DEAD TIME COMPENSATION FOR HIGH-FLUX DEPTH IMAGING

Joshua Rapp[⋆†], *Yanting Ma*[⋆], *Robin M. A. Dawson*[†], *and Vivek K Goyal*[⋆]

[⋆] Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215
[†] The Charles Stark Draper Laboratory, Inc., Cambridge, MA 02139

ABSTRACT

Time-correlated single photon counting (TCSPC) is a powerful technique for lidar depth imaging, allowing for accurate range measurements from very low light levels. However, single-photon detectors used in TCSPC have a dead time after each photon detection, which blocks registration of subsequent photons arriving within that dead time, causing a distortion of the detection time distribution. The most common approach to avoiding dead time distortion is to optically reduce the photon arrival rate such that with high probability no photons arrive during the dead time. However, this prevents the high photon flux acquisition necessary for real-time applications such as autonomous navigation. In this paper, we propose a dead time compensation method that enables fast data acquisition with dead time-limited detectors. Specifically, we model dead time-affected detection times as a Markov chain, present a simple method for approximating the stationary distribution, and estimate depths using a log-matched filter matched to that distribution. Our method applies to multimodal imaging systems where a standard camera is used in conjunction with lidar to provide information about scene reflectivity. Simulation results for real 3D scenes show that our method reduces the root mean squared error by several orders of magnitude for the same acquisition time.

Index Terms— Computational 3D imaging, lidar, single photon detection, dead time, Markov chain

1. INTRODUCTION

Autonomous vehicles promise safer and more efficient transportation systems, with anticipated outcomes including fewer fatalities, increased mobility, reduced congestion, and changes in vehicle ownership patterns [1, 2]. A key element for realizing fully autonomous vehicles is the environmental sensing system for mapping, localization, object detection, etc., which directly influences driving decisions [3]. Many proposed systems combine input measurements from radar, lidar, and conventional cameras, which each have advantages for particular tasks and environmental conditions [4–6]. Lidar is an especially useful modality, as it provides high transverse and longitudinal resolution and can operate in a wide variety of lighting conditions. One form of lidar currently under development for autonomous vehicles is based on time-correlated single photon counting (TCSPC). Due to its sensitivity to individual photons and high temporal resolution, TCSPC can reliably produce accurate range measurements from repeated illumination and detection, even from extremely low light levels [7–9]. Importantly, TCSPC lidar has been shown to sufficiently sense weakly-reflective objects at long distances using eye-safe illumination levels [10], which are necessary conditions for vehicular deployment.

One of the key downsides to the single photon detectors used in TCSPC is that after each detection they have a dead time, a period during which no other incident photons can be registered. While the photon arrival statistics are well understood, the detection-dependent dead times generally make the detection process much more difficult to analyze. A common approach to avoid the complications of dead time is to lower the photon flux at the detector to ensure that at most 5% of illuminations generate photon arrivals [11]. The arrival rate is then low enough that it is unlikely for the dead time after a detection to block a subsequent photon from being detected. However, this mode of operation is undesirable for autonomous vehicles, which must collect data as rapidly as possible in order to make real-time navigation decisions. An alternative approach is to acquire data at high flux and naïvely ignore the effects of dead time on data acquisition. These effects are often called "pileup," since early photons block later photons, resulting in a skewing of the probability density function (PDF) of detections towards earlier times [12]. Still, ignoring dead time is inadvisable for vehicular lidar, as scene patches returning high photon flux would be inaccurately perceived to be closer than their true depth. Unfortunately, classical approaches to dead time compensation assume synchronous dead time models [13–17], which do not apply to modern TCSPC systems [18].

In [19], we characterized the detection time distribution affected by dead time as the stationary distribution of a Markov chain, based on which we proposed a depth estimator with improved accuracy provided that the scene reflectivity and the ambient light level is known. In this work, we present a simple approximation to the transition probabilities of the Markov chain, and the limiting distribution obtained from the approximated transition probability matrix is used in the depth estimator. The main contribution of this paper is a framework for incorporating the depth estimator into forming full depth images, such as those required for autonomous navigation. Supplemental reflectivity information has previously been used to improve depth sensing tasks such as superresolution [20], object segmentation [21], and vehicle detection [22], including aiding in TCSPC lidar depth imaging [23]. We specifically describe how a multimodal sensing system supplementing TCSPC lidar with a grayscale camera can provide the reflectivity information for our depth estimator. To reduce computational time, we quantize reflectivity uniformly and show by a numerical example how depth estimation accuracy is affected by the number of quantization levels. With this approach, we demonstrate the potential for fast, photon-efficient depth acquisition that could be incorporated into real-time imaging systems.

