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

TRACKING BEYOND THE UNAMBIGUOUS RANGE WITH MODULO SINGLE-PHOTON LIDAR

Fernandez-Menduina, Samuel; Rapp, Joshua; Mansour, Hassan; Greiff, Marcus; Parsons, Kieran

TR2024-021 March 19, 2024

Abstract

In single photon lidar (SPL), the laser repetition rate sets the maxi- mum distance that can be recovered unambiguously. Conventional SPL extends this maximum recordable depth by reducing the repetition rate; however, the slower acquisition speed limits the number of received photons, which may be insufficient to track fast-moving objects. Inspired by recent successes in modulo sensing, we lever- age the smoothness of typical trajectories to achieve long-range tracking beyond the unambiguous range. Although SPL naturally acquires modulo time-of-flight measurements, it introduces several challenges—including random sampling times, multiple noise sources, and absolute distance uncertainty—that are not addressed by the current modulo sensing literature. Hence, we propose an interpolation and denoising method that operates directly over the modulo samples. We further disambiguate the absolute distance based on the changing reflectivity fall-off. Monte Carlo simulations considering realistic trajectories under practical conditions show that, when properly unwrapped, the normalized mean squared error of our depth estimate decreases by over 20 dB with respect to a lidar setup whose repetition period leads to no ambiguity.

IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) 2024

© 2024 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

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

TRACKING BEYOND THE UNAMBIGUOUS RANGE WITH MODULO SINGLE-PHOTON LIDAR

S. Fernández-Menduiña,†, J. Rapp*, H. Mansour*, M. Greiff*, K. Parsons**

*Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA

†University of Southern California, Los Angeles, CA 90089, USA

ABSTRACT

In single photon lidar (SPL), the laser repetition rate sets the maximum distance that can be recovered unambiguously. Conventional SPL extends this maximum recordable depth by reducing the repetition rate; however, the slower acquisition speed limits the number of received photons, which may be insufficient to track fast-moving objects. Inspired by recent successes in modulo sensing, we leverage the smoothness of typical trajectories to achieve long-range tracking beyond the unambiguous range. Although SPL naturally acquires modulo time-of-flight measurements, it introduces several challenges—including random sampling times, multiple noise sources, and absolute distance uncertainty—that are not addressed by the current modulo sensing literature. Hence, we propose an interpolation and denoising method that operates directly over the modulo samples. We further disambiguate the absolute distance based on the changing reflectivity fall-off. Monte Carlo simulations considering realistic trajectories under practical conditions show that, when properly unwrapped, the normalized mean squared error of our depth estimate decreases by over 20 dB with respect to a lidar setup whose repetition period leads to no ambiguity.

Index Terms— Lidar, modulo sensing, single-photon detection, non-uniform sampling, modulo single-photon lidar

1. INTRODUCTION

Conventional single-photon lidar (SPL) systems operate by illuminating a scene with a periodically pulsed laser and detecting the photons scattered back using a time-resolved single-photon avalanche diode (SPAD). Photon detection times are measured with respect to the most recent pulse emission time, which enables direct time-of-flight depth imaging from as little as one photon per scene pixel [1, 2]. Importantly, the periodic illumination causes a tradeoff: a higher repetition rate results in more photon detections but decreases the depth range that can be measured unambiguously. This trade-off becomes critical in tracking applications [3] such as autonomous driving [4] that require both a high frame rate and long-distance tracking capability.

Most SPL approaches set the unambiguous range to be larger than an expected upper bound on the depth [2, 5], or they measure only relative depths for a scene that is at a long distance but lies entirely within one unambiguous range [6, 7]. Alternatively, some methods aim to make absolute distance measurements by modifying the hardware to illuminate with either non-periodic pulse trains [8, 9], multiple repetition rates [7], or different repetition rates for each pixel [10]. Unfortunately, these approaches are designed for imaging static scenes and are unsuitable for single-point tracking of a moving object.

