MITSUBISHI ELECTRIC RESEARCH LABORATORIES
https://www.merl.com

Unambiguous Range Extension for1 Doppler Single-Photon Lidar

Kitichotkul, Ruangrawee; Rapp, Joshua; Ma, Yanting; Mansour, Hassan

TR2026-050 May 07, 2026

Abstract

Single-photon lidar (SPL) measures the time delay between a laser pulse emission and a single9 photon detection. SPL typically uses a periodic sequence of illumination pulses to improve the10 accuracy of distance estimates, and the radial velocity of a moving target can also be determined11 from the Doppler shift in the pulse repetition frequency. Although increasing the repetition12 frequency increases the photon detection rate, it also decreases the unambiguous range. Several13 alternative pulsing modes have been proposed to break this tradeoff for SPL measurements of14 static scenes, e.g., using multiple repetition rates or random pulse trains to yield both high count15 rates and a long unambiguous range. In this work, we show that velocity can still be estimated16 despite the use of non-periodic range extension pulse patterns. We derive a general model for the17 photon acquisition process with a moving target and propose maximum likelihood estimators18 (MLEs) to recover the distance, velocity, and photon flux levels. Each range extension mode19 requires a tailored initialization scheme in order for our MLE solver to converge to the global20 optimum. We demonstrate through simulations and experiments that non-periodic SPL can21 simultaneously achieve long-range and high-accuracy distance and velocity estimates.

Optics Express 2026

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Unambiguous Range Extension for Doppler Single-Photon Lidar

RUANGRAWEE KITICHOTKUL,[1,2] JOSHUA RAPP,[2,*] YANTING MA,[2] AND HASSAN MANSOUR[2]

[1]*Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215, USA*
[2]*Mitsubishi Electric Research Laboratories, Cambridge, MA 02139, USA*
[*]*rapp@merl.com*

Abstract:

Single-photon lidar (SPL) measures the time delay between a laser pulse emission and a single photon detection. SPL typically uses a periodic sequence of illumination pulses to improve the accuracy of distance estimates, and the radial velocity of a moving target can also be determined from the Doppler shift in the pulse repetition frequency. Although increasing the repetition frequency increases the photon detection rate, it also decreases the unambiguous range. Several alternative pulsing modes have been proposed to break this tradeoff for SPL measurements of static scenes, e.g., using multiple repetition rates or random pulse trains to yield both high count rates and a long unambiguous range. In this work, we show that velocity can still be estimated despite the use of non-periodic range extension pulse patterns. We derive a general model for the photon acquisition process with a moving target and propose maximum likelihood estimators (MLEs) to recover the distance, velocity, and photon flux levels. Each range extension mode requires a tailored initialization scheme in order for our MLE solver to converge to the global optimum. We demonstrate through simulations and experiments that non-periodic SPL can simultaneously achieve long-range and high-accuracy distance and velocity estimates.

1. Introduction

There is a growing demand for precise 3D scene understanding from weak signals caused by low-power illumination sources or long-range measurements. Single-photon lidar (SPL) is a strong candidate technology due to its sensitivity to extremely low light levels [1–3] and demonstrated capability at distances over 1 km [4–8]. SPL systems are comprised of three main components: a laser that transmits pulses with durations of picoseconds; a detector such as a single-photon avalanche diode (SPAD) sensitive to individual photons; and a time tagger performing time-correlated single-photon counting (TCSPC), i.e., recording the times between pulse transmissions and photon detections. The pulsed laser illumination is typically repeated to accumulate more photons and improve the performance of statistical estimators.

Many interesting applications for SPL are complicated by motion of either the scene, the lidar system, or both, including vehicular lidar for assisted or autonomous driving [9, 10], airborne lidar for terrestrial scanning [11], or ground-based tracking of satellites in space [12]. Lidar that directly estimates velocity in addition to distance can provide key advantages in moving-object detection, odometry, point cloud registration, and trajectory planning [13–19]. For pulsed lidar, velocity estimation has typically required curve fitting to a sequence of distance measurements and computing a temporal derivative [20]; if lateral motion is also present, additional point cloud registration is required to ensure the same points are being tracked over time [21]. Instead, the authors recently introduced Doppler SPL [22], which takes advantage of the use of periodic pulse sequences to directly estimate a velocity from the detection times. The key observation is that a periodic sequence can be considered a type of amplitude modulation [23], where the pulse repetition frequency experiences a Doppler shift when the target moves towards or away from the lidar system. Estimating the change in pulse frequency between the transmission and received light enables radial velocity estimation [24].

One limiting factor of the Doppler SPL framework is its reliance on a periodic pulse sequence, which results in a tradeoff between accuracy and unambiguous range [25, 26]. The number of *signal photons*, i.e., light reaching the detector that originated at the laser, is often the main determinant of accuracy in SPL. It is thus helpful to increase the repetition frequency f_r, since all else being equal, transmitting more pulses will lead to more signal photon detections. However, the unambiguous range z_r is directly related to the repetition period $t_\mathrm{r} = 1/f_\mathrm{r}$ as $z_\mathrm{r} = ct_\mathrm{r}/2$, where c is the speed of light. To avoid ambiguity, systems employing periodic pulsing require that the largest distance in the scene z_max is less than the unambiguous range z_r [2]. Otherwise, if a target is farther than the unambiguous range, multiple pulses will be in flight before the reflected light is detected. Since the time of flight (TOF) of a photon is measured with respect to the most recent pulse emission time, the distance measurement would be aliased.

There are three main strategies employed to break the tradeoff and enable long unambiguous range with higher pulse rates. The *multiple repetition rate* (MRR) mode uses a sequence of different pulse repetition periods, and the unaliased distance is recovered either by the Chinese remainder theorem (CRT) applied to aliased distance estimates for each t_r [6, 27–29] or by the method of correlation and accumulation [26]. The *random* mode employs a random pulse pattern, and the distance is recovered by correlating the recorded timestamps with the pulse pattern [30–33]. The *interleaved* approach combines aspects of the previous two, generating a longer, nonperiodic pulse pattern by interleaving multiple periodic sequences with different pulse rates [34]. Importantly, all of the estimation methods have assumed static scenes, and performance degrades for moving targets [32].

In this paper, we develop a framework to not only make these range extension modes compatible with moving targets, but also to enable radial velocity estimation. We refer to all three strategies as *non-periodic* because they do not use a single periodic pulse pattern. We observe that the velocity of a moving target causes a time warping of the Poisson process intensity function describing the received light, a phenomenon that is not limited to periodic pulsing. We thus derive for each range extension mode the intensity function that governs the distribution of photon detection times in the presence of target motion with a constant radial velocity. We propose a maximum likelihood estimator (MLE) for joint estimation of the distance and velocity, as well as the signal and background photon flux. The MLE is then implemented for each mode with a tailored initialization approach followed by further numerical optimization. One important component of the initialization schemes for the MRR and interleaved modes is a new CRT algorithm robust to non-integer repetition periods and pulse width uncertainty. We validate our framework via simulations that show statistically efficient estimators achieving the Cramér–Rao Bound (CRB) under many conditions. We further perform laboratory experiments highlighting our ability to break the tradeoff between accuracy and unambiguous range. Despite significant background count rates, we show experimental results with sub-centimeter distance accuracy and velocity root mean square error (RMSE) less than 5 cm/s.

2. Measurement Models

Under conventional SPL acquisition, the laser emits n_r short pulses with a repetition period t_r over a total acquisition time $t_\mathrm{a} = n_\mathrm{r}t_\mathrm{r}$. Each pulse travels at the speed of light c, reflects off the target, and some scattered light returns to the lidar system. The time tagger records time stamps of individual photon detections. So long as detector dead times can be neglected [35], the time stamps of a single-photon detector are described as a Poisson process [36]. For periodic pulsing and static targets, the intensity function for the Poisson process observed for SPL has been well established [2].

We recently introduced Doppler SPL for moving targets, showing how the intensity function under periodic illumination experiences a Doppler shift in the repetition frequency [22]. Suppose the target is at initial distance z at time $t = 0$, moving with constant radial velocity v, which is

reasonable over short acquisition times. Let $h(t)$ denote the temporal pulse profile, normalized such that $\int_{-\infty}^{\infty} h(t)\,\mathrm{d}t = 1$. The time stamps then follow an inhomogeneous Poisson process with the intensity function

$$\lambda(t) = S\left[\sum_{n=0}^{n_\mathrm{r}-1} h\left(t - \frac{c}{c-v}\tau - n\frac{c+v}{c-v}t_\mathrm{r}\right)\right] + b, \quad t \in [0, t_\mathrm{a}). \tag{1}$$

The signal flux S is the mean number of detected laser photons per pulse, absorbing the effects of laser power, reflectivity, quantum efficiency, distance fall-off, etc. The background intensity b is the mean number of background detections from ambient light and dark counts per unit time. The initial delay $\tau = 2z/c$ denotes the TOF if the target were static and remained at distance z for the entire acquisition.

For static scenes ($v = 0$), the intensity function is a version of the illumination pattern that is time-shifted by the delay, scaled in amplitude by the reflectivity, and offset by the background. Histograms of *relative* detection times (i.e., modulo t_r) are sufficient for estimating the distance from the time shift. As derived in [22, Supp. 1], targets moving at non-zero velocity v cause an additional *time warping* of the pulse pattern, changing from the initial TOF as $\tau' = \tau c/(c-v)$ and repetition period $t_\mathrm{r}' = t_\mathrm{r}(c+v)/(c-v)$. Since this warping depends on the unknown velocity v, Doppler SPL requires the *absolute* detection times $(T_i)_{i=1}^N$ elapsed since the beginning of the acquisition at $t = 0$, where N is the number of detections. The time stamps $(T_i)_{i=1}^N$ enable joint estimation of all scene parameters—S, b, τ, and v—through frequency analysis and maximum likelihood estimation.

