^{-\Phi(a)}. \quad (70c)$$

Substituting Eq. (70c) back into the left-hand side of Eq. (68), we obtain

$$1 - \int_a^b \varphi(u) e^{-\int_a^u \varphi(v) dv} du = 1 + e^{-\Phi(b)} - e^{-\Phi(a)} \quad (71a)$$

$$= 1 + e^{-\Phi(b)} - e^0 \quad (71b)$$

$$= 1 + e^{-\Phi(b)} - 1 \quad (71c)$$

$$= e^{-\Phi(b)} \quad (71d)$$

$$= e^{-\int_a^b \varphi(u) du}. \quad (71e)$$

□

Lemma 6 (Eq 2.22 [8]). Let \mathcal{T} be the set of timestamps from a Poisson process with flux function $\varphi(t)$. The conditional

probability density function P_n of τ_n given $\tau_{n-1}, \dots, \tau_1$ is

$$P_n(\tau_n | \tau_{n-1}, \dots, \tau_1) = \varphi(\tau_n) e^{-\int_{\tau_{n-1}}^{\tau_n} \varphi(u) du}. \quad (72)$$

Lemma 7 (III, T7 α [7]). Consider a m -variate point process $\{\tau_n, m_n\}_{n \in \mathbb{N}}$ such that $\{\tau_n\}_{n \in \mathbb{N}}$ is a point process and $m_n \in \{0, \dots, m-1\}$. For $0 \leq i < m$, the point process $N^{(i)}(t)$ is defined by

$$N^{(i)}(t) = \sum_{n \geq 1} \mathbf{1}_{\tau_n \leq t} \mathbf{1}_{m_n = i}. \quad (73)$$

If the processes $\Phi(t)$ and $A^{(i)}(t)$ are given by

$$\Phi(t) = \sum_{n \geq 0} \frac{P_{n+1}(\tau_{n+1} = t, m_{n+1} = i | \mathbf{F}_{\tau_n})}{1 - \int_0^{t - \tau_n} P_{n+1}(\tau_{n+1} = u + \tau_n | \mathbf{F}_{\tau_n}) du} \mathbf{1}_{\tau_n \leq t < \tau_{n+1}} \quad (74)$$

and

$$A^{(i)}(t) = \int_0^t \Phi(u) du, \quad (75)$$

then the process $M^{(i)}(t) = N^{(i)}(t) - A^{(i)}(t)$ is an \mathbf{F} -martingale.

Lemma 8 (III, T7 β [7]). If $M^{(i)}(t) = N^{(i)}(t) - A^{(i)}(t)$ is a \mathbf{F} -martingale, then

$$E\left[\int_0^t p(u) dN^{(1)}(u)\right] = E\left[\int_0^t p(u) \Phi(u) du\right] \quad (76)$$

for all non-negative continuous functions $p(t)$.

Corollary 5. If $M^{(i)}(t) = N^{(i)}(t) - A^{(i)}(t)$ is a \mathbf{F} -martingale, then

$$E\left[\int_0^t p(u) dN^{(1)}(u)\right] = E\left[\int_0^t p(u) \Phi(u) du\right] \quad (77)$$

for all continuous functions $p(t)$.

Proof. Let us express $p(t)$ as $p^+(t) - p^-(t)$ where $p^+(t)$ and $p^-(t)$ are non-negative continuous functions. The result follows from the linearity of expectation and Lemma 8. \square

Lemma 9. If the history $\mathbf{F}_{\tau_{j-1}}$ is the sequence $(\tau_n, m_n)_{n < j}$, the conditional probability P_j of the j^{th} timestamp arriving at time t and being assigned $m_j = 1$ is

$$P_j(\tau_j = t, m_j = 1 | \mathbf{F}_{\tau_{j-1}}) = \mathbf{1}_{t > \tau_{j-1} + d} \varphi(t) e^{-\int_{\tau_{j-1}}^t \varphi(u) du}. \quad (78)$$

Proof. From the definition of conditional probability, we have:

$$P_j(\tau_j = t, m_j = 1 | \mathbf{F}_{\tau_{j-1}}) = P_j(\tau_j = t | \mathbf{F}_{\tau_{j-1}}) P_j(m_j = 1 | \tau_j = t, \mathbf{F}_{\tau_{j-1}}) \quad (79)$$

We also have

$$P_j(\tau_j = t | \mathbf{F}_{\tau_{j-1}}) = \varphi(t) e^{-\int_{\tau_{j-1}}^t \varphi(u) du} \quad (80)$$

and

$$P_j(m_j = 1 | \tau_j = t, \mathbf{F}_{\tau_{j-1}}) = \mathbf{1}_{t > \tau_{j-1} + d}. \quad (81)$$

Equation (81) says that $m_j = 1$ with probability 1 if τ_j arrives after the dead time and $m_j = 0$ otherwise. Substituting Eqs. (80) and (81) into Eq. (79) gives us Eq. (78). \square

C.3 Dead-time compensation step of Algorithm 1

In this section we assume that the flux function is a linear combination of Fourier basis functions:

$$\varphi(t) = \sum_{f \in \mathcal{F}_{\text{used}}} A_f \cos(2\pi ft + \phi_f) = \sum_{f \in \mathcal{F}_{\text{used}} \cup -\mathcal{F}_{\text{used}}} C_f p_f(t), \quad (82)$$

where we denote the complex amplitude of the Fourier basis as C_f and use $-\mathcal{F}_{\text{used}}$ to denote the frequencies in $\mathcal{F}_{\text{used}}$ multiplied by -1 .

The presence of dead time distorts our probing measurements, *i.e.*, $\langle p, \varphi \rangle_{\Omega_{\text{active}}} \neq \langle p, \varphi \rangle$ in general. However, it is still possible to recover the amplitudes $\langle p, \varphi \rangle$. Namely, we have:

$$\langle p_l, p_k \rangle_{\Omega_{\text{active}}} = \sum_{i=1}^{N(t_{\text{exp}})} \int_{\tau_{i-1}+d}^{\tau_i} p_l(u) p_k(u) du, \quad (83)$$

where $\tau_0 = -d$. We can compute $\langle p_l, p_k \rangle_{\Omega_{\text{active}}}$ either analytically or numerically. With some abuse of notation, substituting Eq. (83) into Eq. (57) gives us:

$$p_l(\mathcal{T}) = \sum_{k=1}^K C_k \langle p_l, p_k \rangle_{\Omega_{\text{active}}} - \underbrace{\int_0^{t_{\text{exp}}} p_l(u) dM}_{=-e_l} \quad (84a)$$

$$= \begin{bmatrix} \langle p_l, p_1 \rangle_{\Omega_{\text{active}}} \\ \langle p_l, p_2 \rangle_{\Omega_{\text{active}}} \\ \vdots \\ \langle p_l, p_K \rangle_{\Omega_{\text{active}}} \end{bmatrix}^T \begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_K \end{bmatrix} + e_l, \quad (84b)$$

where K is the cardinality of $\mathcal{F}_{\text{used}} \cup -\mathcal{F}_{\text{used}}$. Combining the K equations associated with each probing function p_l yields a linear system of equations for recovering the complex Fourier amplitudes C_k :

$$\begin{bmatrix} p_1(\mathcal{T}) \\ p_2(\mathcal{T}) \\ \vdots \\ p_K(\mathcal{T}) \end{bmatrix} = \begin{bmatrix} \langle p_1, p_1 \rangle_{\Omega_{\text{active}}} & \langle p_1, p_2 \rangle_{\Omega_{\text{active}}} & \cdots & \langle p_1, p_K \rangle_{\Omega_{\text{active}}} \\ \langle p_2, p_1 \rangle_{\Omega_{\text{active}}} & \langle p_2, p_2 \rangle_{\Omega_{\text{active}}} & \cdots & \langle p_2, p_K \rangle_{\Omega_{\text{active}}} \\ \vdots & \vdots & \ddots & \vdots \\ \langle p_K, p_1 \rangle_{\Omega_{\text{active}}} & \langle p_K, p_2 \rangle_{\Omega_{\text{active}}} & \cdots & \langle p_K, p_K \rangle_{\Omega_{\text{active}}} \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \\ \vdots \\ C_K \end{bmatrix} + \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_K \end{bmatrix}. \quad (85)$$

Therefore we can solve for these amplitudes by solving the above least squares problem in the complex domain.

D. Derivations in Special Case of Negligible Dead Time

In this section, we provide an alternative perspective on the flux probing equation when dead time is short enough relative to the photon inter-arrivals that it can be ignored. Under such conditions, photon detections can be modelled as an inhomogeneous Poisson process. Leveraging the independence of photon detections in such settings, the probing measurement can also be viewed as a Monte Carlo estimator. The following proofs are all **courtesy of anonymous Reviewer 1**.

D.1 Alternate proof of Proposition 3

We show that treating probing as Monte Carlo estimation provides an alternate proof for the mean and variance of the probing measurement $p(\mathcal{T})$. We restate Proposition 3 (without the normality claims) here:

Proposition 3 (Mean and Variance of Probing Measurements). The probing measurements $p(\mathcal{T})$ have mean $\langle p, \varphi \rangle$ and variance $\langle p^2, \varphi \rangle$.

Proof Sketch of Proposition 3. The proof proceeds in two steps:

Step 1: We invoke the law of total expectation,

$$E[X] = E[E[X|Y]], \quad (86)$$

where X and Y are random variables, to show that the mean is $\langle p, \varphi \rangle$.

Step 2: We invoke the law of total variance:

$$\mathbf{var}[X] = \mathbf{var}[E[X|Y]] + E[\mathbf{var}[X|Y]], \quad (87)$$

where X and Y are random variables, to show that the variance is $\langle p^2, \varphi \rangle$.

□

Proof of Proposition 3.

Step 1: When dead time is negligible, the sequence of timestamps \mathcal{T} in the interval $[0, t_{\text{exp}}]$ follows an inhomogeneous Poisson process with flux equal to $\varphi(t)$. The probability of an event occurring at time $t \in [0, t_{\text{exp}}]$ is equal to $\varphi(t)/E[N]$ [8] where $E[N]$ is the expected number of photon detections in the interval $[0, t_{\text{exp}}]$.

Applying Eq. (86),

$$E[p(\mathcal{T})] = E[E[p(\mathcal{T})|N]], \quad (88)$$

where N is the number of photon detections in the interval $[0, t_{\text{exp}}]$. Since the timestamps in \mathcal{T} are independent of each other, for the inner expectation we have:

$$E[p(\mathcal{T})|N] = E\left[\sum_{\tau \in \mathcal{T}} p(\tau)|N\right] = \sum_{i=1}^N E[p(\tau_i)] = N \frac{\langle p, \varphi \rangle}{E[N]}. \quad (89)$$

Therefore, the total expectation becomes:

$$E[p(\mathcal{T})] = E\left[N \frac{\langle p, \varphi \rangle}{E[N]}\right] = E[N] \frac{\langle p, \varphi \rangle}{E[N]} = \langle p, \varphi \rangle. \quad (90)$$

Step 2: Similarly, applying Eq. (87), we obtain

$$\mathbf{var}[p(\mathcal{T})] = \mathbf{var}[E[p(\mathcal{T})|N]] + E[\mathbf{var}[p(\mathcal{T})|N]]. \quad (91)$$

The first term of the RHS of Eq. (91) can be simplified using Eq. (89) and the fact that $\mathbf{var}[N] = E[N]$:

$$\mathbf{var}[E[p(\mathcal{T})|N]] = \mathbf{var}\left[N \frac{\langle p, \varphi \rangle}{E[N]}\right] = \mathbf{var}[N] \frac{\langle p, \varphi \rangle^2}{E[N]^2} = \frac{\langle p, \varphi \rangle^2}{E[N]}. \quad (92)$$

For the second term of the RHS of Eq. (91), we have

$$\mathbf{var}[p(\mathcal{T})|N] = E[p^2(\mathcal{T})|N] - E[p(\mathcal{T})|N]^2 \quad (93a)$$

$$= E[p^2(\mathcal{T})|N] - N^2 \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2}. \quad (93b)$$