In this paper, we explore the feasibility of modulo SPL operation, i.e., maximizing the illumination rate at the expense of the depth aliasing that occurs when recording detection times that are modulo the repetition period. To track targets beyond the unambiguous range, we can take advantage of the continuity of realistic trajectories to unwrap the modulo by applying existing methods from the modulo sensing literature [11, 12, 13]. However, modulo SPL measurements introduce several key challenges: 1) samples of the trajectory are non-uniformly spaced in time, 2) not all samples are informative due to ambient light, and informative samples are themselves noisy, and 3) modulo sensing methods unwrap sequences up to a constant ambiguity equal to an integer multiple of the modulo threshold. In this work, we address these problems through the following contributions:

1. We generalize the SPL acquisition model to account for target motion and modulo detection times;
2. We propose an interpolation and denoising algorithm that operates over modulo samples, allowing recovery of uniform, denoised modulo samples of the trajectory for use with existing unwrapping methods;
3. We demonstrate recovery of the absolute target position by finding the global offset that best fits the inverse-squared fall-off of reflectivity with distance.

Monte Carlo simulations show that for signal-to-background ratio (SBR) as low as 1 and targets moving at up to $40\,\mathrm{m/s}$, our algorithm unwraps the shape of the underlying signal in 99.75% of trials and finds the absolute depth in 99.5% of trials. Conditioned on correct unwrapping, our method reduces the normalized mean-squared error (NMSE) by over 20 dB compared to a full-range lidar whose repetition period leads to no ambiguity.

2. DATA ACQUISITION AND MODELING

We consider a target moving in one dimension and aim to acquire its position over time. Our basic setup is depicted in Fig. 1(a).

2.1. Acquisition system

Illumination. The laser illuminates the target with pulses at times $nt_{\mathrm{r}},\ n = 0, \ldots, n_{\mathrm{r}} - 1$. The laser pulse shape $s(t)$ is approximately Gaussian with root mean square (RMS) duration $t_{\mathrm{p}} \ll t_{\mathrm{r}}$. The unambiguous range is $z_{\mathrm{r}} = ct_{\mathrm{r}}/2$, where $c \approx 3 \times 10^8\ \mathrm{m/s}$ is the speed of light.

Detection. For each photon the SPAD detects, the *absolute* detection time T_ℓ, for $\ell = 0, \ldots, N_{\mathrm{d}} - 1$, is recorded with time-stamping

This work was conducted during an internship at MERL. SFM was funded in part by the Fulbright Commission in Spain.

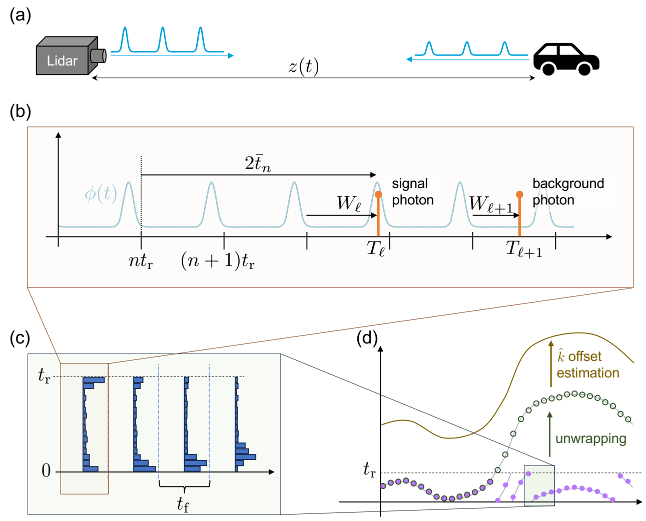

Fig. 1: Overview of modulo SPL acquisition and processing. (a) The target occupies a time-varying position $z(t)$. The laser illuminates the target with a sequence of n_r pulses with repetition period t_r. (b) Only a small fraction of pulses generate a photon detection. The time of flight $\bar{t}_n$ of a photon emitted at nt_r and detected at time T_ℓ may be longer than $t_r/2$ but can only be recorded as a wrapped timestamp W_ℓ. (c) Because the detection sequence is noisy and random, we group detection times into uniformly-spaced frames of duration t_f, and for each frame we construct a histogram, which is used for denoising. (d) We unwrap the trajectory from the uniform modulo samples and estimate the global offset from the reflectivity fall-off.

electronics as illustrated in Fig. 1(b). The output of the SPL system is then the detection time sequence $\{T_\ell\}_{\ell=0}^{N_d-1}$, where the total number of detections N_d is a random variable. Since there is no way to determine which illumination pulse generated a detected photon, the usual assumption that no aliasing occurs encourages the use of the *wrapped* detection time sequence $\{W_\ell\}_{\ell=0}^{N_d-1}$, where $W_\ell = T_\ell \mod t_r$, i.e., the time measured with respect to the most recent illumination time. The absolute detection times are typically ignored unless detector dead times are considered because they cause a time dependence in the sequence $\{T_\ell\}_{\ell=0}^{N_d-1}$ [14, 15].