In this paper, we recognize that the time warping is not limited to periodic pulse trains. Instead, we can consider an arbitrary sequence of P pulses at times $\Pi := \{t_p : 0 \le t_p < \bar{t}_\mathrm{r}\}_{p=1}^P$, where each t_p denotes a pulsing time and the maximum unambiguous delay is $\bar{t}_\mathrm{r}$. Figure 1 shows an example of how a non-periodic pulse train can be generated by triggering a laser with an arbitrary waveform generator (AWG). Then a generic Poisson process intensity function is

$$\lambda(t) = S\left[\sum_{\ell=0}^{L-1}\sum_{p=1}^{P} a_p h\left(t - \frac{c}{c-v}\tau - \frac{c+v}{c-v}(t_p + \bar{t}_\mathrm{r}\ell)\right)\right] + b, \quad t \in [0, t_\mathrm{a}). \tag{2}$$

Due to the limitations of our experimental apparatus (that we will describe in Section 5), we include a scalar $a_p \in \mathbb{R}_+$ modeling uneven power in each laser pulse, normalized such that $\sum_{p=1}^P a_p = P$, and allow for a sequence with unambiguous delay $\bar{t}_\mathrm{r}$ to be repeated L times. The periodic intensity in (1) is a special case of the arbitrary intensity in (2) with $\Pi = \{0\}$, $a_p = 1$, $\bar{t}_\mathrm{r} = t_\mathrm{r}$, and $L = n_\mathrm{r}$.

2.1. Range Extension Modes

Under periodic illumination, the unambiguous range $z_\mathrm{r} = ct_\mathrm{r}/2$ is set by the repetition period. When the distance z exceeds z_r, the measurement model (1) remains valid. However, the inability to determine which laser pulse was the source of each detected signal photon leads to range ambiguity. To avoid range ambiguity, one can use periodic pulsing with a sufficiently long repetition period, which we call the *baseline* mode. Alternatively, if a higher pulse rate causes the maximum scene distance to be greater than the unambiguous range, we call the periodic SPL operation *ambiguous*. To break this tradeoff, we now detail the three main approaches to unambiguous range extension as examples of arbitrary pulse patterns.

2.1.1. Multiple-Repetition-Rate Mode

The MRR mode collects multiple sets of SPL measurements, each with a different repetition period [6, 27–29]. A modulo or relative distance is estimated for each t_r, and the absolute

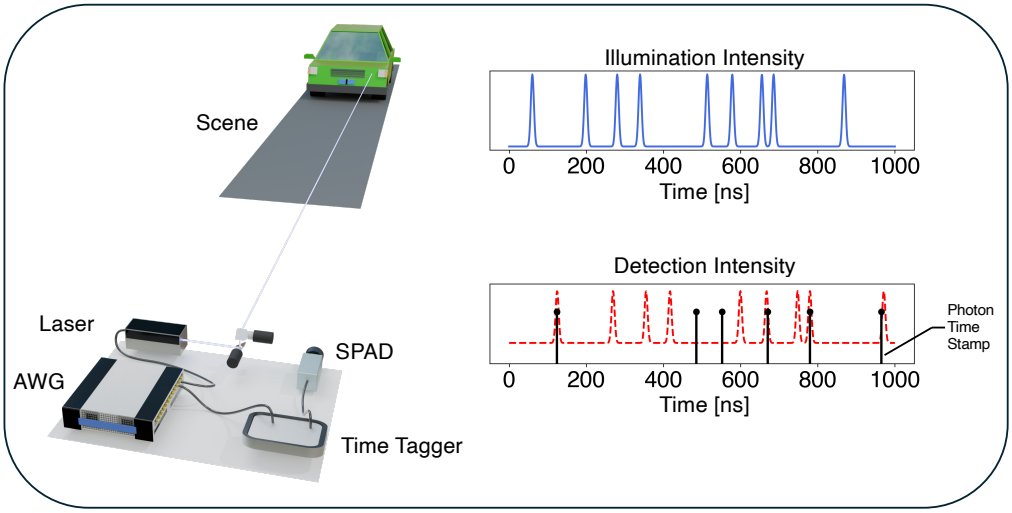

Fig. 1. Acquisition overview. The lidar system consists of a laser, arbitrary waveform generator (AWG), time tagger, and single-photon avalanche diode (SPAD). The laser transmits a non-periodic sequence of pulses toward the scene, yielding a long unambiguous range. The SPAD detects photons both reflected from the target and due to ambient light. The detection intensity relates to the illumination intensity by a time delay (distance), time warping (velocity), amplitude scaling (reflectivity), and an amplitude offset (background).

distance is recovered from the relative distances using the robust Chinese remainder theorem (CRT) [37–41] or the correlation-and-accumulation approach [26]. Suppose the lidar system acquires K sequential measurements with repetition periods $(t_\mathrm{r}^k)_{k=1}^K$, each using n_r pulses, where the beginning of each subacquisition is $(t_0^k)_{k=1}^K$. Assuming the repetition periods are integers, then the unambiguous delay $\bar{t}_\mathrm{r}$ is extended to the least common multiple (LCM) of the repetition periods [6, 27–29]. The repetition periods in real experiments deviate from integer numbers, so based on the uncertainty of our system, we approximate the unambiguous TOF by rounding the repetition periods to the nearest integer number of nanoseconds as

$$\bar{t}_\mathrm{r} \approx \mathrm{LCM}([t_\mathrm{r}^1], \ldots, [t_\mathrm{r}^K]), \tag{3}$$

where the bracket operation $[\cdot]$ indicates rounding. Because each subacquisition uses periodic pulses, the detection times in each subacquisition follow the model described in (1). However, the initial TOF for the kth subacquisition shifts to $\tau^k = \tau + (2v/c)t_0^k$ due to the target's movement. If each subacquisition follows immediately after the previous one, then $t_0^1 = 0$ and $t_0^k = t_0^{k-1} + n_\mathrm{r} t_\mathrm{r}^{k-1}$ for any $k > 1$.

2.1.2. Random Mode

The random mode employs a random pulse pattern Π, generated by a pseudo-random number generator [30, 31] or a chaos laser [42, 43]. The distance is recovered by correlating the recorded timestamps with the pulse pattern, and the pulse pattern length determines the unambiguous range. A further benefit of randomized pulse sequences is that they have increased robustness to interference from other lidar systems [44].

We generate random pulse patterns computationally using a pseudo-random number generator, following [30, 31]. Our hardware implementation requires a minimum separation of 9 ns between pulses, so we used Poisson disk sampling [45] for experiments. Although there is no need

to impose a strict minimum pulse separation in simulation, uniformly sampling pulsing times over $[0, \bar{t}_\mathrm{r})$ often yields clusters with separations smaller than the pulse width [46]. Instead, we simulate pulse patterns by partitioning the pattern duration $\bar{t}_\mathrm{r}$ into P equal non-overlapping sub-intervals and sampling each pulsing time uniformly within each sub-interval. Empirically, this approach yields fewer overlapping pulses while being faster than Poisson disk sampling.

We observe uneven pulse power in experiments, which we include in the detection model (2). Although individual pulses within a pattern have different amplitudes, we find that the amplitude pattern remains approximately constant over time. In experiments, the amplitudes $a_1, \ldots, a_P$ are known through a calibration procedure, described in Section 5.1.3. In simulations, we set $a_p = 1$ for all $p \in \{1, \ldots, P\}$.

2.1.3. Interleaved Mode

Rather than using a completely arbitrary (i.e., random) pulse pattern, the interleaved mode transmits a structured pattern combining multiple periodic sequences with different pulse rates simultaneously [34]. Given repetition periods $(t_\mathrm{r}^k)_{k=1}^K$, the pulse pattern is $\Pi = \bigcup_{k=1}^K \{nt_\mathrm{r}^k \mid n \in \mathbb{N} \text{ and } nt_\mathrm{r}^k < \bar{t}_\mathrm{r}\}$, where the pattern duration $\bar{t}_\mathrm{r}$ is the least common multiple of the repetition periods $(t_\mathrm{r}^k)_{k=1}^K$. Like the MRR mode, the periodic sub-sequences enable velocity estimation through frequency analysis, as we discuss in Section 3.2.4. Similar to the random mode, the pulse amplitudes fluctuate, but exhibit a consistent pattern. Again, we calibrate the pulse amplitudes in experiments and set $a_p = 1$ for all $p \in \{1, \ldots, P\}$ in simulations.

3. Estimation Methods

3.1. Maximum Likelihood Estimation

Given any parametric Poisson process intensity function $\lambda(t)$, the log-likelihood is straightforward to derive [36]. It is thus natural to consider maximum likelihood estimation of the parameters. However, even for the simple case of periodic SPL for static targets, the negative log-likelihood is highly non-convex [2, 47]. There are many approaches for solving such non-convex optimization problems. We follow the example of periodic Doppler SPL and adopt a two-step implementation of the MLE. First, we determine a reasonably efficient initialization procedure that avoids a brute-force grid search over all four parameters. Second, we refine the initial estimate using numerical optimization methods, assuming the negative log-likelihood is locally convex.

The numerical optimization procedure is effectively the same for all modes. Given the initialization, we alternate between optimizing (S, b) and (τ, v) using the L-BFGS-B [48] algorithm in each step, as we observe that the alternating scheme improves convergence of the MLE solver to the global optimum. We also bound τ by $[\widehat{\tau} - 10\sigma, \widehat{\tau} + 10\sigma]$, where σ is the pulse width, to ensure that the estimate from numerical optimization is close to the initialization $\widehat{\tau}$. The log-likelihood is highly non-concave in τ, so large deviations in τ from the initialization often result in unreliable solutions.

The initialization procedure differs between modes, depending primarily on whether there is a periodic component to the pulse pattern. We now elaborate on the estimation procedure for each mode.