Expanding the first term of Eq. (93b),

$$E[p^2(\mathcal{T})|N] = E \left[\sum_{i=1}^N p(\tau_i) \cdot \sum_{i=1}^N p(\tau_i) \right] \quad (94a)$$

$$= E \left[\sum_{i=1}^N p^2(\tau_i) + 2 \sum_{i=1}^N \sum_{j=1}^{i-1} p(\tau_i) p(\tau_j) \right] \quad (94b)$$

$$= E \left[\sum_{i=1}^N p^2(\tau_i) \right] + 2E \left[\sum_{i=1}^N \sum_{j=1}^{i-1} p(\tau_i) p(\tau_j) \right] \quad (94c)$$

$$= \sum_{i=1}^N E[p^2(\tau_i)] + 2 \sum_{i=1}^N \sum_{j=1}^{i-1} E[p(\tau_i) p(\tau_j)] \quad (94d)$$

$$= N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} + 2 \sum_{i=1}^N \sum_{j=1}^{i-1} E[p(\tau_i)] \cdot E[p(\tau_j)] \quad (94e)$$

$$= N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} + 2 \sum_{i=1}^N \sum_{j=1}^{i-1} \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2} \quad (94f)$$

$$= N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} + \frac{2(N-1)N}{2} \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2} \quad (94g)$$

$$= N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} + (N-1)N \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2}. \quad (94h)$$

Combining Eqs. (94h) and (93b),

$$\mathbf{var}[p(\mathcal{T})|N] = N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} + (N-1)N \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2} - N^2 \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2} \quad (95a)$$

$$= N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} - N \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2}. \quad (95b)$$

The expectation of $\mathbf{var}[p(\mathcal{T})|N]$ is

$$E[\mathbf{var}[p(\mathcal{T})|N]] = E \left[N \frac{\langle p^2, \boldsymbol{\varphi} \rangle}{E[N]} - N \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]^2} \right] \quad (96a)$$

$$= \langle p^2, \boldsymbol{\varphi} \rangle - \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]}. \quad (96b)$$

Finally, by substituting Eqs. (96b) and (92) into Eq. (91), we obtain

$$\mathbf{var}[p(\mathcal{T})] = \mathbf{var}[E[p(\mathcal{T})|N]] + E[\mathbf{var}[p(\mathcal{T})|N]] \quad (97a)$$

$$= \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]} + \langle p^2, \boldsymbol{\varphi} \rangle - \frac{\langle p, \boldsymbol{\varphi} \rangle^2}{E[N]} \quad (97b)$$

$$= \langle p^2, \boldsymbol{\varphi} \rangle. \quad (97c)$$

□

D.2 Alternate derivation of the covariance of probing measurements

We show that the Monte Carlo interpretation of flux probing also provides an alternate proof of Proposition 4 for the covariance, which we restate here:

Proposition 4 (Covariance of Probing Measurements). For arbitrary probing functions $p_1(t)$, $p_2(t)$ we have that $\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\varphi(u)du$.

Proof Sketch of Proposition 4. Let $h(t) = p_1(t)p_2(t)$. We invoke the law of total covariance

$$\mathbf{Cov}(X, Y) = \mathbf{Cov}(E[X|Z], E[Y|Z]) + E[\mathbf{Cov}(X, Y|Z)], \quad (98)$$

and conditional covariance

$$\mathbf{Cov}(X, Y|Z) = E[XY|Z] - E[X|Z]E[Y|Z], \quad (99)$$

where X , Y , and Z are random variables, to prove the proposition. \square

Proof of Proposition 4. Invoking Eq. (98), we have:

$$\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = \mathbf{Cov}(E[p_1(\mathcal{T})|N], E[p_2(\mathcal{T})|N]) + E[\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N)]. \quad (100)$$

The first term of the RHS of Eq. (100) can be simplified using Eq. (89):

$$\mathbf{Cov}(E[p_1(\mathcal{T})|N], E[p_2(\mathcal{T})|N]) = \mathbf{Cov}\left(N \frac{\langle p_1, \varphi \rangle}{E[N]}, N \frac{\langle p_2, \varphi \rangle}{E[N]}\right) \quad (101a)$$

$$= \frac{\langle p_1, \varphi \rangle \langle p_2, \varphi \rangle}{E[N]^2} \mathbf{Cov}(N, N) \quad (101b)$$

$$= \frac{\langle p_1, \varphi \rangle \langle p_2, \varphi \rangle}{E[N]^2} \underbrace{\mathbf{var}[N]}_{= E[N]} \quad (101c)$$

$$= \frac{\langle p_1, \varphi \rangle \langle p_2, \varphi \rangle}{E[N]}. \quad (101d)$$

Equation (99) allows us to expand $\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N)$:

$$\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N) = E[p_1(\mathcal{T})p_2(\mathcal{T})|N] - E[p_1(\mathcal{T})|N]E[p_2(\mathcal{T})|N]. \quad (102)$$

We now proceed by simplifying the first term of the RHS of Eq. (102):

$$E[p_1(\mathcal{T})p_2(\mathcal{T})|N] = E\left[\sum_{i=1}^N p_1(\tau_i) \cdot \sum_{i=1}^N p_2(\tau_i)\right] \quad (103a)$$

$$= E\left[\sum_{i=1}^N p_1(\tau_i)p_2(\tau_i) + 2\sum_{i=1}^N \sum_{j=1}^{i-1} p_1(\tau_i)p_2(\tau_j)\right] \quad (103b)$$

$$= E\left[\sum_{i=1}^N h(\tau_i)\right] + 2E\left[\sum_{i=1}^N \sum_{j=1}^{i-1} p_1(\tau_i)p_2(\tau_j)\right] \quad (103c)$$

$$= \sum_{i=1}^N E[h(\tau_i)] + 2\sum_{i=1}^N \sum_{j=1}^{i-1} E[p_1(\tau_i)p_2(\tau_j)] \quad (103d)$$

$$= N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} + 2\sum_{i=1}^N \sum_{j=1}^{i-1} E[p_1(\tau_i)] \cdot E[p_2(\tau_j)] \quad (103e)$$

$$= N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} + 2\sum_{i=1}^N \sum_{j=1}^{i-1} \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2} \quad (103f)$$

$$= N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} + \frac{2(N-1)N}{2} \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2} \quad (103g)$$

$$= N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} + (N-1)N \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2}, \quad (103h)$$

where $h(t) = p_1(t)p_2(t)$.

From Eq. (89), it follows that the second term of Eq. (102) is equal to

$$E[p_1(\mathcal{T})|N]E[p_2(\mathcal{T})|N] = N^2 \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2}. \quad (104)$$

We can now express $E[\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N)]$ as

$$E[\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N)] = E\left[N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} + (N-1)N\frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2} - N^2\frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2}\right] \quad (105a)$$

$$= E\left[N\frac{\langle h, \boldsymbol{\varphi} \rangle}{E[N]} - N\frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]^2}\right] \quad (105b)$$

$$= \langle h, \boldsymbol{\varphi} \rangle - \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]}. \quad (105c)$$

Combining Eqs. (100), (101) and (105c), we finally obtain

$$\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})) = \mathbf{Cov}(E[p_1(\mathcal{T})|N], E[p_2(\mathcal{T})|N]) + E[\mathbf{Cov}(p_1(\mathcal{T}), p_2(\mathcal{T})|N)] \quad (106a)$$

$$= \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]} + \langle h, \boldsymbol{\varphi} \rangle - \frac{\langle p_1, \boldsymbol{\varphi} \rangle \langle p_2, \boldsymbol{\varphi} \rangle}{E[N]} \quad (106b)$$

$$= \langle h, \boldsymbol{\varphi} \rangle \quad (106c)$$

$$= \int_0^{t_{\text{exp}}} h(u)\boldsymbol{\varphi}(u)du \quad (106d)$$

$$= \int_0^{t_{\text{exp}}} p_1(u)p_2(u)\boldsymbol{\varphi}(u)du. \quad (106e)$$

□

E. Experiments

We present additional information about the experiments in the paper and include one additional experiment in Section E.2.3 to evaluate our method’s ability to passively reconstruct ultrafast laser pulse trains.

E.1 Experimental hardware

The following sensors and light sources are used in our experiments:

- *Free-running SPAD*: We use a single-pixel PDM Series Fast-Gated SPAD from Micro Photon Devices, operated in asynchronous mode. The SPAD has a $50 \times 50 \mu\text{m}$ active area, a timing jitter of 68 ps, a 231 ns dead time and a quantum efficiency of 18% at 660 nm.
- *Time-to-digital converter (TDC)*: The SPAD’s output is converted into a stream of timestamps by a PicoQuant PicoHarp300 system. Except where noted, the PicoHarp is operated in asynchronous mode and is not connected to any laser sources. The PicoHarp’s timing resolution is $Q = 4$ ps and its overall timing jitter is 16 ps.
- *Galvo mirrors*: For experiments where 2D imaging was required, we used a pair of galvo mirrors from Thorlabs (Model GVS012) to acquire measurements over a 2D field of view with our single-pixel SPAD.
- *Laser scanning video projector*: We used an AnyBeam MEMS raster scanning laser projector (model HD301M1-H2-01) for some of our imaging experiments. Its nominal frame rate is 60 Hz, its resolution is 1280×720 and its nominal brightness is 30 ANSI lumens. Video was supplied via an HDMI port.
- *Miscellaneous bulbs*: We used three lamps for additional illumination: (1) a Kasa Smart Wi-Fi LED Multi-Color Light Bulb (800 lumens), (2) an off-the-shelf compact fluorescent lamp, and (3) our laboratory’s ceiling lights, consisting of twelve T8 replacement LED lamps (SYLVANIA OCTRON, FO28/835XP/XL/SS/ECO3, 2600 lumens, driven by SYLVANIA QUICKTRONIC high-efficiency ballasts).
- *Spinning fan*: We used a handheld, battery-operated spinning fan operating at a frequency of approximately 54 Hz.
- *Color filters*: For acquisition of color ultra-wideband video we used red, green and blue transparent color correction lighting gel filters from SAKOLLA.
- *Avalanche Photodiode*: Thorlabs Model DET10A2.
- *Collimated Laser Diode*: Thorlabs Model CPS532-C2.
- *Ultrafast low-power pulsed lasers*: We used two Alphasalas picosecond lasers for most of our passive ultra-wideband imaging experiments, both having a wavelength of approximately 665 nm (models PLD670-100 and PLD670-50). The lasers have a user-selectable power setting, a user-selectable pulse repetition rate between 1 MHz to 100 MHz, and a nominal pulse FWHM of 83 ps for the power setting we employed.
- *Ultrafast high-power pulsed laser*: We also used a class-4 picosecond laser with a wavelength of 532 nm (NKT Photonics Katana 05HP). The repetition rate was 10 MHz and its FWHM was 58 ps under the experimental conditions we employed, resulting in a flux function with frequency support up to 17.24 GHz.
- *Experimentally-verified top imaging speed*: The upper bound is limited by the speed of our fastest light source—*i.e.*, the NKT Photonics Katana 05HP—not our method or our SPAD.

E.2 Passive ultra-wideband sensing of 1D signals

E.2.1 Additional details on Figure 1 experiment

We provide additional details about the experiment shown in Figure 1.

Scene: The scene is shown in Figure 5 (row 2, second image from right). The scene illumination consisted of the laser projector, our smart light bulb and our two low-powered picosecond laser sources. At any given instant, all light sources behaved as diffuse point emitters either by design (*i.e.*, smart bulb) or because of a diffuser placed in front of them. The SPAD was focused on a single point on a planar target as shown in the photo.

Illumination: We used the projector to play a five-second MGM video clip at a nominal frame rate of 60 frames per second. A Thorlabs ground-glass diffuser was placed in front of the laser projector so that light is diffused from each projector pixel as would happen with light bouncing off of a line-of-sight (LOS) wall. A similar diffuser was placed in front of the pulsed laser sources, which were operating at 3 MHz and 40 MHz, respectively. The smart bulb was experimentally determined to flicker at approximately 900 Hz. The smart bulb was located two meters above the optical table shown in the photo.