2.2. Target

Depth. The target occupies the position $z(t)$ at time t. We define the time-of-flight (TOF) $\bar{t}_n$ as the time it takes the nth pulse to reach the object, so $z(nt_r + \bar{t}_n) = c\bar{t}_n$.

We model a one-dimensional trajectory with its position $z(t)$, velocity $v(t)$, and acceleration $a(t)$. We assume the trajectory follows a Singer model [16], where $a(t)$ is a correlated random process, whose correlation profile can be adjusted to mimic different maneuvers by changing the power of the innovation σ_s^2 and the decay β. Typical trajectories, including those following the Singer model, cannot be expected to be bandlimited. However, given the high sampling rate of lidar systems, we assume that the lack of bandlimitedness has negligible impact.

Reflectivity. We assume the beam divergence is negligible, so equal laser power hits the target regardless of its distance. Furthermore, we consider a Lambertian surface albedo, which implies that the light reflects isotropically, producing fall-off proportional to the inverse-square of the distance [17]. Then, the reflectivity for the nth pulse is

$$\alpha_n := \alpha_{\text{ref}}/(\bar{t}_n)^2, \tag{1}$$

where α_{ref} includes constant attenuation effects from the albedo, view angle, detector efficiency, and fall-off measured with respect to reference TOF, $\bar{t}_{\text{ref}}$. In some scenarios, e.g., with a known target, it may be possible to calibrate α_{ref} in advance of tracking.

2.3. Generalized probabilistic measurement model

Photons arrive at the detector as a realization of a Poisson process, but existing modeling (e.g., [2, 5]) assumes a static scene and $z < z_r$. Accounting for aliasing and motion, the illumination $\sum_{n=0}^{n_r-1} s(t - nt_r)$ results in a Poisson process intensity function

$$\phi(t) := b + \sum_{n=0}^{n_r-1} \alpha_n s(t - nt_r - 2\bar{t}_n), \tag{2}$$

where b is the background intensity that combines ambient light and dark counts, i.e., false detections due to thermal noise in the SPAD. Assuming the target moves sufficiently slowly so there is approximately one (potentially wrapped) pulse $s(t)$ in each repetition period, then during the nth illumination period, the photon detection rate is

$$\Phi_n := \int_{nt_r}^{(n+1)t_r} \phi(t)\,\mathrm{d}t = \alpha_q S + B, \quad q \le n, \tag{3}$$

where $S := \int_0^{t_r} s(t)\,\mathrm{d}t$, $B := bt_r$, and $\alpha_q S/B$ is the SBR. Note that a detection in period n may be produced by a photon from a previous illumination $q \le n$. We assume operation in the low-flux regime, i.e., $\Phi_n \ll 1\ \forall n$, so that SPAD dead times can be ignored [14]. Thus, photon detection constitutes an inhomogenous Poisson process with time-varying rate function $\phi(t)$. Wrapped detection times are random variables with density $f_W^{(n)}(t) = \phi(t)/\Phi_n$, for $t \in [nt_r, (n+1)t_r)$. Both reflectivity and TOF vary as a function of n, so these random variables are not identically distributed, but we assume they are independent.

Unlike conventional modeling, the formulation (2) allows not only for the case when $\bar{t}_n > t_r/2$, but also when the pulse shape $s(t)$ lies across a wrapping boundary. These events represent the points when the detector introduces the modulo discontinuity. Although the distribution of the detection times has to be modified to account for the wrapping effect [18], since we assume $t_p \ll t_r$, the effect of the pulse wrapping onto itself is negligible.