3.2. Solver Initialization

3.2.1. Baseline Mode

Following [22], the log-likelihood for a periodic Doppler SPL measurement $(T_i)_{i=1}^N$ is

$$\mathcal{L}(S, b, \tau, v) = -n_\mathrm{r} S - t_\mathrm{a} b + \sum_{i=1}^{N} \lambda(T_i), \tag{4}$$

where $\lambda(t)$ is the baseline intensity function (1). Because the target motion induces a Doppler shift in the repetition frequency, we first estimate the received frequency f_{r}' via Fourier analysis tailored for single-photon detections [49]. We constrain the search space for the Doppler shift by assuming the maximum achievable speed is $|v_{\max}| = 150\,\mathrm{m/s}$; higher velocities can be measured in principle, so long as $|v| \ll c$. Given the estimated Doppler shift $\widehat{f_{\mathrm{r}}'}$, we compute the initial velocity as $\widehat{v} = c(f_{\mathrm{r}} - \widehat{f_{\mathrm{r}}'})/(f_{\mathrm{r}} + \widehat{f_{\mathrm{r}}'})$. We next un-warp the effect of the estimated velocity on the detection times and apply a *censoring* procedure, which cross-correlates the detection times with a short rectangular window to yield initial estimates of S, b, and τ [47, 50].

3.2.2. Multiple-Repetition-Rate Mode

The log-likelihood for the MRR mode is

$$\mathcal{L}(S, b, \tau, v) = \sum_{k=1}^{K} \mathcal{L}^k \left(S, b, \tau + (2v/c)t_0^k, v\right), \tag{5}$$

where each component $\mathcal{L}^k$ is the log-likelihood for Doppler SPL (4) with repetition period t_{r}^k. We first separately compute preliminary estimates $\left\{\left(\widehat{S}^k, \widehat{b}^k, \widehat{\tau}^k, \widehat{v}^k\right)\right\}_{k=1}^{K}$ for each t_{r}^k using the initialization scheme for the baseline Doppler SPL mode described in Section 3.2.1. We then combine these preliminary estimates into the MLE initializations. The initial S, b, and v for MRR are determined by averaging individual estimates of the corresponding scene parameters across repetition periods, e.g., $\widehat{v} = (1/K)\sum_{k=1}^{K} \widehat{v}^k$.

For the delay, we assume each preliminary estimate $\widehat{\tau}^k$ approximates a modulo measurement of the true time of flight. We note that for a moving scene, the time of flight changes due to the different start times of each sequence. We thus apply the CRT to the motion-corrected TOF estimates to disambiguate the range:

$$\widehat{\tau} = \mathrm{CRT}\left[(\widehat{\tau}^1 - 2\widehat{v}t_0^1/c, t_{\mathrm{r}}^1), \ldots, (\widehat{\tau}^K - 2\widehat{v}t_0^K/c, t_{\mathrm{r}}^K)\right], \tag{6}$$

where $\mathrm{CRT}[(\tau^1, t_{\mathrm{r}}^1), \ldots, (\tau^K, t_{\mathrm{r}}^K)]$ denotes the robust CRT solution to the system of equations we will describe in (14). Further details of the CRT and our algorithm are discussed in Section 3.4.

3.2.3. Random Mode

The log-likelihood for the random mode is

$$\mathcal{L}(S, b, \tau, v) = -LPS - t_{\mathrm{a}}b + \sum_{i=1}^{N} \lambda(T_i), \tag{7}$$

where $\lambda(t)$ is the intensity function of the particular random pulse train (2). One challenge of the random mode is that there are no periodic components to the pulse pattern, so we cannot take advantage of the Fourier initialization used for Doppler SPL. Rather than perform a multi-dimensional grid search over τ and v, we initialize the velocity parameter to $0\,\mathrm{m/s}$.

For the other parameters, S, b and τ, we initialize estimates using the cross-correlation approach typically used for random pulse train measurements [31], combined with the censoring procedure [50]. First, we estimate τ using matched filtering:

$$\widehat{\tau} = \underset{\tau \in [0, \bar{t}_{\mathrm{r}})}{\arg\max} \sum_{i=1}^{N} \log\left(\widetilde{S}\left[\sum_{p=1}^{P} a_p h(T_i \bmod \bar{t}_{\mathrm{r}} - \tau - t_p)\right] + \widetilde{b}\right) \tag{8}$$

which is the maximum likelihood estimate of τ if $v = 0$ and the flux parameters $\widetilde{S}$ and $\widetilde{b}$ are known. However, even inaccurate flux estimates can yield good τ estimates, so long as $\widetilde{S}$ and $\widetilde{b}$

are non-zero [22], so we arbitrarily set $\widetilde{S} = 1$ and $\widetilde{b} = 10^{-3}$ in simulation and experiments. Given $\widehat{\tau}$, we then estimate initial S and b by counting the number of detections, shifted back by $\widehat{\tau}$, in the time windows around the pulsing times:

$$\widehat{S} = \frac{1}{LP} \sum_{i=1}^{N} \mathbf{1}_{\mathcal{W}}(T_i \bmod \bar{t}_{\mathrm{r}} - \widehat{\tau}), \tag{9}$$

where $\mathcal{W} := \bigcup_{p=1}^{P} [t_p - t_{\mathrm{win}}/2, t_p + t_{\mathrm{win}}/2]$ is the union of time windows around the pulsing times. We set the window size $t_{\mathrm{win}} = 4\sigma + 4v_{\max}t_{\mathrm{a}}/c$, where $v_{\max} = 150\,\mathrm{m/s}$ is a maximum speed anticipated, to account for potential motion blur due to non-zero velocity. Then the background intensity is initialized by subtracting the estimated number of signal detections from the total number of detections:

$$\widehat{b} = (N - LP\widehat{S}) \Big/ t_{\mathrm{a}} \,. \tag{10}$$

We observe that the log-likelihood is often sufficiently smooth in v, allowing gradient-based methods to reach global optimal points, despite v being initialized to $0\,\mathrm{m/s}$. However, as we will demonstrate in Section 4.3, the MLE with this initialization can fail in settings with high-speeds and/or long acquisition times.

3.2.4. Interleaved Mode

The log-likelihood for interleaved mode takes the same general form (7) as for an arbitrary pulse sequence. However, as described in Section 2.1.3, the pulsing pattern Π is comprised of periodic sub-sequences of pulses, and we can exploit the structure of this pattern. First, we extract the likely signal detections for each t_{r}^k by searching for the largest cluster of relative detection times $(T_i \bmod t_{\mathrm{r}}^k)_{i=1}^N$ in a window of size $4\sigma + 4v_{\max}t_{\mathrm{a}}/c$, similar to the censoring estimator [50]. Using those subsets of detections, we compute preliminary estimates $\left\{\left(\widehat{S}^k, \widehat{\tau}^k, \widehat{v}^k\right)\right\}_{k=1}^K$ for each t_{r}^k as in the baseline and MRR approaches. As in MRR mode, we average the signal flux and velocity estimates over all K repetition periods to obtain initial estimates $\widehat{S}$ and $\widehat{v}$. We apply the CRT to TOF estimates to disambiguate the range:

$$\widehat{\tau} = \mathrm{CRT}\left[(\widehat{\tau}^1, t_{\mathrm{r}}^1), \ldots, (\widehat{\tau}^K, t_{\mathrm{r}}^K)\right]. \tag{11}$$

The background intensity is initialized using the same estimator as the random mode (10).

In experiments, the true pulsing times slightly deviate from the nominal interleaved patterns given by the superposition of the repetition periods. While the nominal t_{r}^ks are sufficient to initialize using the procedure described above, for gradient-based maximization of the log-likelihood (7), we use the calibrated times for each pulse, rather than assuming exact periodicity holds.

3.3. Cramér–Rao Bound

The CRB lower bounds the covariance matrix of any unbiased estimator by the inverse of Fisher information matrix (FIM) [51]. We present here an overview of the CRB derivation for an arbitrary pulse train, which generalizes all other range extension modes. A full derivation of the CRB for the S, b, τ, and v parameters is contained in Supplement 1. Simulation results in Section 4 show that our MLE implementation can often achieve the CRB.

For simplicity, we assume $a_p = 1$ and $L = 1$. We also assume the pulse shape $h(t)$ is zero outside of the support $[-\delta, \delta]$ for some small δ and that the pulses do not overlap. The 4×4

FIM for all parameters is then approximately given as

$$\mathcal{I} \approx \begin{bmatrix} \mathcal{I}_S & \mathcal{I}_{Sb} & 0 & 0 \\ \mathcal{I}_{Sb} & \mathcal{I}_b & 0 & 0 \\ 0 & 0 & \mathcal{I}_\tau & \mathcal{I}_{\tau v} \\ 0 & 0 & \mathcal{I}_{\tau v} & \mathcal{I}_v \end{bmatrix}, \tag{12}$$

where

$$\mathcal{I}_S = P \int_{-\delta}^{\delta} \frac{h(t)^2}{Sh(t)+b} \mathrm{d}t \tag{13a}$$

$$\mathcal{I}_b = P \int_{-\delta}^{\delta} \frac{1}{Sh(t)+b} \mathrm{d}t + \frac{t_\mathrm{a} - P(2\delta)}{b} \tag{13b}$$

$$\mathcal{I}_{Sb} = P \int_{-\delta}^{\delta} \frac{h(t)}{Sh(t)+b} \mathrm{d}t \tag{13c}$$

$$\mathcal{I}_\tau = S^2 \left(1 + \frac{v}{c}\right)^2 P \int_{-\delta}^{\delta} \frac{h'(t)^2}{Sh(t)+b} \mathrm{d}t \tag{13d}$$

$$\mathcal{I}_v = S^2 \sum_{p=1}^{P} \left[\frac{2}{c}\left(\frac{\tau}{2} + t_p\right)\right]^2 \left(\int_{-\delta}^{\delta} \frac{h'(t)^2}{Sh(t)+b} \mathrm{d}t\right) \tag{13e}$$

$$\mathcal{I}_{\tau v} = S^2 \left(1 + \frac{v}{c}\right) \sum_{p=1}^{P} \left[\frac{2}{c}\left(\frac{\tau}{2} + t_p\right)\right] \left(\int_{-\delta}^{\delta} \frac{h'(t)^2}{Sh(t)+b} \mathrm{d}t\right). \tag{13f}$$

Any unbiased estimate $(\widehat{S}, \widehat{b}, \widehat{\tau}, \widehat{v})$ has its root mean square errors lower bounded by the diagonal entries of the inverse FIM, e.g., $\sqrt{\mathbb{E}[(\widehat{S}-S)^2]} \geq \sqrt{(\mathcal{I}^{-1})_{11}}$.