Acquisition: We acquired a single timestamp stream over a one-second acquisition interval. A total of 770k photons were detected. To make flux reconstruction even more challenging, we applied Bernoulli thinning with $p = 0.1$ to the captured timestamps, resulting in a single stream of 77k timestamps in total. This stream was used for probing and for the results shown in Figure 1.

Frequency scanning and detection: For frequencies between 0 Hz and 3 MHz we probed with a step of 1 Hz, taking into account the 77k timestamps spanning the full one-second acquisition interval. For frequencies between 3 MHz and 10 GHz we probed with a step of 6 Hz using only the approximately 7.7k timestamps acquired in the first 100 ms of the acquisition. This was done for computational efficiency.² For each probed range, we set the probability α of false alarm to $\frac{1}{\#freqs}$, where $\#freqs$ is the number of probed frequencies. This ensures that only one false alarm is returned by the CFAR detector on average.

Flux function visualizations: To plot the reconstructed flux function at the four timescales shown in Figure 1, we sampled it over four different temporal windows—1 s, 5.5 ms, 1 μ s, and 2.7 ns. For temporal windows less than one second, we centered them at $t = 500.983$ ms. This time instant was chosen to highlight the asynchronous relation of the two laser pulses. The flux function was evaluated at 10k regular samples over each temporal window. These samples were then low-pass filtered with frequency cutoffs at 5 kHz, 100 kHz, 5 GHz, and 1.8 THz, respectively, to avoid aliasing in each plot.

Figure 6 shows an enlarged and more detailed view of the flux spectrum shown in Figure 1, where the detected frequencies and the CFAR thresholds employed are shown, for both the 0-to-3 MHz and the 3 MHz-to-10 GHz bands.

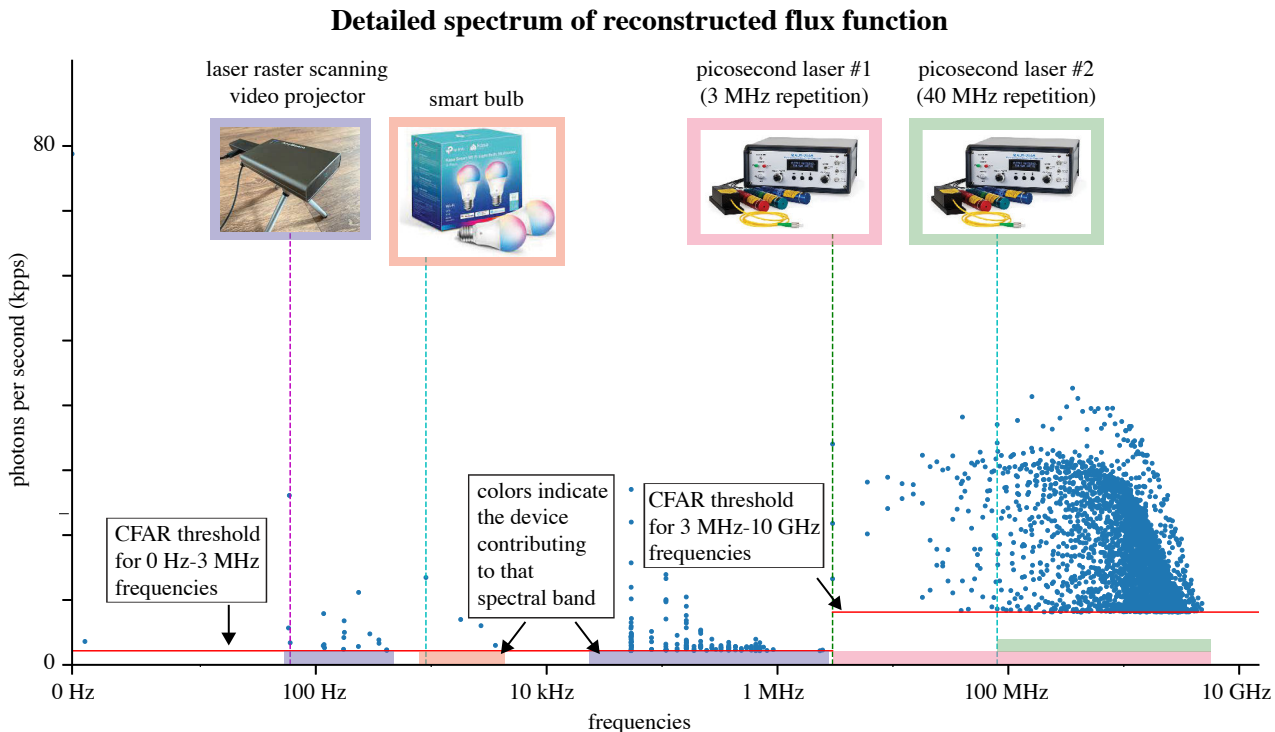


Figure 6: Detailed spectrum of the reconstructed flux function for experiment in Figure 1.

E.2.2 Passive ultra-wideband 1D sensing over room-size distances

In this experiment we demonstrate passive reconstruction of multiple concurrent high-frequency flux variations over room-sized distances.

²Probing the 3 MHz-to-10 GHz bandwidth with 6 Hz steps took 100 minutes on a 383 GB Dual 14-Core CPU running at 2.2GHz and a 24 GB Titan RTX GPU. Note that the computational cost increases quadratically with exposure time because it increases the number of timestamps processed and also reduces the step size employed in frequency scanning.

Scene: Our SPAD was focused at a spot on one of our laboratory’s diffusely-painted walls, approximately one meter away from the SPAD. The spot, as well as the entire laboratory, was illuminated simultaneously by two emitters:

- a very bright, diffusely-reflecting laser spot approximately four meters away from where the SPAD was focused; this laser spot was created by aiming our high-power pulsed laser to a second diffusely-painted wall in the laboratory, and
- our laboratory’s ceiling lights, the closest of which was approximately two meters away from where the SPAD was focused; these lights exhibited a very faint high-frequency flicker due to their electronic high-efficiency ballasts.

Acquisition: We used an exposure of two seconds for timestamp stream acquisition. The total number of photons detected was 940k.

Frequency scanning and detection: We probed frequencies from 0 Hz to 100 Hz with a step size of 0.5 Hz to avoid detection of frequencies corresponding to the sinc lobe centered at DC. In addition, we probed frequencies from 100 Hz to 17 GHz with a step size of 0.3 Hz. As in the experiment of Figure 1, we set the probability α of false alarm to $\frac{1}{\text{\#freqs}}$ for each probed range, where #freqs is the number of probed frequencies.

The recovered spectrum of the flux function is shown in Figure 7. It is evident from the spectrum that we can passively reconstruct flux variations well above 10 GHz with our SPAD: the maximum frequency detected was 16.9 GHz. Moreover, despite the fact that the 85.1 kHz flicker emitted our laboratory’s LED lighting was very faint, it was still recoverable despite our laser’s simultaneous emissions. Lastly, note that the spectrum contains a single outlier at 85 Hz, as its amplitude was just above the bound set by the CFAR detector.

Flux function visualization: We plot the reconstructed flux function at two timescales—a 60 μ s temporal window (Figure 8, left) and a 2.5 ns window (Figure 8, right). The former window captures the component of the flux function due to the flickering ceiling lamps; the latter shows the reconstruction of a single laser pulse whose estimated FWHM is 58.9 ps. To create these plots for each window, we evaluate the flux function at 1k regular samples, with frequency cutoffs at 8.3 MHz and 200 GHz, respectively, to avoid aliasing.

Spectrum of room-scale flux function

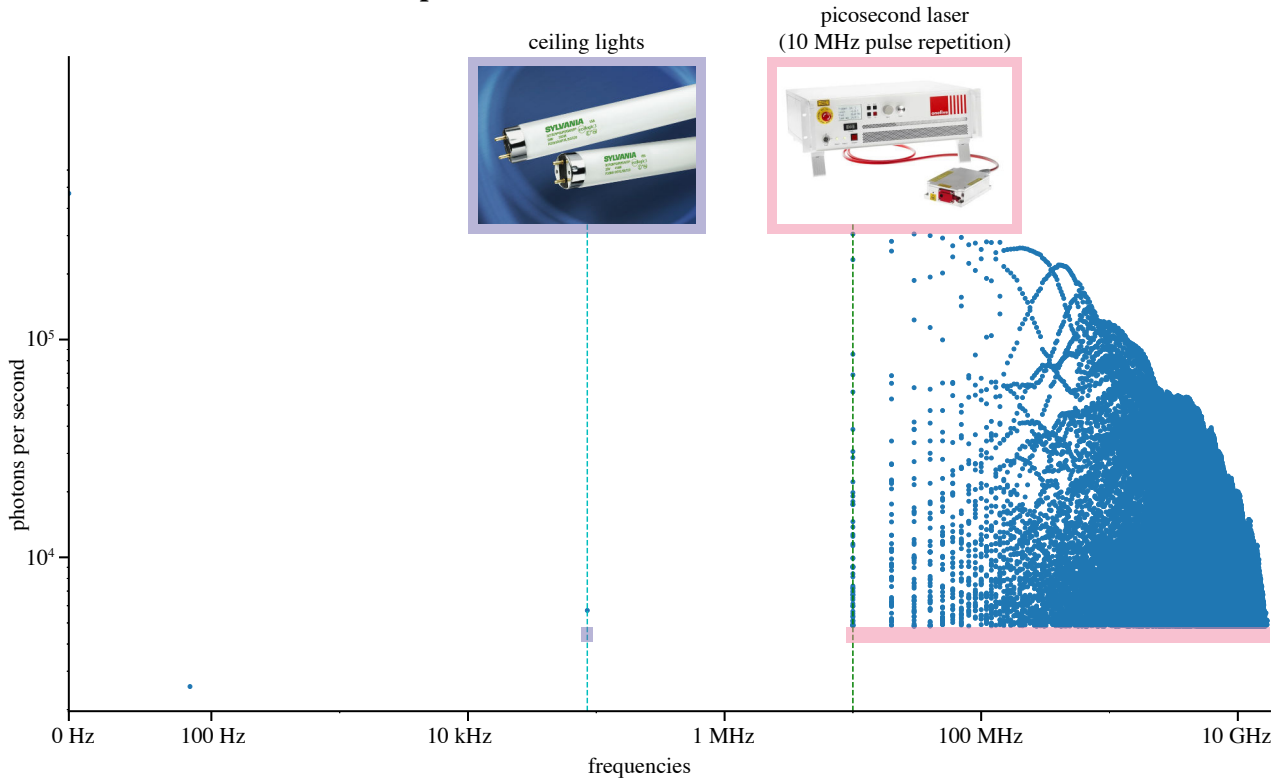


Figure 7: Spectrum of the reconstructed flux function from our room-scale passive ultra-wideband imaging experiment. Our method successfully reconstructs flux due to the flickering LED ceiling lamps (85.1 kHz), the dominant frequency of the pulsed laser (10 MHz), as well as the laser’s higher harmonics (up to 16.9 GHz).




Figure 8: Flux function visualization for our room-scale experiment. **Left:** Flux component due to the LED ceiling lamps’ faint (2.4%) flicker at 85.1 kHz. **Right:** Flux component due to a picosecond laser pulse (58.9 ps FWHM).

E.2.3 Passive reconstruction of laser pulse trains

Here we demonstrate passive unsynchronized reconstruction of picosecond-scale laser pulse trains with high accuracy.

Scene: Our scene consists of a planar target, approximately one meter away from the SPAD, illuminated by a diffused low-power pulsed laser operating at 20 MHz. The SPAD is focused on a single point on the planar target. We use the same setup as in Figure 1, but with only one laser source turned on.

Active histogram-based imaging: To acquire a conventional photon-count histogram that captures the shape of the laser pulse, we synchronized the SPAD to the laser and operated the PicoHarp as a Time-Correlated Single-Photon Counter. We used a one second exposure for this procedure, detecting 267k photons in total.

Flux probing: We used the first 2657 raw (*i.e.*, absolute-time) timestamps from the PicoHarp as a timestamp stream for the purpose of probing. This corresponds to an acquisition time of approximately 10 milliseconds.