3. MODULO SENSING

Modulo sensing aims to acquire samples of a signal minimizing quantization noise [12, 19] and clipping distortion [11]. Since the noise introduced by the detection-time quantization is negligible compared to other sources of uncertainty, we focus on the latter. Consider a signal $g(t)$ and an analog-to-digital converter with dynamic range $[-\lambda, \lambda]$. The goal is to recover $g(t)$ from its uniform and modulo reduced samples $\mathcal{M}_\lambda(g[n]) := (g[n] + \lambda \mod 2\lambda) - \lambda$. The recovery process comprises two steps: 1) recovering the unwrapped samples from the wrapped samples and 2) applying classical reconstruction techniques to recover $g(t)$. Theoretical results show that recovery of noiseless bandlimited signals is guaranteed with sufficient oversampling [11], while in practice the oversampling requirement depends on the noise power and recovery algorithm used [13].

Unwrapping algorithms follow two main approaches. Temporal domain techniques exploit the redundancy between samples [12, 20]; for instance, unlimited sampling [11, 21, 22] computes higher-order finite differences between samples and generalizes classical phase unwrapping using the first-order difference (FDPU) [23, 24]. Joint temporal-spectral domain techniques exploit finite support in both the temporal and the frequency domains [25, 13, 26]. Several works have demonstrated depth estimation with modulo measurements [27, 28, 29], but these methods encode depth in the phase or frequency of a signal while the modulo is applied to the amplitude.

Limitations for SPL. In modulo SPL, the modulo is applied directly to the TOF. Furthermore, target tracking via SPL introduces a set of differences with respect to the typical modulo sensing setup. First, existing unwrapping methods assume and often require uniform samples, e.g., for a well-characterized spectrum; however, the stochastic nature of photon detection yields samples in SPL that are non-uniformly spaced in time. Second, SPL measurements contain uninformative background detections as well as pulse-shape uncertainty for signal photons. Finally, the modulo operator leaves an offset ambiguity equal to an integer multiple of the modulo threshold. Existing solutions limit applicability to signals with no DC component [13] or ensure the first few samples are unaffected by the modulo operation [12], which cannot be guaranteed in a tracking scenario.

4. MODULO SINGLE PHOTON LIDAR

4.1. Interpolation and Denoising of Modulo Samples

Although we aim to track moving targets, we assume that the displacement is negligible within short time frames of size t_{f}. Since the RMS duration of $s(t)$ already induces some uncertainty, we set the duration $t_{\mathrm{f}} = p_{\mathrm{f}}\, c t_{\mathrm{p}}/v_{\max}$, where $v_{\max}$ is an upper bound on the maximum speed of the target and p_{f} is a parameter controlling the bias-variance trade-off. We split the original sequence into n_{f} possibly overlapping frames of length t_{f}. Inside the mth frame, we have K_m identically distributed samples, where K_m is a random variable.

As shown in Fig. 1(c), we form a histogram of detection times in each frame to estimate the depth. In conventional SPL, the maximum-likelihood (ML) depth is approximated by the log-matched filter assuming $B = 0$ [30, 2]. Applying a censoring approach that uses a window to reject likely background counts before filtering typically improves results [5]. To account for the modulo, we update the censoring to wrap around the modulo boundary, and the log-matched filter becomes

$$D_m := \operatorname*{arg\,max}_{t \in (0, t_{\mathrm{r}}]} \sum_{\ell \in \mathcal{I}_m} \log \left(\sum_{i=-\infty}^{\infty} s(W_\ell - t - i t_{\mathrm{r}}) \right), \quad (4)$$

where $\mathcal{I}_m$ contains the indices of the detections within the mth frame. Since $t_{\mathrm{p}} \ll t_{\mathrm{r}}$, we approximate (4) by the index of the maximum of the circular convolution between the pulse and the histogram of the photons received during the frame.

By assigning the estimate D_m to the center of frame m and repeating the process for $m = 0, \ldots, n_{\mathrm{f}} - 1$, we obtain a uniformly spaced sequence of n_{f} samples whose amplitude lies in $[0, t_{\mathrm{r}})$, as shown by the purple dots in Fig. 1(d). This sequence represents the modulo samples denoised and interpolated from the raw detections. Thus, we can apply conventional unwrapping algorithms and denote the unwrapped measurements as $\{\tilde{U}_m\}_{m=0}^{n_{\mathrm{f}}-1}$ (green circles in Fig. 1(d)). Although we focus on SPL, many existing modulo unwrapping algorithms may benefit from our approach to denoising and interpolating by grouping samples together, in particular if noise power is large and oversampling factor is high.