3.4. Robust Chinese Remainder Theorem

3.4.1. Problem Formulation

For periodic SPL, when $\tau > t_\mathrm{r}$, we measure only the relative TOF $\tau \bmod t_\mathrm{r}$. Without any additional information, the integer number of repetition periods between transmit and receive is unknown and causes the range ambiguity. The MRR and interleaved modes resolve this ambiguity by using sub-sequences with different repetition periods. For each mode, we can use the repetition periods $(t_\mathrm{r}^k)_{k=1}^K$ to formulate a system of congruences for the relative TOF $\tau^k := \tau \bmod t_\mathrm{r}^k$ from each repetition period t_r^k as:

$$\tau = n_k t_\mathrm{r}^k + \tau^k, \quad k \in \{1, \ldots, K\}, \tag{14}$$

where $n_k := \lfloor \tau/\tau^k \rfloor$ applies the floor operation. According to the CRT [52], if $t_\mathrm{r}^1, \ldots, t_\mathrm{r}^K$ are integers having a common pairwise greatest common divisor (GCD) and $\tau \in [0, \bar{t}_\mathrm{r})$ where $\bar{t}_\mathrm{r} = \mathrm{LCM}(t_\mathrm{r}^1, \ldots, t_\mathrm{r}^K)$, then there is a unique solution τ to this system of congruences.

When the estimated relative TOFs $(\widehat{\tau}^k)_{k=1}^K$ are not exact and possibly non-integers, robust CRT algorithms can still recover τ under some assumptions about the estimation errors [37–41]. Specifically, Wang and Xia [37] showed that the integers $n_1, \ldots n_K$ can be recovered exactly if $|\widehat{\tau}^k - \tau^k| < M/4$ for all k, where M is the common pairwise GCD of the moduli $t_\mathrm{r}^1, \ldots, t_\mathrm{r}^K$. Therefore, the MRR and interleaved modes effectively extend the unambiguous range to $\bar{t}_\mathrm{r}$. In our simulations where the repetition periods can be set to integers precisely, we use Wang et al.'s maximum likelihood approach for robust CRT with a Gaussian noise model [40] to estimate τ

from relative TOF estimates $(\widehat{\tau}^k)_{k=1}^K$. In this case, the Gaussian noise translates to uncertainty due to the instrument response function, (i.e., pulse width, timing jitter, SPAD response time, etc.) of the lidar system.

3.4.2. Real Number Moduli

In real SPL systems, we often cannot set repetition periods to be exact integers, e.g., 100 ns, but rather real numbers, e.g., 100.087 . . . ns, due to hardware limitations. While many robust CRT algorithms in the literature allow for a real-number solution τ and remainders $(\widehat{\tau}^k)_{k=1}^K$, they all require that the moduli $(t_\mathrm{r}^k)_{k=1}^K$ are integers. Fitting the repetition periods to a time grid such that they are approximately integers on the grid is one solution to make the problem compatible to existing robust CRT algorithms. However, as we show in Supplement 1, we find solutions from these algorithms to be sensitive to small fractional deviations from integer moduli.

Instead, when $(t_\mathrm{r}^k)_{k=1}^K$ are positive real numbers that are not necessarily integers, we propose to solve the robust CRT problem through a mixed-integer program (MIP) enforcing consistency between the relative TOF estimates. The ground truth relative TOFs $(\tau^k)_{k=1}^K$ satisfy the following system of linear equations:

$$\begin{aligned} \tau &= n_1 t_\mathrm{r}^1 + \tau^1, \\ &\vdots \\ \tau &= n_K t_\mathrm{r}^K + \tau^K. \end{aligned}$$

The system has K equations and contains $K + 1$ unknowns. We first solve for $n_1, \ldots, n_K$ by leveraging the fact that they must be integers. Taking the difference between any pair of equations j and k, where $j \neq k$ and $j, k \in \{1, \ldots, K\}$, results in

$$\tau^j - \tau^k = -n_j t_\mathrm{r}^j + n_k t_\mathrm{r}^k. \tag{15}$$

We aim to find $n_1, \ldots, n_K$ satisfying these $K(K-1)/2$ equations. Since we only have estimates of the relative TOFs $(\widehat{\tau}^k)_{k=1}^K$, we propose to minimize the sum of squared residues, resulting in an MIP:

$$\min_{\substack{n_1,\ldots,n_K \\ n_k \in \{0,1,\ldots,n_k^{\max}\} \\ \forall k \in \{1,\ldots,K\}}} \left\| \begin{bmatrix} \vdots \\ \tau_j - \tau_k \\ \vdots \end{bmatrix} - \begin{bmatrix} & & \vdots & & \\ 0 & \ldots & -t_\mathrm{r}^j & \ldots & t_\mathrm{r}^k & \ldots & 0 \\ & & \vdots & & \end{bmatrix} \begin{bmatrix} n_1 \\ \vdots \\ n_K \end{bmatrix} \right\|_2^2, \tag{16}$$

where $n_k^{\max} = \lfloor \bar{t}_\mathrm{r} / t_\mathrm{r}^k \rfloor$ with $\bar{t}_\mathrm{r} = \mathrm{LCM}([t_\mathrm{r}^1], \ldots [t_\mathrm{r}^K])$ and $[\cdot]$ denotes rounding to the nearest integer. Each matrix row corresponds to the difference between the j^{th} and the k^{th} equations, so the j^{th} entry is $-t_\mathrm{r}^j$, the k^{th} entry is t_r^k, and other entries are zeros. If we consider an example where $K = 2$, as in our simulations and experiments, then the MIP in (16) simplifies to

$$\min_{\substack{n_1,n_2 \\ n_k \in \{0,1,\ldots,n_k^{\max}\},\ k \in \{1,2\}}} (\widehat{\tau}^1 + n_1 t_\mathrm{r}^1 - \widehat{\tau}^2 - n_2 t_\mathrm{r}^2)^2. \tag{17}$$

We solve the MIP (16) using the SCIP solver [53] through CVXPY [54, 55]. Given a solution $\widehat{n}_1, \ldots, \widehat{n}_K$ from the program (16), our unambiguous TOF estimate is

$$\widehat{\tau} = \frac{1}{K} \sum_{k=1}^K \widehat{n}_k t_\mathrm{r}^k + \widehat{\tau}^k. \tag{18}$$

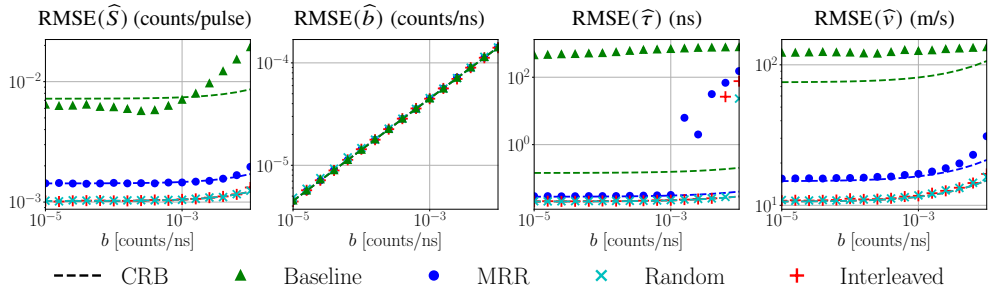

Fig. 2. SPL simulation results with the acquisition time fixed to 0.5 ms. The dashed lines indicate the CRBs and the markers indicate the simulation results for each mode. The range extension modes benefit from detecting many more photons than the baseline mode.

4. Simulation Results

We compare the performance of different SPL range extension modes through simulations. We assume there is no pulse power variation, so $a_p = 1$ for any p in the intensity function (2) for the random and interleaved modes. Unless otherwise specified, the following settings are used: $S = 0.01$ counts/pulse, $\tau = 1077$ ns, $v = 30$ m/s. The pulse profile $h(t)$ is assumed to be a Gaussian with standard deviation $\sigma = 0.1$ ns. For MRR and interleaved modes, we use two repetition periods: $t_{\mathrm{r}}^1 = 100$ ns and $t_{\mathrm{r}}^2 = 104$ ns, resulting in the unambiguous delay of $\bar{t}_{\mathrm{r}} = 2600$ ns. We also include the periodic Doppler SPL as the baseline with $t_{\mathrm{r}} = 2600$ ns to match the unambiguous range. The number of pulses and the acquisition time in the random mode is always set to match the number of pulses in the interleaved mode. We report the root mean square error (RMSE) for Monte Carlo experiments performed with 10 000 trials per combination of parameters.

4.1. Robustness to Background for Fixed Acquisition Time

We first test the robustness to background while keeping the acquisition time fixed. We vary $b \in [10^{-5}, 10^{-2}]$ and set $t_{\mathrm{a}} = 0.5$ ms, resulting in 5 to 5000 mean background detections. The number of laser pulses varies by mode, with 192 pulses of the baseline mode, 4900 pulses for the MRR mode, and 9614 pulses for the random and interleaved modes.

According to (13), all non-zero elements of the Fisher information matrix depend on the number of pulses P. The results in Fig. 2 correspondingly show three tiers of performance as a function of the number of transmitted pulses. In the bottom tier is the baseline mode, which produces too few signal photons for reliable estimation of anything other than the background. The baseline S estimate often beats the CRB, but that is largely due to the bias of having an estimator constrained to be non-negative [47]. In the middle tier is the MRR mode, which has substantially better accuracy than the baseline. However, there is still a noticeable gap when compared to the top-tier random and interleaved modes, because those modes can transmit more pulses over the same acquisition time. For low b, all three range-extension modes are at or near the CRB for all four parameters.