Frequency scanning and detection: We probed frequencies from 0 Hz to 11 GHz with a step of 60 Hz. This frequency upper bound was motivated by the simulations in Section G.1 which indicated that a maximum frequency of 11 GHz is sufficient to resolve pulses with a FWHM close to our laser’s 83 ps FWHM. As in the experiments described previously, we set the probability α of false alarm to $\frac{1}{\text{\#freqs}}$, where #freqs is the number of probed frequencies.

Pulse train visualization: We sample the flux function at 1 million equidistant points across one 50-nanosecond period of the laser. This corresponds to a sampling rate of 20 THz. Figure 9 (top) depicts the reconstructed laser pulse train obtained from our method. We note that reconstruction of such a train is not possible using a conventional photon-count histogramming technique because their reliance on relative times effectively time-wraps the entire pulse train into a single pulse. Figure 9 (bottom-left) shows a zoomed-in view of an individual pulse from the pulse train, reconstructed by flux probing (estimated FWHM: 76 ps, peak location: 4.656 ns) along with the reconstruction obtained from histogramming (estimated FWHM: 72 ps, peak location: 4.648 ns). Figure 9 (bottom-right) shows a zoomed-in view of these reconstructions. Overall, the shapes are very similar even though probing uses $100\times$ fewer photons. Additionally, we observe a slight 8 ps shift between the pulses reconstructed by probing versus histogramming. This shift is well below the 16 ps jitter of our TCSPC and thus not significant.

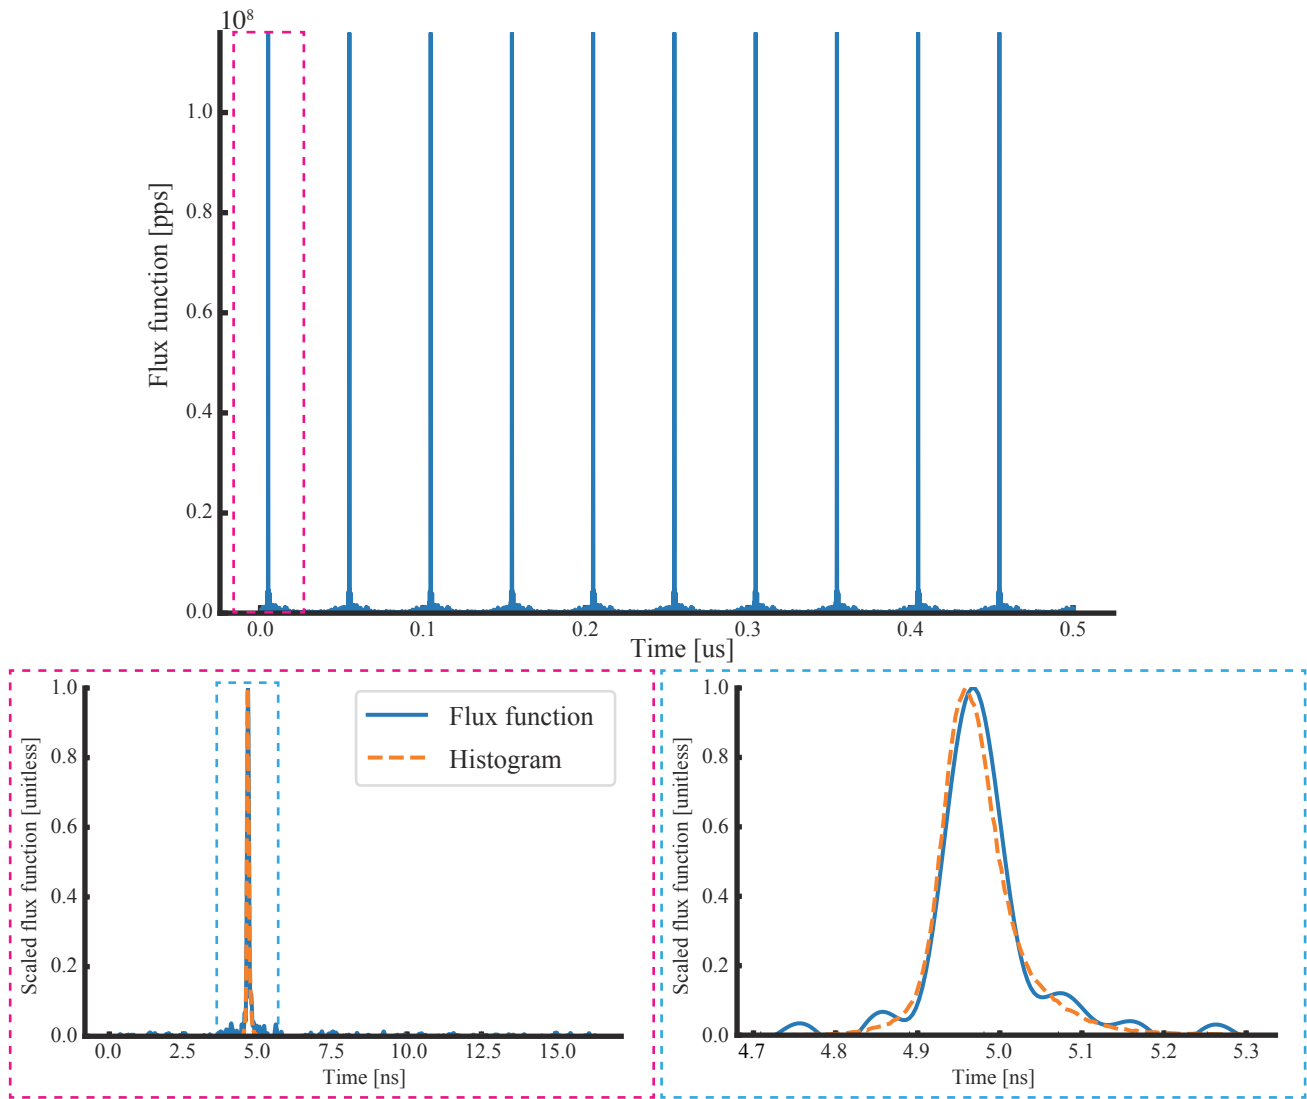


Figure 9: Passive reconstruction of picosecond laser pulse trains. **Top:** Reconstructed laser pulse train. **Bottom left:** One reconstructed pulse from the train, along with the single pulse reconstructed by active photon histogramming. **Bottom right:** Zoomed-in view of the laser pulses shown on the left.

E.3 Passive ultra-wideband sensing of 2D signals

E.3.1 Fan experiment: freezing time at all timescales

We provide additional details on the fan experiment shown in Figure 5 (top).

Scene: The scene itself is shown in Figure 10. It consisted of a battery-operated spinning fan and two white cardboard panels placed at a right angle. The scene was approximately 400 mm × 250 mm in size and approximately one meter away from the SPAD.

Illumination: We operated the laser at a pulse repetition rate of 20 MHz for this experiment. A Thorlabs ground-glass diffuser was placed in front of the laser so that the whole scene was illuminated by a diffuse ultrafast point source.

Scanning and timestamp stream collection: Since we only have a single-pixel SPAD available, we used the galvo mirrors to scan the field of view with a spatial resolution of 3 mm, resulting in an image with a resolution of 151 by 91 pixels. We used an exposure time of $t_{\text{exp}} = 0.1$ s for timestamp acquisition at each galvo position. Photon counts ranged from approximately 13k photons at the brightest pixels (*i.e.*, on the bottom left point of the fan) to 5k photons at pixels on the background white plane, to far fewer photons elsewhere in the scene.

Temporal pixel alignment: In order to simulate concurrent imaging, the flux functions of all pixels must be computationally shifted relative to each other to create the effect of a common absolute clock. We use the laser sync to provide temporal alignment of the flux functions within the laser repetition period. However, to ensure alignment of the pixels’ absolute times, we need to also have a correspondence between pulses, *i.e.*, that the i -th pulse of pixel a corresponds to the j -th pulse of pixel b . For that, we use a collimated DC laser diode that is interrupted by the fan blades at each rotation. This provides a timing signal that is sufficient for such a temporal pixel alignment. In particular, we consider a pixel’s pulse 0 to be the pulse closest to the time instant when the fan interrupts the DC laser beam. To detect the on-off light signal from the DC laser, we use the avalanche photodiode (APD) connected to an oscilloscope that generates temporal markers at each on-off transition, which are then provided to the TCSPC. After reconstructing the flux function at each pixel independently, we use those markers to align all the flux functions onto a common absolute timeline.

Frequency scanning and detection: For the sake of efficiency, the frequency-scanning step of the algorithm in Figure 3 was performed for only one pixel in the scene, and the set of flux frequencies detected for that pixel were used to probe flux at all other pixels as well. Specifically, we selected the brightest pixel on the fan blades and scanned frequencies from 0 Hz to 10 GHz with a step size of 6 Hz. As usual, we set the probability α of false alarm to $\frac{1}{\text{\#freqs}}$, where \#freqs is the number of probed frequencies.

We used the probing measurements at the brightest pixel on the fan blade to determine the fan’s spinning frequency, f_{fan} , and the laser’s dominant frequency, f_{laser} . Note that since the fan’s two blades block the APD twice per full rotation, flux probing detected the frequency $f_{\text{fan}} = 108$ Hz, which was twice the fan’s actual rotation. Using f_{fan} and f_{laser} , we expanded the set of probed frequencies to also include (1) higher harmonics of the laser’s dominant frequency f_{laser} (*i.e.*, integer multiples of f_{laser} up to 15 GHz); (2) higher harmonics of the fan’s spinning frequency (*i.e.*, integer multiples of f_{fan} up to 100 kHz); and (3) linear combinations of f_{fan} , f_{laser} and their harmonics, computed by convolving the harmonics of f_{fan} and f_{laser} .

Rendering ultra-wideband video: The reconstructed flux functions are continuous, and include variations from millisecond to picosecond timescales. To render discrete video frames at the finest resolution, we integrate the flux function over consecutive intervals of four picoseconds, corresponding to a frame rate of 250 billion frames per second. At the coarsest timescale, frames are computed by integrating the flux function over one-millisecond intervals.

Acquired ultra-wideband videos: Since the flux functions reconstructed at each pixel are ultra-wideband, it is possible to “zoom in” and “zoom out” in time across vastly different timescales:

- *1 kfps*: The fan blades are spinning and perform five full rotations in 0.1 seconds. At this timescale, the fan appears to be illuminated by temporally- and spatially-constant illumination from the laser. Figure 11 shows 9 images at this timescale. Note the shadow of the fan’s blade on the background plane (middle row).
- *250 Gfps*: At the finest timescale, each laser pulse can be seen traversing the scene individually, with one laser pulse sweeping through the scene in one nanosecond. The fan blades appear to be frozen, and the pulse appears to propagate over them. Note that we can render such a video sequence for each individual laser pulse, as the fan blades slowly change position. Figure 12 shows 9 frames at this timescale, where a single laser pulse traverses the scene while the fan blades appear frozen. Figure 13 shows the results obtained from histogramming the relative timestamps as obtained from the TCSPC. Our method provides sharper transients and can also recover the blades of the fan, whereas conventional photon histogramming blurs them. Figure 14 shows a different laser pulse propagating through the scene 2.2 milliseconds later, hence the different positions of the fan’s blades.

To further illustrate the capability of “zooming in” and “zooming out” in time, we capture ultra-wideband video in color and incrementally “zoom in” to smaller and smaller timescales to observe the following effects:

- *10 fps*: At this frame rate the fan’s blades spin too fast to resolve, so the fan appears to be static and its blades are not distinguishable. The scene appears to be illuminated by a temporally- and spatially-constant light source even though it is actually a pulsed laser.
- *300 fps*: The fan’s blades are motion blurred at this timescale, making them difficult to distinguish. The scene appears to be illuminated by a temporally- and spatially-constant light source.
- *1 kfps*: Fan blades can be seen spinning, performing five full rotations in 0.1 seconds. At this timescale, we can observe the moment at which the edge of the fan blade starts to block part of the green DC laser beam (Figure 15 bottom right). The reflected light from the green DC laser appears as a spot in the back white plane.
- *2 kfps*: Fan blades appear to decelerate. The trajectory of the green laser’s inter-reflection starts to become noticeable on the rear plane, though it is still fairly rapid. The scene still appears to be illuminated by temporally- and spatially-constant illumination from the pulsed laser.
- *10 kfps*: With a further slowdown of the fan blades, the trajectory of the reflected green laser on the back plane becomes distinctly visible. The scene still appears to be illuminated by temporally- and spatially-constant illumination from the pulsed laser.
- *100 kfps*: The fan blades slow down even more, appearing to move very slowly.
- *30 Mfps*: At this timescale, the flicker from the pulsed laser becomes visible, while the light emitted from the green laser is too dim to notice. The fan blades appear to be stationary.
- *200 Mfps*: At this timescale, the flicker from the pulsed laser appears to exhibit slight spatial variations.
- *10 Gfps*: Spatial propagation of the laser’s pulses is now clearly visible.
- *250 Gfps*: At the finest timescale, each pulse can be seen travelling through the scene individually, with one laser pulse sweeping through the scene in one nanosecond. The fan blades appear to be frozen, and the pulse appears to propagate over them (Figure 16).