4.2. Recovering Absolute Position via Reflectivity

Modulo sensing algorithms are able to unwrap a sequence only up to an integer multiple of t_{r}. We recall, however, that the reflectivity fall-off in (1) depends on the absolute TOF. Although the reflectivity estimate alone is too noisy to reconstruct the trajectory, we can leverage the changing fall-off to recover the global offset without constraining the signal, e.g., to have finite energy. As with the depth, we perform censoring and reflectivity estimation for a sequence of n_α frames, although the frame duration, amount of overlap, and censoring window length may differ. We obtain a sequence of values $\hat{\alpha}_j$, for $j = 0, \ldots, n_\alpha - 1$, and interpolate the unwrapped TOF to the same temporal positions, which we denote by $\tilde{U}_j$ with some abuse of notation. We then define a set of candidate TOFs as $U_{j,k} = \tilde{U}_j + k t_{\mathrm{r}}$, where $k \in \mathbb{N}$ sets the global offset of the signal. For each k, we compute the reflectivity decay as $\theta_{j,k} = 1/U_{j,k}^2$. When the noise in the estimated TOFs is negligible compared to the error in the estimated reflectivity, and modeling the latter as a Gaussian i.i.d. process, we compute

$$\hat{k} := \operatorname*{arg\,min}_{k=0,1,\ldots} \sum_{j=0}^{n_\alpha - 1} \left(\hat{\alpha}_j - \alpha_{\mathrm{ref}} \theta_{j,k}\right)^2 . \quad (5)$$

Then, we reconstruct the final, fully unwrapped estimates of the TOFs by adding $\hat{k} t_{\mathrm{r}}$ to the unwrapped sequence before interpolation, $U_m := \tilde{U}_m + \hat{k} t_{\mathrm{r}}$. If α_{ref} is known, it can be used directly in Eq. (5); otherwise, we estimate it as $\hat{\alpha}_{\mathrm{ref},k} = \sum_{j=0}^{n_\alpha-1} \hat{\alpha}_j \theta_{j,k} \big/ \sum_{j=0}^{n_\alpha-1} \theta_{j,k}^2$, which requires target motion to resolve the ambiguity.

5. EXPERIMENTAL VALIDATION

Metrics. We compute the proportion of correct unwrappings as

$$\mathrm{CU} := \frac{\sum_{m=0}^{n_{\mathrm{f}}-1} \min(\psi_m(\bar{t}_m), \psi_m(U_m))}{\sum_{m=0}^{n_{\mathrm{f}}-1} \max(\psi_m(\bar{t}_m), \psi_m(U_m))}, \quad (6)$$

where $\psi_m(x) := |\Delta[(\bar{t}_m \mod t_{\mathrm{r}})/t_{\mathrm{r}} - x/t_{\mathrm{r}}]|$ computes discontinuities with respect to the true TOF at the center of the mth frame, which we denote by $\bar{t}_m$ with some abuse of notation. The operator $[\cdot]$ rounds its argument and Δ denotes the first-order difference. We also compute the fraction of times that $\hat{k}$ was correct (success rate, SR). We assess distortion via the normalized mean squared error:

$$\mathrm{NMSE} := 10 \log_{10} \left(\sum_{m=0}^{n_{\mathrm{f}}-1} (\bar{t}_m - U_m)^2 \Big/ \sum_{m=0}^{n_{\mathrm{f}}-1} (\bar{t}_m)^2 \right), \quad (7)$$

which summarizes the error for both TOF and distance. We also compute the conditional NMSE when both $\mathrm{CU} = 1$ and $\hat{k}$ is correct.

Results. To simulate measurements, we use a short repetition period, $t_{\mathrm{r}} = 7.5\,\mathrm{ns}$. Following [2], we also set $t_{\mathrm{p}} = 270\,\mathrm{ps}$ and $\alpha_{\mathrm{ref}} = 0.315$, which accounts for quantum efficiency and the reflectivity of the object. Furthermore, we set $\bar{t}_{\mathrm{ref}} = 3.75\,\mathrm{ns}$, $S = 10^{-2}$, $B = \alpha_{\mathrm{ref}} S/\mathrm{SBR}_0$, and we vary SBR_0 for different realizations. We consider a set of 4 Singer models, combining $\beta \in \{5, 75\}$ with $\sigma_{\mathrm{s}} \in \{5 \times 10^3, 4 \times 10^3\}$. The target switches between models with