Among the middle and top tiers, an important distinction is the robustness to very high background, noticeable in the τ estimation. While the MRR and interleaved modes separate the signal photons by repetition period and then apply the CRT, the random mode uses cross-correlation with all time stamps simultaneously, which may explain the somewhat better robustness.

We further present in Supplement 1 the run time performance for the baseline and the three

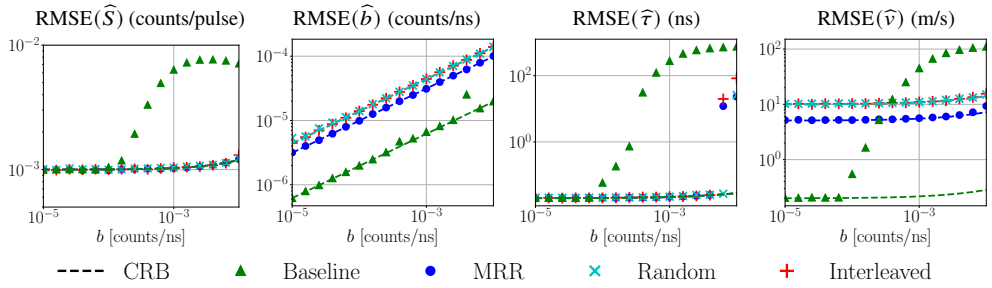

Fig. 3. SPL simulation results with the mean number of signal detections fixed to 100. The dashed lines indicate the CRBs and the markers indicate the simulation results for each mode. The CRB decreases for b and v as the acquisition time increases.

range extension modes. The run times for all modes are similar, increasing with the number of photon detections.

4.2. Robustness to Background for Fixed Number of Pulses

We consider an alternative experiment where the number of pulses is fixed while we vary the background $b \in [10^{-5}, 10^{-2}]$. The acquisition time is varied to ensure 10 000 pulses emitted, i.e., 100 mean signal photons, resulting in the acquisition times 26 ms for the baseline mode, 1.02 ms for the MRR mode, and 0.52 ms for the random and interleaved modes.

We observe in Fig. 3 an interesting bifurcation in the estimator performance. The accuracy of the $\widehat{S}$ and $\widehat{\tau}$ estimators does not directly depend on the acquisition time. All modes have the same number of pulses and thus the same CRB, and so long as b is low enough, all modes can achieve the same accuracy. However, the baseline mode is more sensitive to the background flux due to the longer acquisition time, which leads to a higher number of background detections.

On the other hand, the $\widehat{b}$ and $\widehat{v}$ accuracy depends directly on t_a. The background rate is easier to estimate from more observations over a longer acquisition time. Similarly, the Doppler shift is easier to detect over a long acquisition, which allows for a larger displacement of the target [22]. As a result, the CRB for the baseline mode is substantially lower for both parameters. Actually achieving the CRB is still a challenge, however. At high background rates, the range extension modes meet their respective CRBs and outperform the baseline.

4.3. Random Mode Limitation at High Speed

Although the random mode exhibits excellent robustness to background, our implementation of the MLE has a significant limitation in that there is no efficient initialization procedure available for the velocity parameter. We demonstrate this limitation by simulating SPL measurements for target velocities varying over 0 to 100 m/s. The other parameters are fixed to $S = 1\times10^{-4}$ counts/pulse, $b = 1\times10^{-5}$ counts/ns, $\tau = 1077$ ns, and $t_\mathrm{a} = 10.0$ ms. Fig. 4 shows the results of parameter estimation comparing the random and interleaved modes. While the interleaved mode achieves the CRB for all tested velocities, the random mode accuracy degrades when the velocity exceeds 30 m/s. Since the random-mode likelihood (7) is extremely non-concave and we initialize $\widehat{v} = 0$, the numerical optimization to find the MLE converges to the correct result only if the initial τ estimate is within about one pulse width from the true value. As the target velocity increases, the motion is more likely to cause errors in the τ initialization.

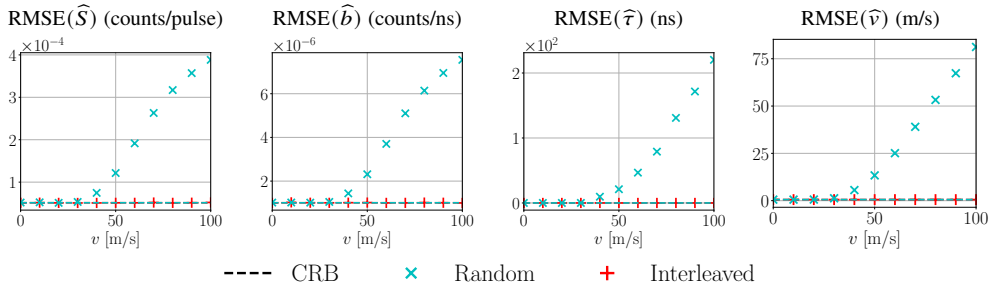

Fig. 4. Performance evaluation versus target velocity. Since the random mode initializes the velocity at $v = 0\,\mathrm{m/s}$, estimation accuracy degrades at high velocity.

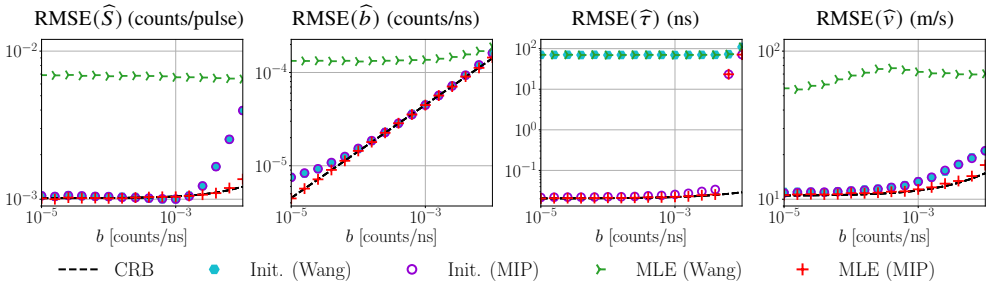

Fig. 5. CRT algorithm performance for non-integer repetition periods. The CRT method of Wang et al. [40] fails when the moduli are too far from integer values, whereas our mixed-integer program (MIP) formulation of the CRT is unaffected. Initializing the MLE solver with the Wang method results in worse performance for all parameters, whereas initializing the MLE solver with our MIP-based method results in parameter estimates achieving the CRB.

4.4. CRT Algorithm Performance for Non-integer Moduli

In the previous simulations, we assumed integer-valued repetition periods for the MRR and interleaved modes. To justify our introduction of a new MIP-based algorithm for the CRT, we investigated the impact of the choice of CRT algorithm on the parameter estimates. We repeated the experiments from Section 4.1 for the interleaved mode, except the repetition periods were chosen to be $t_\mathrm{r}^1 = 99.9\,\mathrm{ns}$ and $t_\mathrm{r}^2 = 104.1\,\mathrm{ns}$. The approximate unambiguous delay is still $\bar{t}_\mathrm{r} = 2600\,\mathrm{ns}$ based on the LCM of the rounded values. To avoid overlaps, we eliminated any pulses for which the previous pulse time was less than 10σ (i.e., 1 ns). We then tested the estimation performance using the method of Wang et al. [40], which is the maximum likelihood CRT estimator for integer moduli, and our MIP-based formulation of the CRT.

Figure 5 shows the performance of our initialization scheme, using either the Wang or MIP method for τ initialization, as well as the results from refinement with our MLE solver. The initial S, b, and v estimates are identical, since those parameters do not depend on the CRT output. However, our MIP-based CRT algorithm outperforms the Wang method by more than three orders of magnitude for the delay parameter τ. The method of Wang et al. does not correctly solve the system of congruences in (14) when the repetition periods deviate from integer values on the order of the pulse width σ or greater. Since the initial τ estimate from Wang et al. is too far from the global optimum, refinement by our MLE solver cannot improve the τ estimate; moreover, the S, b, and v estimates are degraded as the solver converges to an incorrect local minimum.

On the other hand, our MIP-based CRT algorithm is agnostic to whether the repetition periods are integers or not, and it provides a sufficiently close initialization for the MLE refinement to achieve the CRB for all parameters. This simulation justifies the need for a CRT algorithm that can handle non-integer moduli, which is especially common for experimental data.

5. Experiments

We validate the real-world feasibility of our methods by building upon the experimental setup initially used for Doppler SPL [22]. We first describe the physical setup, pulse patterns, and calibration procedures used in our experiments.

5.1. Experimental Setup

5.1.1. Lidar System

Our SPL system is configured as depicted in Fig. 1, and a photograph of the setup is contained in Supplement 1. The core equipment consists of a gain-switched diode laser (PicoQuant Prima, 443 nm), silicon SPAD (ID Quantique ID100-50-ULN), and time tagger (Time Tagger Ultra, Swabian Instruments). To achieve non-periodic pulse patterns, we externally trigger the laser with an arbitrary waveform generator (AWG, Keysight M8195A) programmed to produce trigger pulses with duration 2 ns and amplitude −800 mV. A separate pulse is sent from the AWG to the time tagger at the start of each pulse pattern for synchronization.

5.1.2. Pulse Patterns

The AWG has a sample rate of 60 GHz and a memory capacity of 256 000 samples, which limits the pattern length to around 4 µs. To achieve longer acquisition times, we thus repeat the pulse pattern, i.e., ($L > 1$). Although the Prima laser can achieve periodic pulse rates upwards of 200 MHz ($t_\mathrm{r} = 5$ ns), we observed in initial tests that closely-spaced, low-intensity, non-periodic pulse sequences caused the laser pulse power to vary between pulses by a factor of 10 or more. We found this variability could be reduced to an acceptable level by setting the minimum pulse spacing to be at least 9 ns and the laser intensity to 90% of its maximum.