For more visualizations please refer to the supplemental video.

Additional details about the color ultra-wideband video experiment: We captured the scene using color filters. Specifically, we used red, green, and blue color filters, scanning the scene sequentially by placing one color filter in front of the SPAD at a time. To account for the time-consuming scan, we operated the fan with an AC power supply. Because the AC power leads were reversed relative to the fan’s internal battery, the fan spun counter-clockwise for this experiment. To reconstruct the flux function, we repeated the same procedures outlined in the previous experiment. Due to the imperfect response of the color filters, we applied manual scaling to the red, green, and blue channels. In particular, we multiplied the red channel by 1, the green channel by 0.75, and the blue channel by 0.4.

Fig




Figure 10: Setup for the spinning fan experiment. Left: Synchronization setup. A collimated laser diode illuminates the fan from the back providing an on-off signal to the APD. Right: Photo of the experimental setup captured from the SPAD's viewpoint.



Figure 11: Images show various stages of a fan spinning, rendered at 1 kfps.



Figure 12: Images show a picosecond laser pulse propagating over two planes and a fan. Note how the fan blades appear as frozen. The effective frame rate is 250 Gfps.




Figure 13: Images show a picosecond laser pulse propagating over two planes and a fan obtained from histogramming using the frequency of the laser. Note how the fan blades appear blurred.



Figure 14: Images show a picosecond laser pulse propagating over two planes and a fan. The fan blades are at a different position demonstrating the *unique* capability of our approach to *freeze time at all timescales*.




Figure 15: Color images of a fan spinning rendered at 1 kfps. Note how we can accurately recover the reflection of the green laser from the fan blade to the back plane (top left) and the moment at which the edge of the fan blade blocks the green laser beam (bottom right).




Figure 16: Color images depicting the propagation of a picosecond laser pulse across two planes and a fan. At this timescale, only the red pulsed laser contributes to the illumination of the scene.

E.3.2 Coca-Cola bottle experiment

We provide additional details about the Coca-Cola bottle experiment shown in Figure 5 (rows 2-4).

Scene: The scene itself is shown in Figure 5 (row 2, left). It consisted of a Coca-Cola bottle filled with water and a small amount of milk, as well as a CFL light bulb connected to an AC outlet. The scene was approximately 392 mm×121 mm in size and approximately one meter away from the SPAD.

Illumination: We operated the laser at a 20 MHz pulse repetition rate for this experiment. Due to its 60 Hz AC power supply, the CFL bulb flickered at 120 Hz. Both light sources were turned on during acquisition.

Scanning and timestamp stream collection: As in the ultra-wideband video acquisition experiment of Section E.3.1, we used the galvo mirrors to scan the field of view with a spatial resolution of 3 mm, resulting in an image with a resolution of 392 by 121 pixels. We used an exposure time of $t_{\text{exp}} = 0.1$ s for timestamp stream acquisition at each galvo position. Photon counts ranged from approximately 37k photons at the brightest pixels (*i.e.*, near-specular transmission from the bottle’s surface to the SPAD) to 17k photons at pixels on the CFL bulb, to far fewer photons elsewhere in the scene.

Temporal pixel alignment: Since it was not possible to acquire timestamps simultaneously from all pixels in the field of view, we required an external mechanism to align the acquired flux functions across the field of view. To do this, we used synchronization signals from the laser and the bulb to align the estimated flux functions at picosecond and millisecond timescales. Specifically, we used the sync signal of the laser and an oscilloscope connected to the AC to produce timing markers that were detected by the PicoHarp. After reconstructing the flux function at each pixel independently, we used those markers to align all the flux functions onto a common absolute timeline.

Frequency scanning and detection: For the sake of efficiency, the frequency-scanning step of the algorithm in Figure 3 was performed for only two pixels in the scene and the flux frequencies detected for these pixels were then used to probe the flux frequencies at all other pixels. Specifically, we selected (1) a pixel on the CFL bulb and scanned frequencies from 0 Hz to 10 kHz with a step size of 10 Hz and (2) a pixel on the Coca-Cola bottle and scanned frequencies from 10 kHz to 10 GHz with a step size of 6 Hz. As usual, we set the probability α of false alarm to $\frac{1}{\#\text{freqs}}$.

Rendering ultra-wideband video: The reconstructed flux functions are continuous, and include variations from millisecond to picosecond timescales. To render discrete video frames at the finest resolution, we integrate the flux function over consecutive intervals of 5 ps, corresponding to a frame rate of *200 billion frames per second*. To improve visibility, we scale each video frame’s intensity by a factor of two. For coarser timescales, frames are computed by integrating the flux functions over coarser time intervals. This makes it possible to “zoom in” and “zoom out” in time across vastly different timescales. We demonstrate this with an accompanying video in which we incrementally “zoom in” to smaller and smaller timescales in order to observe the following effects:

- *3 kfps*: A CFL bulb flickers 12 times in 0.1 s. At this timescale, the bottle appears to be illuminated by temporally- and spatially-constant illumination from the laser.
- *3.6 Mfps*: At this timescale, the light emitted by the bulb is constant and the bottle still appears to be illuminated by temporally- and spatially-constant illumination from the laser.
- *300 Mfps*: The light emitted by the bulb per frame is too dim so the bulb appears dark. In contrast, the laser’s pulsing becomes evident, with 20 laser pulses illuminating the scene in 1 microsecond. The timescale is still too coarse, however, for the propagation of the individual laser pulses through the scene to be noticeable. For this reason, the laser light appears as a spatially-invariant flicker that illuminates the entire scene simultaneously.
- *600 Mfps*: As the frame rate increases further, the laser pulses appear to slow down. Here, ten laser pulses illuminate the scene in 500 microseconds.
- *2.4 Gfps*: At even higher frame rates, the laser pulses arrive even more slowly. Here, only five laser pulses illuminate the scene in 250 microseconds. The illumination begins to gain some spatial variability, due to differences in time of flight.
- *200 Gfps*: At the finest timescale, each pulse can be seen traversing through the scene individually, with one laser pulse propagating through the bottle in two nanoseconds. Figure 17 shows 10 frames at this timescale, where a single laser pulse traverses the scene. Note that even though we have reconstructed all two million individual traversals of laser pulses through the scene during the acquisition period, they are spaced too far apart—50 nanoseconds between them—to include more than one traversal in the supplementary video.




Figure 17: Passive transient imaging. Images show a picosecond laser pulse propagating inside a coca-cola bottle filled with water and a small amount of milk. Note how we reconstruct the light bouncing off the cap of the bottle and propagating all the way back to where it entered (last row right).

E.4 Recovering passive non-line-of-sight video

We provide additional details about the passive NLOS experiment shown in Figure 5 (rows 2, 5 and 6).

Scene: The scene is shown in Figure 5 (row 2, second image from right) except that the two lasers were turned off. A Thorlabs ground-glass diffuser is placed in front of the laser projector so that light is diffused from each projector pixel as would happen with light bouncing off of a line-of-sight (LOS) wall. The SPAD was focused on a single point on a planar target as shown in the photo.

Illumination: We used the projector to play a five-second MGM video clip in a loop at a nominal frame rate of 60 fps. The projector was the only light source in this experiment.

Acquisition: We acquired a single timestamp stream over an exposure time $t_{\text{exp}} = 5$ seconds. A total of 959k photons were detected.

Frequency scanning and detection: In order to improve the computational efficiency of our method, we split the acquired timestamps into five one-second streams, each corresponding to a one-second exposure. The resulting streams contained 214k, 180k, 162k, 190k, and 213k timestamps, respectively. For each clip, we probed frequencies up to 40 MHz with a step of 1 Hz. We initially set the probability α of false alarm to $\frac{1}{\#\text{freqs}}$, as in our other experiments. We observed, however, that setting the probability of false alarm to a slightly higher value, namely $\alpha = 0.3\%$, yielded a higher-quality reconstruction. We believe this is related to the fact that as the number of photons increase the CFAR bound also increases, making weaker flux frequencies harder to detect (see the discussion in Section 3.1 of the main paper).

Reconstructing individual video frames from the reconstructed flux function: We used the reconstructed flux function to reconstruct the video frames as seen from the projector’s point of view. This required converting the high-frequency 1D flux signal from the SPAD into the 2D image that the projector actually projected by raster-scanning onto the diffuser. We briefly describe this procedure below.

First, we used the detected frequencies to estimate the projector’s actual refresh rate. This was found to be 58.7 Hz, *i.e.*, slightly below the nominal video rate. This refresh rate defines the start time and the end time of each frame. Given these times, we regularly sampled the flux function at 1.177M (1280×920) points. While the resolution of the laser projector is 1280×720 , we sample with a resolution of 1280×920 to compensate for the time taken by the scanning laser to travel from the end of the raster scan to the beginning. This provides the flux of each pixel of each frame, which can then be remapped onto a 2D image plane. Second, we observed that the projector’s refresh rate drifted over time, causing the video frames to go slightly out of sync. Due to the raster scanning pattern of the laser projector, this drift also causes the odd and even rows to be slightly misaligned. To compensate for this, we manually searched around the detected frequency for an “average” refresh rate which we found to be 58.70893, which produced stable frames in the entire acquisition interval. Moreover, only the odd rows of the reconstructed video frames are shown in order to improve visual clarity.

To calculate the average number of photons per frame in each video clip, we divide the total number of collected timestamps by 60.

NLOS video reconstruction results: Figures 18, 19, and 20 compare the results of the above NLOS video reconstruction procedure to the video frames actually projected by the projector, along with the frames estimated by the timestamp binning procedure outlined in Section G.3.3, for three different video frames. Despite the low number of detected photons, the imaging speeds involved, and the video frames’ complexity, flux probing yields very good reconstruction results when compared to the ground truth. In comparison, timestamp binning, which does not reconstruct the flux function, produces poor results and makes clear that the photons detected per frame are not sufficient to obtain a high-quality image. Indeed, the contrast of frames reconstructed by timestamp binning is limited to just one photon because the time it takes for the raster scan and laser beam to move from one pixel to another is less than the SPAD’s dead time. As a result, no more than one photon can ever be detected per pixel in a single video frame.

The above results point to the somewhat surprising property of our approach: even though the flux function varies considerably from video frame to video frame in this experiment, the timestamps produced in *all* frames of the video contribute to the reconstruction quality of a given frame’s flux function, not just the timestamps acquired in that particular frame. Please see the supplemental video for a better visualization of the reconstructions.




Figure 18: NLOS video reconstruction. Example reconstructed frames from probing and timestamp binning for the video sequence in $[1, 2)$ seconds.




Figure 19: NLOS video reconstruction. Example reconstructed frames from probing and timestamp binning for the video sequence in $[2, 3)$ seconds.




Figure 20: NLOS video reconstruction. Example reconstructed frames from probing and timestamp binning for the video sequence in $[4, 5)$ seconds.

E.5 Probing with SPAD arrays

We provide some additional details on the 2D SPAD array experiment shown in Figure 5 (rows 7-8).

Scene: We applied our approach to the dataset in [9] that was graciously provided by its authors. The maximum number of timestamps acquired per pixel is 100k.

Frequency scanning and detection: We sweep frequencies from 0 Hz to 100 kHz with a step size of 5 Hz. As usual, we set the probability α of false alarm to $\frac{1}{\text{\#freqs}}$.