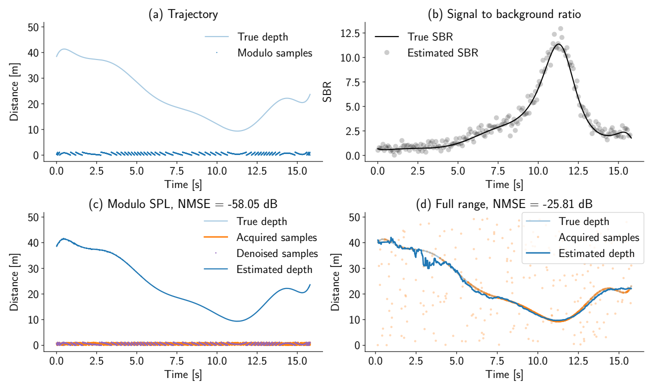

Fig. 2: (a) One realization of a Singer trajectory and its modulo samples. (b) True and estimated SBR during the acquisition period. (c) Acquired samples, denoised samples, and depth estimates using modulo SPL with $t_r = 7.5\,\mathrm{ns}$. (d) Acquired samples and depth estimates using a full-range lidar that avoids ambiguity ($t_r = 330.92\,\mathrm{ns}$). The maximum distance to the target is $41.36\,\mathrm{m}$, its maximum speed is $16.60\,\mathrm{m/s}$, and the minimum SBR is 0.58.

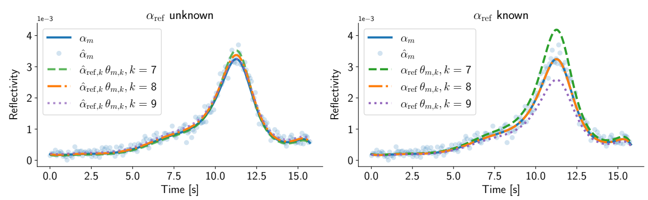

Fig. 3: True reflectivity, estimate of the reflectivity, and candidates $\theta_{m,k}$, each of them scaled by (left) their estimate of α_{ref} and (right) the true value of α_{ref}. The true offset is $k = 8$. While it is possible to recover the global offset in both scenarios, the simultaneous optimization in both α_{ref} and k when the former is unknown shrinks the gap between candidates.

probability 2.5×10^{-3}. We use first-order difference phase unwrapping (FDPU) to recover the sequence.

To compute the depth, we allow for overlapping frames so that the number of samples is the same before and after processing, i.e., $n_{\mathrm{f}} = N_{\mathrm{d}}$, while we consider no overlapping to estimate the reflectivity. Empirically, we observed the best results with overlapping frames and a censoring window of $3t_{\mathrm{p}}$ for depth estimation. For reflectivity, we obtained the best results using non-overlapping frames of length equal to the maximum between $ct_{\mathrm{p}}/v_{\max}$ and $0.015\,\mathrm{s}$, together with a censoring window of $10t_{\mathrm{p}}$. We assume α_{ref} is unknown unless otherwise specified.

Fig. 2 shows an example of how modulo SPL acquisition can lead to lower overall distortion. We plot an instance of a Singer trajectory and compare its recovery using modulo SPL and a full-range (FR) lidar whose laser repetition period is $44\times$ larger ($330.92\,\mathrm{ns}$) so that the trajectory lies within the unambiguous range. We use the same frame-based censoring and estimation for both lidars, but to compute t_{f}, we set $p_{\mathrm{f}} = 1$ for modulo SPL; for the FR lidar, we use $p_{\mathrm{f}} = 50$ to compensate for the higher variance due to fewer signal photon detections. We also depict the true SBR and our estimate, obtained by following the procedure in Sec. 4.2 to recover the reflectivity based on the values of S and B detailed above. Using the