Based on these constraints, we set the pulse patterns as follows:

1. For the *interleaved mode*, we chose repetition periods of 180 ns and 189 ns, yielding a pattern containing 40 pulses and an unambiguous delay $\bar{t}_\mathrm{r} = 3780$ ns (unambiguous range around 566.6 m).

2. For a fair comparison, the *random mode* pattern was generated to likewise contain 40 pulses over 3780 ns while adhering to to the minimum 9 ns between pulses.

3. For the *MRR mode*, photon detections are sequentially acquired using repetition periods of 180 ns and 189 ns. The acquisition time is split roughly in two, with an equal number of pulses transmitted for each t_r.

4. The *baseline mode* was set to match the unambiguous range, periodically transmitting one pulse every 3780 ns.

5. The *ambiguous mode* was set to match the number of pulses per pattern, periodically transmitting 40 pulses every 3780 ns, leading to a repetition period of 94.5 ns ($f_\mathrm{r} \approx 10.58$ MHz).

5.1.3. Pulse Pattern Calibration and Drift Compensation

To compensate for the discrepancy between the intended and actual pulse patterns, we first make a calibration measurement of a static target. We measure the sync time $\bar{t}_{\mathrm{r,cal}}$ and form a

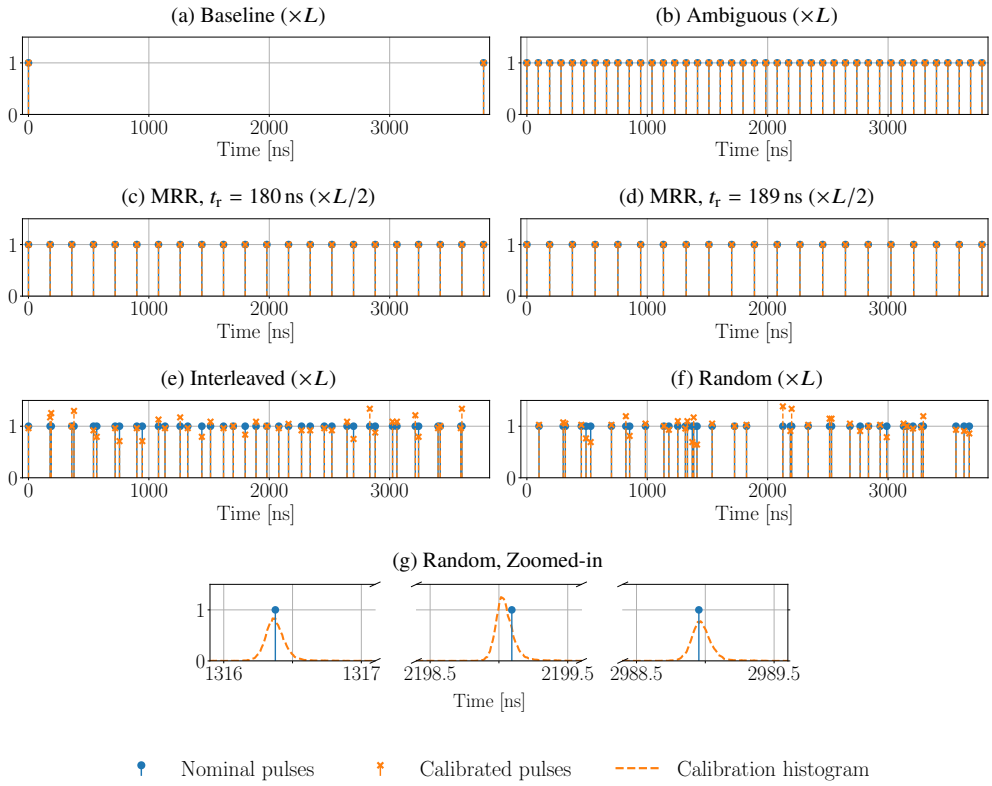

Fig. 6. Pulsing patterns for each mode. The patterns for the baseline mode (a) and ambiguous mode (b) are repeated L times. The MRR mode emits $L/2$ repetitions of the periodic sequence (c) followed by $L/2$ repetitions of the periodic sequence (d). The interleaved (e) and random (f) modes are non-periodic, with variable spacing between pulses. As we highlight in (g), this causes both the pulse amplitude to be non-uniform and the pulse timing to deviate from the nominal laser triggers.

histogram of the calibration time stamps, from which we extract the pulse times $\Pi_{\text{cal}} := \{t_p\}_{p=1}^{P}$ and amplitudes $\{a_p\}_{p=1}^{P}$. Figure 6 shows the trigger times for the various modes. The pulse amplitudes from the calibration measurement are also shown, normalized to unit amplitude and compared to the nominal pulse heights. While the modes transmitting periodic sequences (baseline, ambiguous, MRR) result in uniform pulse amplitudes, the variable pulse spacing for the interleaved and random modes results in some variability in the calibrated pulse amplitudes. In Fig. 6(g), we zoom into the histograms around several of the triggers to highlight that not only the pulse amplitude but also the pulse timing varies from the nominal trigger times. We observed that the amount of deviation depends on the pulse pattern but appeared to be repeatable over the entire acquisition time.

Despite nominally fixing the pulse times in our experiments, we observe that even for periodic pulse trains, the repetition period drifts over time. This drift is immaterial for most static TOF estimation from relative detection times, but previous work has noted this drift when analyzing absolute time stamps [24, 56]. To compensate for the oscillator drift, we record during each experiment the time stamps $(U_\ell)_{\ell=0}^{L-1}$ for all synchronization pulses in addition to the photon time

stamps. We estimate the unambiguous delay for the measurement as

$$\bar{t}_{\mathrm{r,meas}} = \frac{1}{L-1}\sum_{\ell=1}^{L-1}(U_\ell - U_{\ell-1}) = \frac{U_{L-1} - U_0}{L-1} \tag{19}$$

and rescale the pulse times as

$$\Pi_{\mathrm{meas}} := \{t_p(\bar{t}_{\mathrm{r,meas}}/\bar{t}_{\mathrm{r,cal}})\}_{p=1}^{P}. \tag{20}$$

5.1.4. Motion Stage

We constructed a custom linear motion stage from T-slot framing to perform experiments with a dynamic, controlled target. A data acquisition (DAQ, NI USB-6343) device reads position reference measurements from a draw-wire encoder (Phidgets) connected to the sliding target. The reference measurements are synchronized to the lidar time stamps by sending marker pulses to the time tagger at the DAQ sampling rate.

For eye safety reasons, our experimental setup is contained in a light-tight enclosure that limits the translation stage to roughly 3 m. To emulate long-range acquisition, we thus add an electronic delay between the sync and SPAD channels via the time tagger hardware. Although we cannot demonstrate true long-range measurements, we can achieve delays corresponding to distances that would cause significant aliasing. In all the experiments presented, we use an electronic delay of 1855 ns yielding an equivalent distance of approximately 278 m.

5.2. Experimental Results

5.2.1. Dynamic Target

We begin by comparing the various range extension modes for a dynamic target, i.e., a white card that moves along the translation stage. The limited memory of the AWG prevents us from loading sequences of different repetition rates at the same time. In combination with the lack of precise consistency in the motion trajectory, we are unable to emulate the MRR mode for dynamic targets. We compare the other four modes for two settings with different amounts of ambient light present in the scene, which is injected into the scene via a dimmable LED light strip. For all modes, we divide the photon time stamp stream into frames of duration $t_{\mathrm{a}} = 50\,\mathrm{ms}$ (20 frames/s) and independently estimate the flux, distance, and velocity for each frame using the MLE tailored to each mode. To fairly compare S and b despite different units and different number of pulses per mode, we show the signal and background flux estimates as the mean number of detections per synchronization period (3780 ns). The SBR is the ratio of those flux estimates. We also show the equivalent distance corresponding to the estimated delay. The RMSE is computed with respect to a reference measurement from the draw-wire encoder.

Lower Background Experiment Figure 7 shows the results of the experiment at the lower background setting, which is consistently estimated across all modes. The baseline mode clearly suffers from the low repetition rate, which leads to a low SBR. The resulting distance and velocity estimates are extremely imprecise, especially when the target is farthest from—and reflects fewer photons back to—the lidar system. The high repetition rate of the ambiguous mode leads to much higher SBRs and good velocity estimation performance; however, the distance is aliased and highly inaccurate. Meanwhile, both the random and interleaved modes have very similar results, with excellent recovery of both the distance (RMSE < 1 cm) and velocity (RMSE < 5 cm/s).

Higher Background Experiment Figure 8 shows the results of the experiment at the higher background setting, which leads to extremely low SBR levels. Increasing the ambient light further emphasizes the tradeoffs of the various modes. The SBR is so low for the baseline mode