Flux probing with non-negligible dead time After detecting the frequencies, we recover the amplitude and phase of each Fourier coefficient using Algorithm 1. We notice that compensating for dead time alters the magnitude of the flux function significantly, increasing the dynamic range of the recovered video frames. This effect is also depicted in Figure 21.

Flux reconstruction results: We use flux probing with non-negligible dead time to recover the flux function at each pixel of the SPAD array. In the supplemental video, we show a 50 kfps video of the spinning fan that was rendered from our continuous flux functions, along with the videos reconstructed from Seets et al. [9]. Despite the fact that our processing involves strictly pixel-wise reconstruction of flux functions, our method recovers sharper videos of the scene. Since our flux functions are continuous, we can render video at higher frame rates as well. In our supplementary videos, we also render slow-motion video (5 Mfps) of the fan for both our approach and the “changepoint” video from [9]. We observe that the motion of the fan from [9] appears discontinuous whereas in our video the motion is smooth and the fan appears sharp as well.

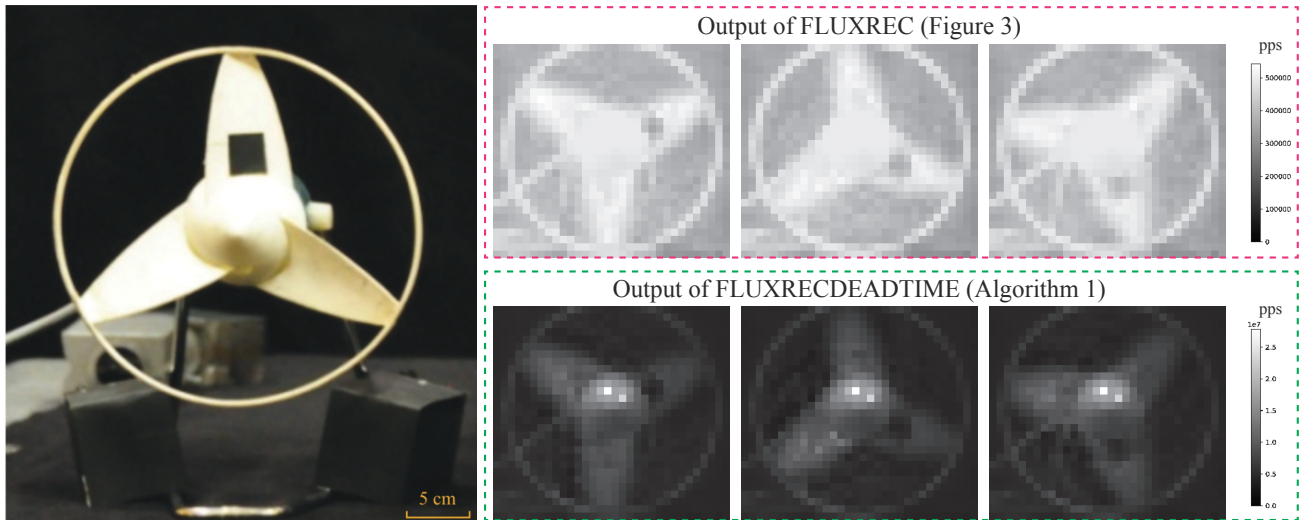


Figure 21: Probing with SPAD arrays. We apply our approach to the dataset from Seets et al. [9], by probing under the assumption of negligible dead time (top right) and non-negligible dead time (bottom right).

F. Supplemental Video

The video presents a summary of our methodology along with video clips showcasing many of the experimental results in Sections 5 and E. A few additional details about specific clips in the video are given below.

Fan experiment: We adjusted the flux of each frame rate sequence by normalizing by the maximum value of the red channel and further processed the sequence by applying tone mapping to the normalized flux, *i.e.*, taking its power to $\frac{1}{2.2}$. The time codes for the corresponding video clip are **05:10 - 05:29**.

Coca-cola bottle experiment: We adjusted the flux of each frame rate sequence by normalizing by its maximum value and brightening the video sequence by a factor of two. The time codes for the corresponding video clip are **05:30 - 06:17**.

A microscope in time: This clip illustrates an interactive visualization of the 1D flux function reconstructed by the experiment in Figure 1. This clip was created by performing an analytical integration and efficient sampling of the function $\hat{\phi}(t)$ defined in Figure 3 using OpenCL, enabling real-time zoom-in and zoom-out in time. To create the interactive plot, the function is sampled at 1024 regular intervals for the time window currently being displayed. Snapshots from this video clip tool are shown in Figures 22 and 23.

- **07:07 - 07:20:** At the span of one second, we can see the 58.7 Hz refresh rate of the laser projector.
- **07:23 - 07:25:** Over tens of milliseconds, we can see the 900 Hz flicker of the light bulb.
- **07:26 - 07:29:** Zooming further, at a couple of milliseconds, we can see the time-varying modulation of the projector’s laser beam.
- **07:32 - 07:34:** At the scale of microseconds, we can observe the multiple pulses of the 3 MHz laser.
- **07:37 - 07:39:** Zooming even further, we can see the pulses from the asynchronous 3 MHz and 40 MHz lasers.
- **07:43 - 07:47:** At the finest timescale, over a couple of nanoseconds, we can resolve an individual laser pulse of the 40 MHz laser.

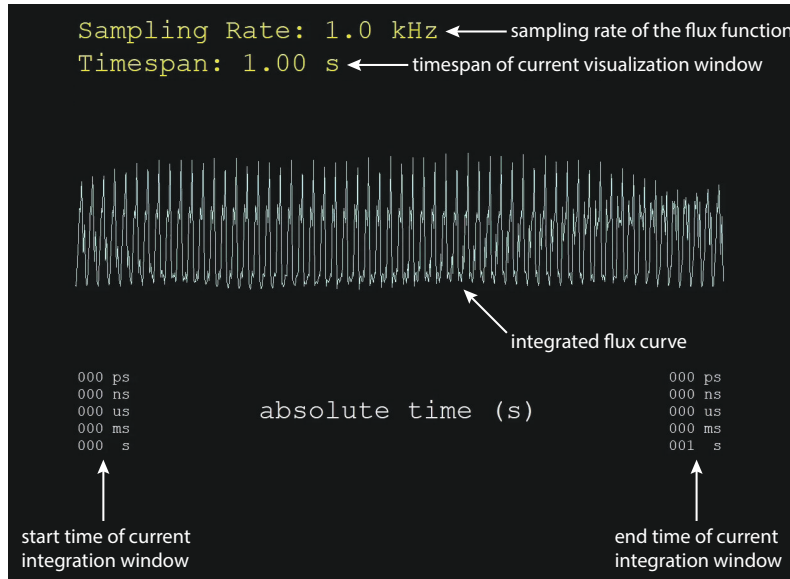


Figure 22: Visualization of interactive flux function using OpenCL.

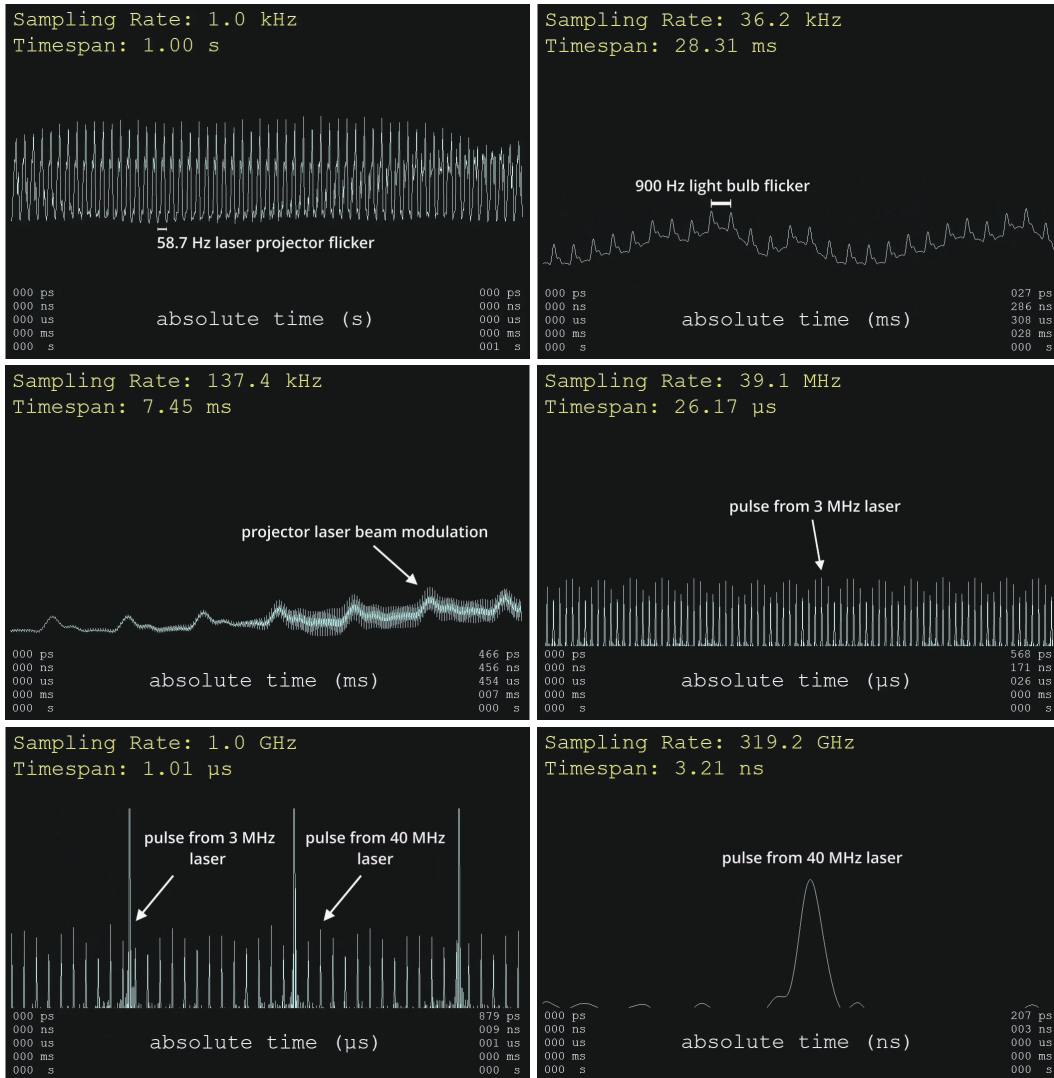


Figure 23: Visualization of interactive flux function using OpenCL.

G. Simulations

In this section, we provide additional details about the simulations in the main paper. These simulations highlight the potential of our theory to push ultra-wideband, low-light imaging far beyond the capabilities of existing imaging systems.

G.1 Simulation of laser pulse trains

G.1.1 Model of laser pulse

We approximate the lobe of the laser pulse as a Gaussian distribution whose standard deviation is $\frac{\text{FWHM}}{2.355}$. This is approximately the standard deviation of a Gaussian with the desired FWHM. We use a repetition rate of 20 MHz so the simulated pulse train has a period of fifty nanoseconds. We position the Gaussian pulse so that its mean is at the period's center, *i.e.*, at $t = 25$ nanoseconds, and truncate it at $[0, 50 \text{ ns}]$ so that consecutive pulses in the train do not overlap. The frequency support of pulses with a FWHM of eight and eighty picoseconds is approximately 12.5 GHz and 125 GHz, respectively. In particular, frequencies outside these ranges have amplitude less than 3% of the peak amplitude.

G.1.2 Timestamp stream generation

We use Çınlar's method [10] to generate individual photon timestamps from the pulse train. This method is preferable over thinning [11] because the narrow temporal support of the laser pulse train makes thinning computationally inefficient.

After timestamp generation, we apply a dead time of 250 ns to remove timestamps that arrive during the SPAD's inactive period. The remaining timestamps are quantized to a resolution of four picoseconds, *i.e.*, the nominal quantization of our experimental system's time-to-digital converter.

For the 80 ps laser pulse train simulation, we use an exposure time $t_{\text{exp}} = 10$ ms and set the average flux to be 200 kpps.