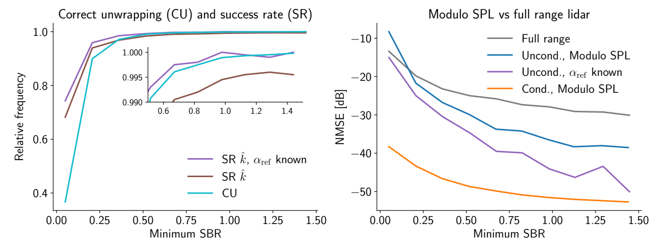

Fig. 4: Results for Monte Carlo simulations (2000 trials) using modulo SPL and a full range lidar (FR) that avoids ambiguity (repetition period $8.62\times$ longer). We varied the minimum SBR of the trajectory and assumed α_{ref} was unknown. We provide (left) the ratio of correct unwrappings (CU, Eq. (6)), and the success rate (SR) for $\hat{k}$, and (right) the normalized mean squared error (NMSE), both unconditional (Uncond.) and conditional on CU $= 1$ and SR $= 1$ (Cond.). We also provide the NMSE and SR of $\hat{k}$ when α_{ref} is known. Average maximum speed is $40\,\mathrm{m/s}$.

same trajectory, we further illustrate estimation of the global offset $\hat{k}$ when α_{ref} is unknown and known (Fig. 3). The offset recovery is significantly easier when α_{ref} is known.

In Fig. 4 we evaluate the modulo SPL performance obtained via Monte Carlo simulations with 2000 trials, for a range of minimum SBR values over the trajectory. We compare the NMSE against a FR lidar with $8.62\times$ longer repetition period. We compute t_{f} using $p_{\mathrm{f}} = 1$ for modulo SPL and $p_{\mathrm{f}} = 10$ for the FR lidar, which minimizes the NMSE for the FR system. We also provide the NMSE (conditional on CU $= 1$ and correct offset recovery) and the success rate for $\hat{k}$ when α_{ref} is known. To ensure the FR lidar can track the trajectory unambiguously, we scale the depth excursion so that it is never folded. The average maximum speed of the target across the trajectories is approximately $40\,\mathrm{m/s}$. At very low SBR, unwrapping errors cause modulo SPL to perform worse than FR lidar. However, as the SBR increases, unwrapping and offset recovery both improve, resulting in less distortion (lower NMSE) for the modulo SPL on average. The performance bottleneck becomes offset recovery, which is further improved if α_{ref} is known in advance. Finally, conditioned on cases in which unwrapping and offset recovery are both correct, the performance is dramatically better than FR lidar, with more than $20\,\mathrm{dB}$ improvement in trajectory recovery.

6. CONCLUSION

In this paper, we addressed the problem of using an SPL system to track a target in one dimension beyond the unambiguous range imposed by the repetition period of the laser. We introduced a method to denoise and interpolate samples that leverages the scenario where the oversampling factor is high. We demonstrated that computational reconstruction of the trajectory from its modulo samples is possible, and that using shorter repetition periods can reduce the distortion in the depth estimates at the expense of slightly more complicated processing. Future work may focus on improving the offset recovery success rate.

7. REFERENCES

[1] A. Kirmani, D. Venkatraman, D. Shin, A. Colaco, et al., “First-photon imaging,” *Science*, vol. 343, no. 6166, pp. 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*, vol. 1, no. 2, pp. 112–125, 2015.

[3] Z. Li, B. Liu, H. Wang, H. Yi, and Z. Chen, "Advancement on target ranging and tracking by single-point photon counting lidar," *Opt. Express*, vol. 30, no. 17, pp. 29907–29922, 2022.

[4] J. Rapp, J. Tachella, Y. Altmann, S. McLaughlin, and V. K. Goyal, "Advances in single-photon lidar for autonomous vehicles: Working principles, challenges, and recent advances," *IEEE Signal Process. Mag.*, vol. 37, no. 4, pp. 62–71, July 2020.

[5] J. Rapp and V. K. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," *IEEE Trans. Comput. Imaging*, vol. 3, no. 3, pp. 445–459, 2017.

[6] Z.-P. Li, X. Huang, Y. Cao, B. Wang, et al., "Single-photon computational 3D imaging at 45 km," *Photonics Res.*, vol. 8, no. 9, pp. 1532–1540, Apr. 2020.

[7] Z.-P. Li, J. Ye, X. Huang, P.-Y. Jiang, et al., "Single-photon imaging over 200 km," *Optica*, vol. 8, no. 3, pp. 344–349, 2021.

[8] 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*, vol. 16, no. 18, pp. 13685–13698, Sept. 2008.