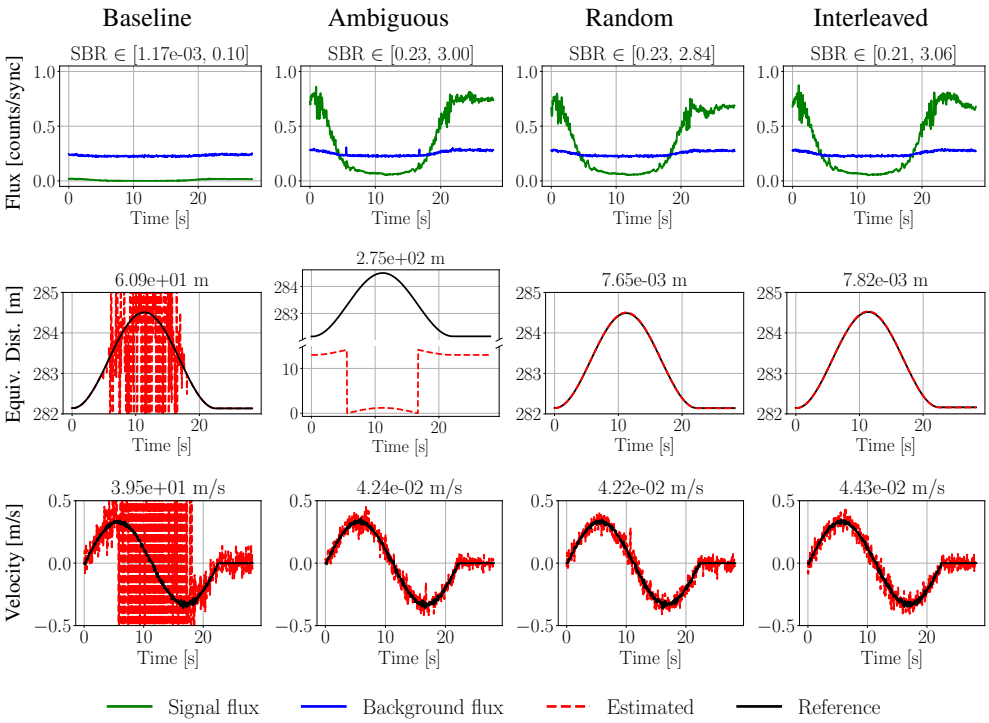

Fig. 7. Dynamic experiment under low ambient light. The baseline mode is inaccurate due to low detection rates, whereas the ambiguous mode estimates the distance incorrectly due to aliasing. Meanwhile, the random and interleaved range extension modes achieve excellent performance for all four parameters. Above each flux plot we show the range of observed SBR levels; above the distance and velocity plots, we show the RMSE computed with respect to the reference measurement.

that distance and velocity estimation is poor everywhere; whereas the ambiguous mode achieves excellent velocity estimates despite the aliased distance. The performance of the random mode is essentially unchanged by the stronger ambient light. Interestingly, the interleaved mode suffers from large distance inaccuracies in a few frames, which may correspond to errors with the CRT initialization. However, the velocity initialization, which does not depend on the CRT, is largely unaffected.

5.2.2. Static Target

Lastly, we investigate the effect of the acquisition time on the performance of the various modes. We use a static target, which allows us to also test the MRR mode by stitching together photon time stamps from each component repetition period. The laser power and ambient light settings are fixed for all modes, as are the ground truth equivalent distance (284.5357 m) and velocity (0 m/s for the static target). However, since the pulse power varies as a function of the pulse spacing, the ground truth S and b values were calibrated separately for each mode and are contained in Table 1.

We show in Fig. 9 the RMSE for S, b, τ, and v across all modes, averaging over 10 acquisitions for each t_a. Since the SBR is fixed for each acquisition, the accuracy of parameter estimation tends to improve for all modes as t_a (and thus the number of signal photon detections) increases. The baseline mode transmits the fewest pulses and thus has the worst performance. The ambiguous,

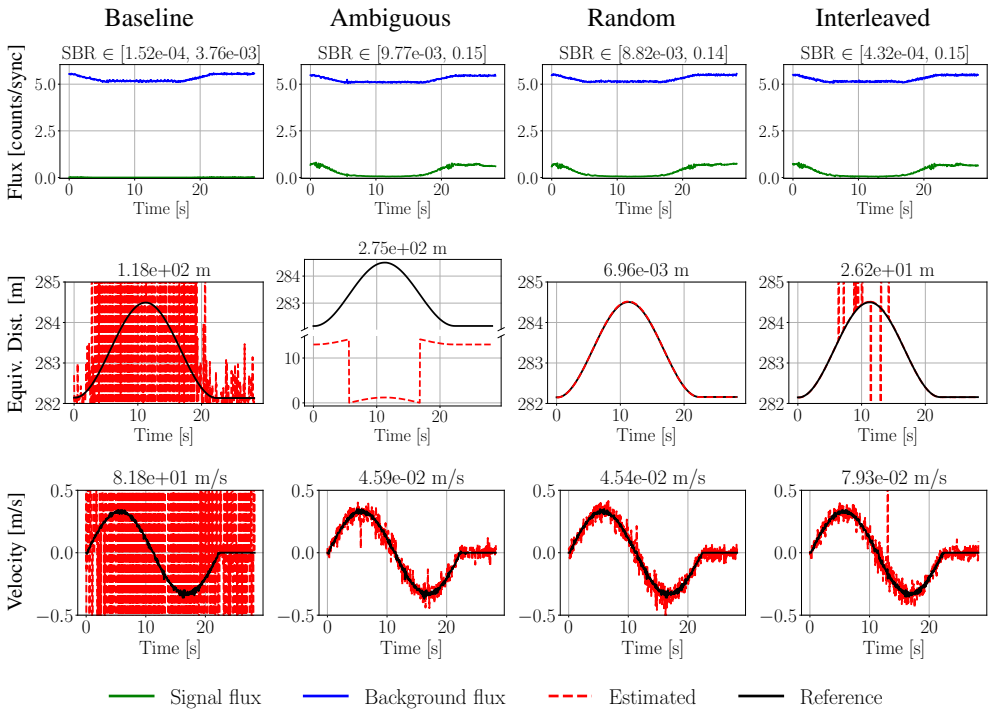

Fig. 8. Dynamic experiment under string ambient light. The baseline and ambiguous modes suffer the same types of errors as in the low-background setting. Both range extension modes perform well overall, although the interleaved mode has errors in a few frames due to the low SBR.

Mode	Baseline	Ambiguous	MRR	Random	Interleaved
S ($\times 10^{-4}$) [counts/pulse]	4.03	13.6	14.0	13.2	13.8
b ($\times 10^{-4}$) [counts/ns]	1.75	1.75	1.74	1.74	1.75

Table 1. Reference flux values for the static measurements.

random, and interleaved modes transmit the same number of pulses and thus have largely identical performance, except in two areas: 1) the ambiguous mode has poor performance for τ due to aliasing, and 2) the ambiguous mode has a slight advantage in v at low t_a, likely because the initialization is simpler and can use all signal photon detections at once. The MRR mode has slightly worse performance compared to the other range extension modes since it transmits about half the number of pulses.

6. Conclusion

In this paper, we investigated how to make both unambiguous range extension and velocity estimation compatible with single-photon lidar. We developed velocity estimation approaches for three modes of range extension using multiple repetition periods, interleaved repetition periods, or random pulse trains. All of these modes break the tradeoff between a high pulse rate (leading to ambiguous distance measurements) and a longer unambiguous rate but very few signal photons.

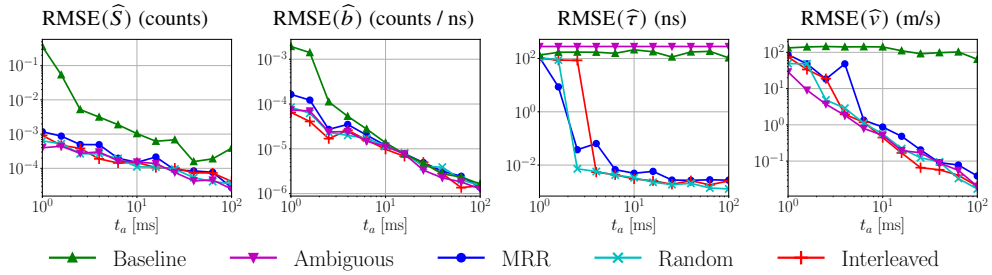

Fig. 9. Static experiment comparing the performance as a function of the acquisition time t_a. A longer acquisition time generally improves accuracy, except when the distance is aliased (ambiguous) or the SBR is too low (baseline).

Of the three modes we tested, all have their own strengths and weaknesses. The MRR mode has straightforward velocity estimation but transmits fewer pulses per acquisition time, which decreases the Fisher information and increases the CRB. Furthermore, our experimental lidar system was not capable of implementing the MRR mode for dynamic targets. The interleaved mode allows higher pulse rates—lowering the CRB—but the MLE solver requires slightly more complicated initialization. The presence of multiple simultaneous repetition rates moreover adds clutter when isolating the signal photons associated with each repetition period. The random mode also allows high pulse rates, but the lack of a simple velocity initialization strategy limits its accuracy with high-speed targets.

Some of these limitations are a result of our maximum likelihood estimator implementation, which used a two-step approach of first initializing and then refining an estimate assuming local convexity. One alternative approach would be to initialize with a grid search across two or more dimensions, which might better find the true global optimum at the cost of higher computational complexity. Another alternative may be to use other estimation implementations such as particle filters or consensus-based optimization [57] to find the global maximum of the log-likelihood from multiple initial points.

Due to eye safety concerns, our experiments were confined within an enclosure that limited the true maximum range between the lidar system and target. Using electronic delays, we were able to emulate a longer time of flight between pulse transmission and signal photon detection. However, future work should consider the additional application-specific challenges for long-range measurement. In particular, the beam divergence or turbulence could affect the distributions of photon time stamps and thus the estimation performance.

Acknowledgment. This work was completed while R.K. was an intern at MERL. J.R. thanks Dr. W. Yerazunis for his help with the design and construction of the laser enclosure.

Disclosures. All authors: Mitsubishi Electric Research Laboratories (E,P).

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

Supplemental document. See Supplement 1 for supporting content.

References

1. A. Kirmani, D. Venkatraman, D. Shin, *et al.*, "First-photon imaging," Science **343**, 58–61 (2014).
2. D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3-D and reflectivity imaging with single-photon detectors," IEEE Trans. Comput. Imaging **1**, 112–125 (2015).
3. Y. Altmann, X. Ren, A. McCarthy, *et al.*, "Lidar waveform-based analysis of depth images constructed using sparse single-photon data," IEEE Trans. Image Process. **25**, 1935–1946 (2016).