For the 8 ps pulse train simulation, we set the average flux to 500 kpps and vary t_{exp} in the range $[0.1 \text{ ms}, 3.75 \text{ ms}]$. This range of exposure times results in timestamp streams with lengths between 50 and 1875 timestamps on average. These streams were used for the comparisons shown in Figure 4 (third from left).

G.1.3 Estimation methods

We compare three methods for flux function estimation:

- **Flux probing:** We probe up to 12.5 GHz with a step size of 100 Hz for the 80 ps laser pulse simulations. We set the CFAR detector's probability of false alarm to be just one frequency on average out of the resulting samples, *i.e.*, $\alpha = \frac{100 \text{ Hz}}{12.5 \text{ GHz}}$. For the 8 ps laser pulse simulations, we probe up to 125 GHz with a step size of $\frac{1}{t_{\text{exp}}}$, where t_{exp} ranges from 0.1 ms to 3.75 ms. We again set the probability of false alarm so that only one false alarm occurs on average, *i.e.*, $\alpha = \frac{1}{t_{\text{exp}}} \frac{1}{125 \text{ GHz}}$.
- **Oracle-based probing:** We regularly sample the ranges $[0, 12.5 \text{ GHz}]$ and $[0, 125 \text{ GHz}]$ with a sampling rate of 20 MHz. This corresponds to the frequency support of the 80 ps and 8 ps ground-truth laser pulse trains, respectively. We then probe these sets of frequencies and use the probing results to reconstruct the laser pulses.
- **Active histogram-based imaging:** We assume the sync frequency is $f_{\text{sync}} = 20$ MHz and map each timestamp to the relative time interval $[0, 50 \text{ ns}]$. These relative timestamps are then binned using Matplotlib's `hist()` routine, with a bin width of 4 ps. To convert raw photon counts to a flux function expressed in photons per second, we multiply each bin's contents by the total number of pulses emitted during the exposure time as well as by the bin width, *i.e.*, we multiply by the scalar $(4 \text{ ps} \times 20 \text{ MHz} \times t_{\text{exp}})^{-1}$.

To produce the plot shown in Figure 4 (third from left), we randomly generated twenty timestamp streams for 74 exposure times regularly sampled from the range $[0.1 \text{ ms}, 3.75 \text{ ms}]$. For each method, we report the root mean squared error (RMSE) between the estimated and ground-truth flux functions. For a fairer comparison, we sample the flux function returned by each method at the center of each histogram bin.

As can be seen in that plot, oracle-based probing always outperforms both histogramming and probing. Probing, on the other hand, outperforms histogramming except when the total number of timestamps is very small (below 100). In that regime, many probed frequencies fall below the CFAR threshold and therefore are not detected. This results in a worse flux function approximation on average, despite the very coarse reconstruction returned by histogramming (*e.g.*, see Figure 4, second from left).

G.2 Simulations with non-negligible dead time

In this section, we perform simulations in the regime where T_{avg} is comparable to dead time.

G.2.1 Flux function

We define the flux function as DC plus a linear combination of N sinusoidal functions (N varies from 1 to 50) with frequencies uniformly sampled without replacement in the interval [5 kHz, 100 kHz] in steps of 100 Hz. We vary the DC, *i.e.*, average photons per second from 100 kpps to 100 Mpps. The amplitudes of each sinusoid is set to the average photons per second divided by N .

G.2.2 Timestamp stream generation

We scale the flux function for a desired photon timestamp count and simulate timestamps using the scaled flux function and thinning [11]. We use an exposure time of 10 ms. We take dead time and quantization into account as described in Section G.1.2.

G.2.3 Estimation methods

We compare two methods for flux function estimation, namely the algorithm shown in Figure 3 (left) of the main paper and Algorithm 1 in Section C.1.

G.2.4 Results

Figure 24 depicts the flux functions reconstructed using flux probing and flux probing with non-negligible dead time for various photon levels and number of frequency components. We observe that although the shape of the flux function reconstructed by both methods is similar, the flux function's amplitude is underestimated when ignoring dead time. This phenomenon has a significant impact in the dynamic range of recovered flux function as can be seen in the real-world experiment of Figure 21. Finally, we highlight the fact that it is possible to reconstruct flux functions in settings where only 4% of the photons are detected.

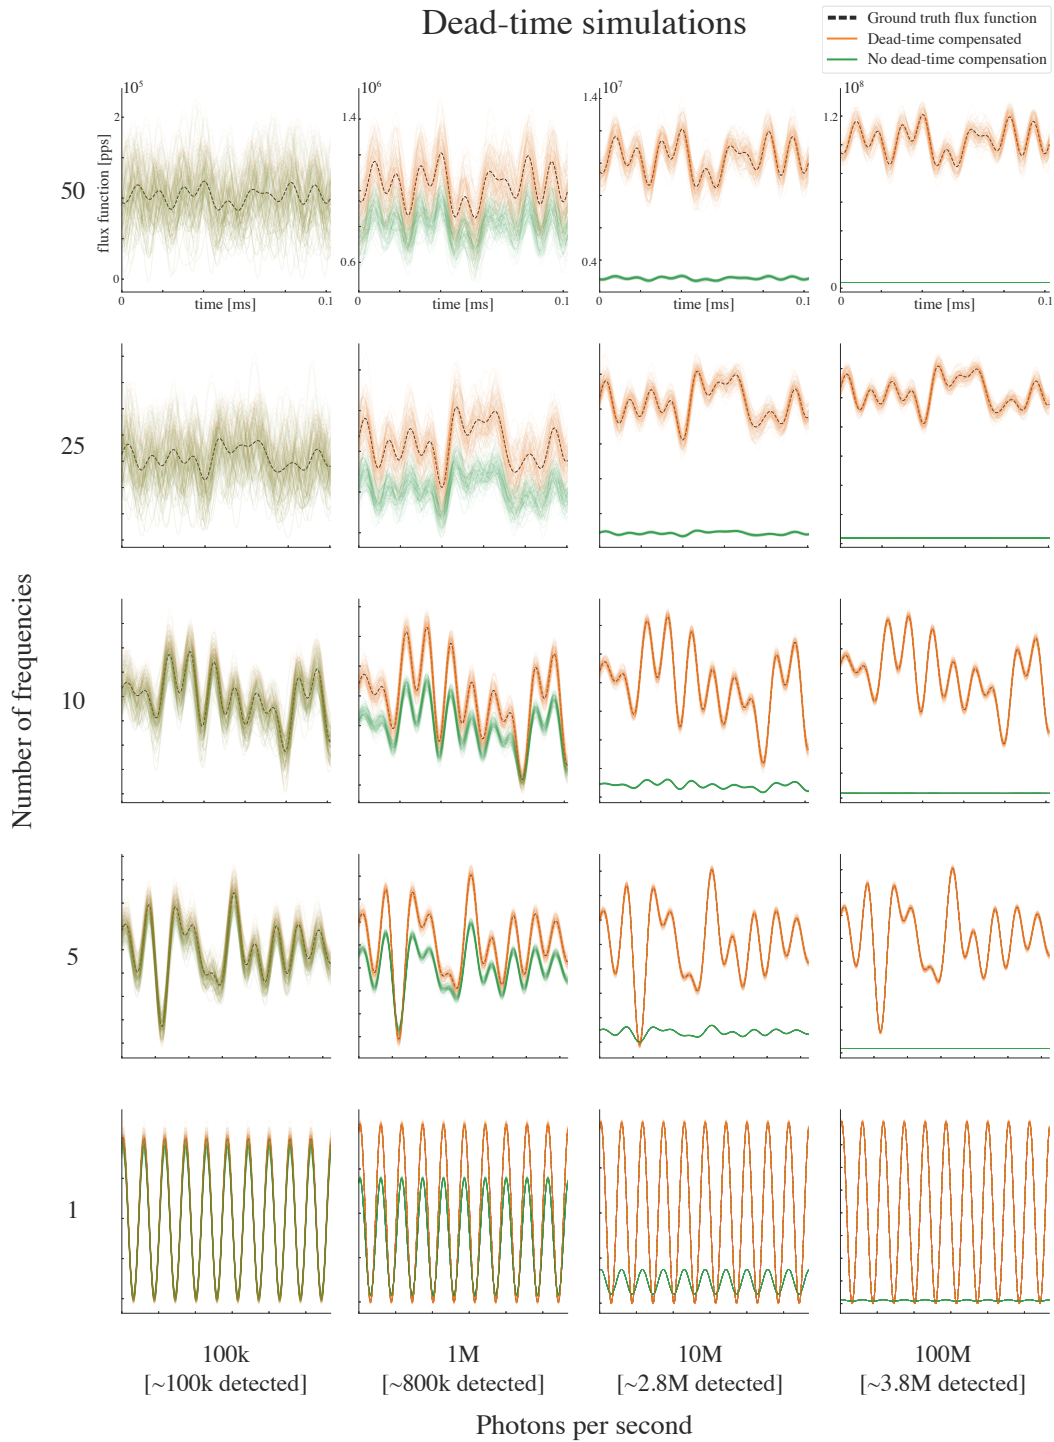


Figure 24: Our theory generalizes to the case where T_{avg} is comparable to dead time. We simulate various flux functions and demonstrate that flux probing with non-negligible dead time allows us to accurately recover the flux function even in the case when 96% of the photons are not detected (last column).

G.3 Non-line-of-sight video simulations

Since we have no access to the actual ground truth for the non-line-of-sight (NLOS) video, we performed simulations of NLOS video reconstruction to provide quantitative results.

G.3.1 Model of flux function

We take 30 frames of video of the MGM lion (Figure 26). Each frame is 720 by 1280 pixels. We vectorize the video to create a 27.648M-dimensional ($720 \times 1280 \times 30$) vector φ , which is our flux function. We treat each element of φ as the flux during a small, 36.17 ns time window that corresponds to the time allocated to each individual projector pixel. We then rescale the vector to achieve an average flux of 2 Mpps.

G.3.2 Timestamp stream generation

We use thinning [11] to generate a single stream of timestamps corresponding to an exposure time of one second and the flux function defined above. To simulate a lower flux level of N pps with $N < 2\,000\,000$, we use Bernoulli thinning with $p = \frac{N}{2\,000\,000}$. We take dead time and quantization into account as described in Section G.1.2.

G.3.3 Estimation methods

We compare three methods for flux function estimation:

- **Flux probing:** We probe up to 7.5 MHz with a step size of 1 Hz. We set the probability of false alarm α to be 0.3% to match the value used in the real-world experiment of Section E.4. We then regularly sample the reconstructed flux function to obtain 27.648M samples spanning the one-second acquisition interval.
- **Oracle-based probing:** We computed the fast Fourier transform of the vector φ and kept the 105735 most significant frequencies according to their amplitude. This corresponds to 0.3% of total frequencies, *i.e.*, the expected number of frequencies that will pass the CFAR detection test. We probe using these frequencies and regularly sample the reconstructed flux function to obtain 27.648M samples spanning the one-second acquisition interval.
- **Timestamp binning:** Each timestamp is mapped to an element of vector φ . Because of the SPAD’s dead time, at most one photon can be detected for each element of φ , *i.e.*, each pixel in the video will contain either zero or one photon. The resulting flux function is therefore an indicator function that describes whether a photon was detected or not.

Each estimated flux function is then reshaped into a $720 \times 1280 \times 30$ video and its values rescaled to the range $[0, 255]$.

G.3.4 Reconstructions

We show the PSNR plot for flux levels in Figure 25. In general, flux probing and oracle-based probing enable recovery of video frames with better contrast than timestamp binning, whose frames are strictly binary. The PSNR plot shows that oracle-based probing performs the best as expected. Note that while timestamp binning has higher PSNR than flux probing for flux levels below 300 kpps, Figures 26 and 27 show that the lion remains visible in frames reconstructed by flux probing whereas this is not the case for frames reconstructed by timestamp binning.

Lastly, the performance of all three methods drops as flux levels increase to the point where dead time becomes significant. This is evident in the reconstructed frames shown in Figure 28, which have lower contrast and a saturated appearance.