[9] N. J. Krichel, A. McCarthy, and G. S. Buller, "Resolving range ambiguity in a photon counting depth imager operating at kilometer distances," *Opt. Express*, vol. 18, no. 9, pp. 9192–9206, 2010.

[10] C. Dai, W.-L. Ye, C. Yu, X. Huang, et al., "Long-range photon-efficient 3D imaging without range ambiguity," *Opt. Lett.*, vol. 48, no. 6, pp. 1542–1545, Feb. 2023.

[11] A. Bhandari, F. Krahmer, and R. Raskar, "On unlimited sampling," in *Proc. Int. Conf. Sampling Theory Appl.*, July 2017, pp. 31–35.

[12] O. Ordentlich, G. Tabak, P. K. Hanumolu, A. C. Singer, and G. W. Wornell, "A modulo-based architecture for analog-to-digital conversion," *IEEE J. Sel. Topics Signal Process.*, vol. 12, no. 5, pp. 825–840, 2018.

[13] E. Azar, S. Mulleti, and Y. C. Eldar, "Robust unlimited sampling beyond modulo," arXiv:2206.14656, 2022.

[14] J. Rapp, Y. Ma, R. M. A. Dawson, and V. K. Goyal, "Dead time compensation for high-flux ranging," *IEEE Trans. Signal Process.*, pp. 3471–3486, Oct. 2019.

[15] J. Rapp, Y. Ma, R. Dawson, and V. Goyal, "High-flux single-photon lidar," *Optica*, vol. 8, no. 1, pp. 30–39, Nov. 2021.

[16] R. A. Singer, "Estimating optimal tracking filter performance for manned maneuvering targets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-6, no. 4, pp. 473–483, 1970.

[17] M. Liao, L. Wang, R. Yang, and M. Gong, "Light fall-off stereo," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, June 2007, pp. 1–8.

[18] P. E. Jupp and K. Mardia, *Directional Statistics*, vol. 2, Wiley Online Library, 2000.

[19] L. G. Ordoñez, P. Ferrand, M. Duarte, M. Guillaud, and G. Yang, "On full-duplex radios with modulo-ADCs," *IEEE Open J. Commun. Soc.*, vol. 2, pp. 1279–1297, 2021.

[20] E. Romanov and O. Ordentlich, "Above the Nyquist rate, modulo folding does not hurt," *IEEE Signal Process. Lett.*, vol. 26, no. 8, pp. 1167–1171, 2019.

[21] S. Fernández-Menduiña, F. Krahmer, G. Leus, and A. Bhandari, "Computational array signal processing via modulo non-linearities," *IEEE Trans. Signal Process.*, vol. 70, pp. 2168–2179, 2022.

[22] A. Bhandari and F. Krahmer, "HDR imaging from quantization noise," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2020, pp. 101–105.

[23] R. W. Schafer, "Echo removal by discrete generalized linear filtering," Technical Report 466, MIT Research Laboratory of Electronics, Cambridge, MA, 1969.

[24] K. Itoh, "Analysis of the phase unwrapping algorithm," *Appl. Optics*, vol. 21, no. 14, pp. 2470, July 1982.

[25] A. Bhandari, F. Krahmer, and T. Poskitt, "Unlimited sampling from theory to practice: Fourier-Prony recovery and prototype ADC," *IEEE Trans. Signal Process.*, vol. 70, pp. 1131–1141, 2022.

[26] S. B. Shah, S. Mulleti, and Y. C. Eldar, "Lasso-based fast residual recovery for modulo sampling," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, 2023, pp. 1–5.

[27] T. Feuillen, M. Alaee-Kerahroodi, A. Bhandari, B. Shankar M. R., and B. Ottersten, "Unlimited sampling for FMCW radars: A proof of concept," in *Proc. IEEE Radar Conf.*, Mar. 2022, pp. 1–5.

[28] G. Shtendel and A. Bhandari, "HDR-TOF: HDR time-of-flight imaging via modulo acquisition," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2022, pp. 3808–3812.

[29] T. Feuillen, B. Shankar M. R. R., and A. Bhandari, "Unlimited sampling radar: Life below the quantization noise," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, June 2023, pp. 1–5.

[30] I. Bar-David, "Communication under the Poisson regime," *IEEE Trans. Inform. Theory*, vol. 15, no. 1, pp. 31–37, 1969.