4. A. M. Pawlikowska, A. Halimi, R. A. Lamb, and G. S. Buller, "Single-photon three-dimensional imaging at up to 10 kilometers range," Opt. Express **25**, 11919–11931 (2017).
5. Z.-P. Li, X. Huang, Y. Cao, *et al.*, "Single-photon computational 3D imaging at 45 km," Photonics Res. **8**, 1532–1540 (2020).
6. Z.-P. Li, J. Ye, X. Huang, *et al.*, "Single-photon imaging over 200 km," Optica **8**, 344–349 (2021).
7. R. H. Hadfield, J. Leach, F. Fleming, *et al.*, "Single-photon detection for long-range imaging and sensing," Optica **10**, 1124–1141 (2023).
8. A. McCarthy, G. G. Taylor, J. Garcia-Armenta, *et al.*, "High-resolution long-distance depth imaging LiDAR with ultra-low timing jitter superconducting nanowire single-photon detectors," Optica **12**, 168–177 (2025).
9. J. Rapp, J. Tachella, Y. Altmann, *et al.*, "Advances in single-photon lidar for autonomous vehicles: Working principles, challenges, and recent advances," IEEE Signal Process. Mag. **37**, 62–71 (2020).
10. H. Ovrén, M. Holmberg, and M. Henriksson, "Long-range time-correlated single-photon counting lidar 3D-reconstruction from a moving ground vehicle," J. Field Rob. pp. 1–15 (2025).
11. Y. Hong, S. Liu, Z.-P. Li, *et al.*, "Airborne single-photon LiDAR towards a small-sized and low-power payload," Optica **11**, 612–618 (2024).
12. Y. Li, W.-L. Ye, Z.-P. Li, *et al.*, "Compact single-photon lidar for satellite laser ranging," Opt. Express **33**, 40876–40890 (2025).
13. J. Zhang and S. Singh, "LOAM: Lidar odometry and mapping in real-time," in *Proc. Robot. Sci. Sys,* (Berkeley, CA, USA, 2014), pp. 1–9.
14. Z. Lin, M. Hashimoto, K. Takigawa, and K. Takahashi, "Vehicle and pedestrian recognition using multilayer lidar based on support vector machine," in *Int. Conf. Mechatron. Mach. Vis. Pract. (M2VIP),* (2018), pp. 1–6.
15. Y. Ma, J. Anderson, S. Crouch, and J. Shan, "Moving object detection and tracking with Doppler LiDAR," Remote. Sens. **11** (2019). Art. no. 1154.
16. X. Peng and J. Shan, "Detection and tracking of pedestrians using Doppler lidar," Remote. Sens. **13** (2021). Art. no. 2952.
17. B. Hexsel, H. Vhavle, and Y. Chen, "DICP: Doppler iterative closest point algorithm," in *Proc. Robot. Sci. Sys,* (New York City, NY, USA, 2022), pp. 1–11.
18. W. Yang, Z. Gong, B. Huang, and X. Hong, "Lidar with velocity: correcting moving objects point cloud distortion from oscillating scanning lidars by fusion with camera," IEEE Rob. Autom. Lett. **7**, 8241–8248 (2022).
19. Y. Wu, D. J. Yoon, K. Burnett, *et al.*, "Picking up speed: Continuous-time lidar-only odometry using Doppler velocity measurements," IEEE Rob. Autom. Lett. **8**, 264–271 (2023).
20. P. Jonsson, J. Hedborg, M. Henriksson, and L. Sjöqvist, "Reconstruction of time-correlated single-photon counting range profiles of moving objects," in *Proc. SPIE Electro-Opt. Remote Sens. Photon. Tech. Appl. IX,* vol. 9649 (2015), pp. 11–18.
21. M. Alipour Sormoli, M. Dianati, S. Mozaffari, and R. Woodman, "Optical flow based detection and tracking of moving objects for autonomous vehicles," IEEE Trans. Intell. Transp. Syst. **25**, 12578–12590 (2024).
22. R. Kitichotkul, J. Rapp, Y. Ma, and H. Mansour, "Simultaneous range and velocity measurement with Doppler single-photon lidar," Optica **12**, 604–613 (2025).
23. F. Heide, W. Heidrich, M. Hullin, and G. Wetzstein, "Doppler time-of-flight imaging," ACM Trans. Graph. **34**, 1–11 (2015). Art. no. 36.
24. R. Kitichotkul, J. Rapp, Y. Ma, and H. Mansour, "Doppler single-photon lidar," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, (2025), pp. 1–5.
25. S. Fernandez-Menduina, J. Rapp, H. Mansour, *et al.*, "Tracking beyond the unambiguous range with modulo single-photon lidar," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, (2024), pp. 6–10.
26. Y. Kang, X. Wang, T. Zhang, and W. Zhao, "Single-photon ranging lidar based on multi-repetition-rate pulse train correlation and accumulation," Opt. Lett. **49**, 1628–1631 (2024).
27. Y. Liang, J. Huang, M. Ren, *et al.*, "1550-nm time-of-flight ranging system employing laser with multiple repetition rates for reducing the range ambiguity," Opt. Express **22**, 4662–4670 (2014).
28. B. Du, C. Pang, D. Wu, *et al.*, "High-speed photon-counting laser ranging for broad range of distances," Sci. Reports **8**, 1–6 (2018). Art. no. 4198.
29. C. Dai, W.-L. Ye, C. Yu, *et al.*, "Long-range photon-efficient 3D imaging without range ambiguity," Opt. Lett. **48**, 1542–1545 (2023).
30. N. J. Krichel, A. McCarthy, and G. S. Buller, "Resolving range ambiguity in a photon counting depth imager operating at kilometer distances," Opt. Express **18**, 9192–9206 (2010).
31. P. A. Hiskett, C. S. Parry, A. McCarthy, and G. S. Buller, "A photon-counting time-of-flight ranging technique developed for the avoidance of range ambiguity at gigahertz clock rates," Opt. Express **16**, 13685–13698 (2008).
32. Y. Yu, B. Liu, Z. Chen, and Z. Li, "A macro-pulse photon counting lidar for long-range high-speed moving target detection," Sensors **20** (2020). Art. no. 2204.
33. Y. Ding, H. Wu, X. Gao, *et al.*, "Coded-pulse-bunch-laser-based single-photon lidar for fast long-distance ranging," J. Opt. Soc. Amer. A. **39**, 206–212 (2022).
34. D. Reilly and G. Kanter, "High speed lidar via GHz gated photon detector and locked but unequal optical pulse rates," Opt. Express **22**, 15718–15723 (2014).
35. J. Rapp, Y. Ma, R. Dawson, and V. K. Goyal, "Dead time compensation for high-flux ranging," IEEE Trans. Signal

Process. **67**, 3471–3486 (2019).
36. D. L. Snyder and M. I. Miller, *Random Point Processes in Time and Space* (Springer Science & Business Media, 2012). Chap. 1–2.
37. W. Wang and X.-G. Xia, "A closed-form robust Chinese remainder theorem and its performance analysis," IEEE Trans. on Signal Process. **58**, 5655–5666 (2010).
38. X. Li, H. Liang, and X.-G. Xia, "A robust Chinese remainder theorem with its applications in frequency estimation from undersampled waveforms," IEEE Trans. on Signal Process. **57**, 4314–4322 (2009).
39. X. Li, X.-G. Xia, and H. Liang, "A robust Chinese remainder theorem with its applications in moving target Doppler estimation," in *Proc. IEEE Radar Conf.*, (2010), pp. 1289–1294.
40. W. Wang, X. Li, W. Wang, and X.-G. Xia, "Maximum likelihood estimation based robust Chinese remainder theorem for real numbers and its fast algorithm," IEEE Trans. on Signal Process. **63**, 3317–3331 (2015).
41. X. Li, X.-G. Xia, W. Wang, and W. Wang, "A robust generalized Chinese remainder theorem for two integers," IEEE Trans. on Inf. Theory **62**, 7491–7504 (2016).
42. Z. Hu, C. Jiang, J. Zhu, *et al.*, "Chaos single photon LIDAR and the ranging performance analysis based on Monte Carlo simulation," Opt. Express **30**, 41658–41670 (2022).
43. Z. Hu, C. Jiang, J. Zhu, *et al.*, "Efficient and robust chaos single photon LiDAR," Opt. & Laser Technol. **174** (2024). Art. no. 110623.
44. P. Du, F. Zhang, Z. Li, *et al.*, "Single-photon detection approach for autonomous vehicles sensing," IEEE Trans. Veh. Technol. **69**, 6067–6078 (2020).
45. R. Bridson, "Fast Poisson disk sampling in arbitrary dimensions," in *SIGGRAPH sketches,* (ACM, San Diego, California, 2007), p. 22.
46. J. Beck, "Irregularities of distribution. I," Acta Math. **159**, 1–49 (1987).
47. R. Kitichotkul, J. Rapp, and V. K. Goyal, "The role of detection times in reflectivity estimation with single-photon lidar," IEEE J. Sel. Top. Quantum Electron. **30** (2024). Art. no. 8800114.
48. C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization," ACM Trans. Math. Softw. **23**, 550–560 (1997).
49. M. Wei, S. Nousias, R. Gulve, *et al.*, "Passive ultra-wideband single-photon imaging," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, (Paris, France, 2023), pp. 8101–8112.
50. J. Rapp and V. K. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," IEEE Trans. Comput. Imaging **3**, 445–459 (2017).
51. S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice-Hall, Inc., 1993). Chap. 3.
52. O. Ore, "The general Chinese remainder theorem," Am. Math. Mon. **59**, 365–370 (1952).
53. T. Achterberg, T. Berthold, T. Koch, and K. Wolter, "Constraint integer programming: a new approach to integrate CP and MIP," in *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems,* vol. 5015 of *Lecture Notes in Computer Science* L. Perron and M. A. Trick, eds. (Springer Berlin Heidelberg, 2008), pp. 6–20.
54. S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," J. Mach. Learn. Res. **17**, 1–5 (2016).
55. A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," J. Control. Decis. **5**, 42–60 (2018).
56. S. Nousias, M. Wei, H. Xiao, *et al.*, "Opportunistic single-photon time of flight," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recog.*, (2025), pp. 15852–15862.
57. H. Zhang, Y. Ma, R. Kitichotkul, *et al.*, "ProxiCBO: A consensus-based method for composite optimization," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, (Barcelona, Spain, 2026). To be published.