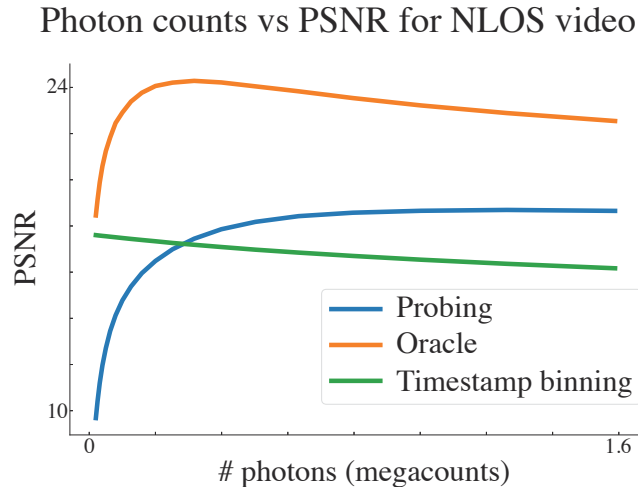


Figure 25: Plot showing number of photons vs PSNR for the three estimation methods.




Figure 26: Simulation comparison for NLOS video reconstructions using 30k photons.

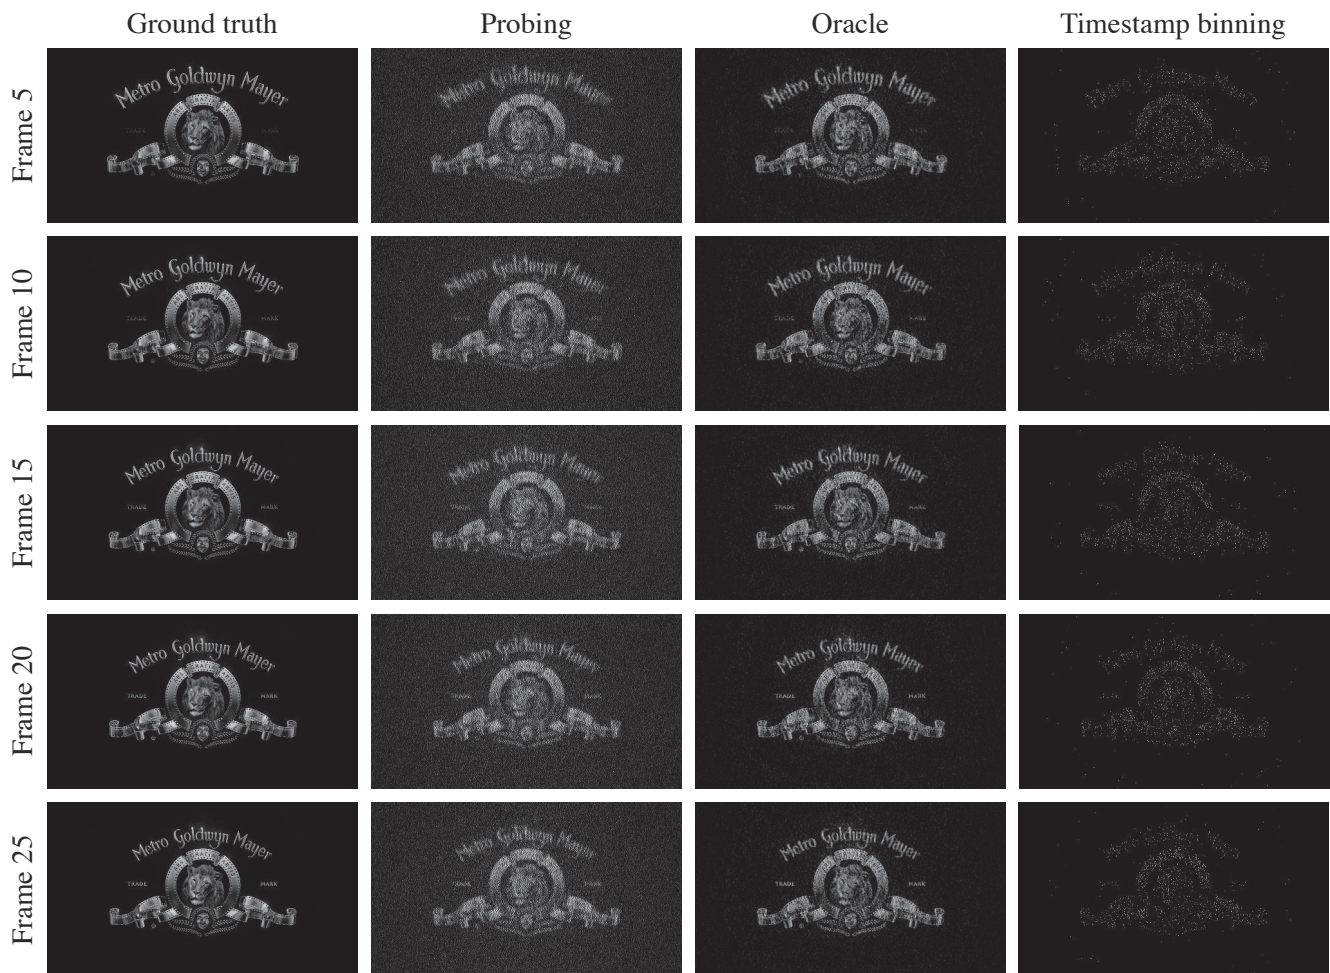


Figure 27: Simulation comparison for NLOS video reconstructions using 200k photons.

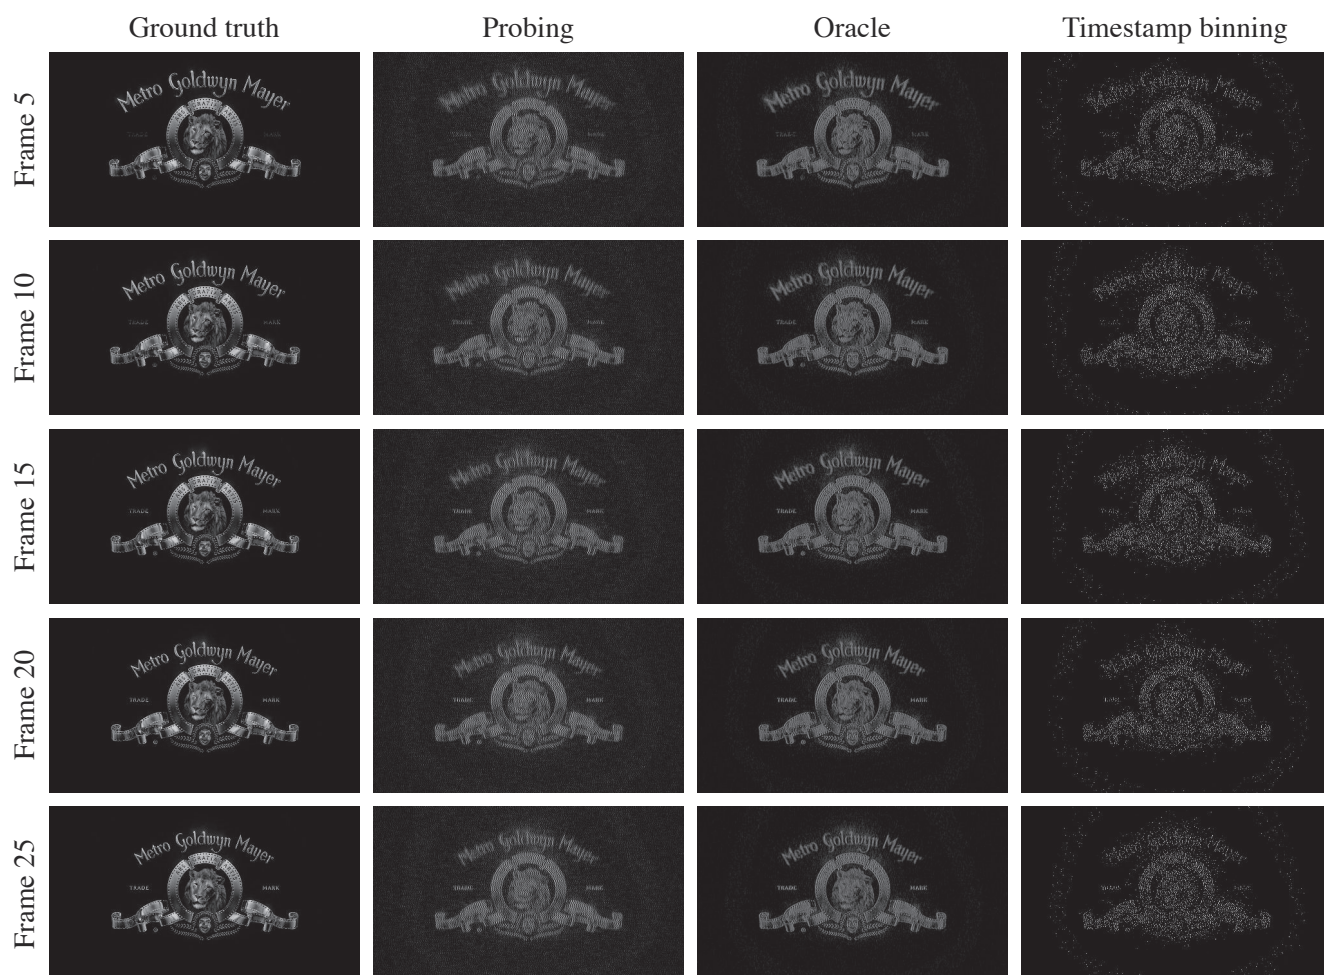


Figure 28: Simulation comparison for NLOS video reconstructions using 1.6M photons.

G.4 Normal distribution of probing measurements

As the number of timestamps increases, probing measurements tend to approximate a normal distribution (Proposition 3). This is confirmed by the simulations shown in Figure 29, where we simulate the distribution of Fourier probing measurements for various numbers of timestamps and probed frequencies, for a given flux function.

More specifically, we generate a stream of 3.5 million timestamps and thin them into separate timestamp stream realizations of either 2, 5, 10, or 50 timestamps that span the entire acquisition interval. We then use these timestamp streams to probe various frequencies, and we estimate the distribution of the resulting probing measurements. Figure 29 overlays their empirical distribution and the probability density function for a Gaussian centered at the ground-truth mean and variance (calculated by probing the ground-truth flux function). For probing measurements with at least 10 timestamps, the joint and marginal distributions of the complex Fourier probing measurements show close agreement with the normal distribution.




Figure 29: Distribution of probing measurements. We generate a timestamp stream according to the flux function shown in the upper left and use a thinning procedure to obtain a collection of Fourier probing measurements for various frequencies (columns) and timestamp stream realizations containing either 2, 5, 10, or 50 timestamps (rows). The empirical joint and marginal distributions of the real and imaginary measurements are shown, along with contour plots of a normal distribution whose mean and variance is calculated using the ground-truth flux function.

References

- [1] S. N. Cohen and R. J. Elliott, *Stochastic Calculus and Applications*, vol. 2. Birkhäuser New York, 2015. 3, 5, 8
- [2] M. Jeanblanc, M. Yor, and M. Chesney, *Mathematical Methods for Financial Markets*. Springer London, 2009. 5, 7, 8
- [3] Eyer, L. and Bartholdi, P., “Variable stars: Which nyquist frequency?,” *Astron. Astrophys. Suppl. Ser.*, vol. 135, no. 1, pp. 1–3, 1999. 6, 15
- [4] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 1991. 8
- [5] A. V. Oppenheim, J. R. Buck, and R. W. Schaffer, *Discrete-time Signal Processing*. Prentice Hall, 2001. 14
- [6] 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. 16, 17
- [7] P. Brémaud, *Point Processes and Queues: Martingale Dynamics*. Springer, 1981. 17, 19
- [8] D. L. Snyder and M. I. Miller, *Random Point Processes in Time and Space*. Springer New York, 2012. 17, 18, 21
- [9] 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. 45
- [10] E. Çinlar, *Introduction to Stochastic Processes*. Prentice Hall, 1975. 48
- [11] P. W. Lewis and G. S. Shedler, “Simulation of nonhomogeneous poisson processes by thinning,” *Nav. Res. Logist. Q.*, vol. 26, no. 3, pp. 403–413, 1979. 48, 49, 51