2. ACQUISITION MODEL

2.1. Photon Arrival Process

TCSPC lidar data is acquired by raster scanning a laser over a scene and detecting back-reflected photons from the repeated pulsed illumination at each scene patch indexed by $(i, j) \in \{(1, \ldots, n_i) \times$

$(1,\ldots,n_j)\}$. The arrival times of photons at the detector from each (i,j) are described by a Poisson process with intensity $\lambda_{i,j}(t)$ [24]. Due to the pulsed illumination, each $\lambda_{i,j}(t)$ is periodic with period t_{r}, and n_{r} is the number of illumination periods in an acquisition. Within one period, the intensity is a mixture of two components

$$\lambda_{i,j}(t) = \lambda^s_{i,j}(t) + \lambda^b_{i,j}(t), \tag{1}$$

where $\lambda^s_{i,j}(t)$ is the arrival intensity of signal photons due to the illumination back-reflected from (i,j) and $\lambda^b_{i,j}(t)$ is the arrival intensity of background photons due to ambient light and dark counts. The signal intensity is time-varying and scene dependent:

$$\lambda^s_{i,j}(t) = \alpha_{i,j}\beta s(t - 2z_{i,j}/c), \tag{2}$$

where $\alpha_{i,j} \in [0,1]$ is the *reflectivity* of (i,j), combining attenuation effects due to object reflectance, radial falloff, view angle, detector quantum efficiency, etc.; $\beta \in [0,\infty)$ is the *illumination gain*, corresponding to the expected number of photon arrivals per illumination from a unit reflectivity object; $s(t)$ is the illumination pulse shape normalized such that $\int_0^{t_r} s(t)dt = 1$; $z_{i,j}$ is the depth of the surface at (i,j); and c is the speed of light. Let $\boldsymbol{z}$ and $\boldsymbol{\alpha}$ denote the full depth and reflectivity images, respectively. In this work, we assume the pulse shape is Gaussian $s(t) \propto \exp[-t^2/(2\sigma_{\mathrm{p}}^2)]$ with standard deviation σ_{p}—although our approach holds for any pre-calibrated $s(t)$—and β and $\lambda^b_{i,j}(t) = \lambda^b$ are constant for all (i,j). The parameters σ_{p}, β, λ^b are assumed to be known, since they are easy to obtain via calibration. The expected number of signal photon arrivals per period is defined as $S_{i,j} = \int_0^{t_r} \lambda^s_{i,j}(t)dt = \alpha_{i,j}\beta$, and the expected number of background arrivals is $B = \lambda^b t_{\mathrm{r}}$. The total flux at each point is given as $\Lambda_{i,j} = S_{i,j} + B$, and the average signal-to-background ratio is $\mathrm{SBR} = (\beta/(n_i n_j B)) \sum_{i=1}^{n_i}\sum_{j=1}^{n_j} \alpha_{i,j}$. The PDF of photon arrival times in one period (i.e., relative to the most recent illumination) is [24]

$$f^{\mathrm{A}}_{i,j}(t) = \lambda_{i,j}(t)/\Lambda_{i,j}. \tag{3}$$

2.2. Markov Chain Model for Detection Times

Whereas photon arrival times at the detector are statistically independent in accordance with the Poisson process model, the actual registration of photons by the detector is not a Poisson process due to the detection-dependent dead time. Modern TCSPC systems are modeled as nonparalyzable and asynchronous, which means that after a photon is detected, the detector is immediately forced to be insensitive to incoming photons and is then reset after a fixed dead time duration t_{d}, irrespective of the position within an illumination period [18, 19]. Note that the next detected photon is the first photon that arrives at the detector after the reset. Because the arrival times are statistically independent and the reset time only depends on the last detection, we can model a sequence of detection times as a Markov chain. The Markov chain has a continuous state space $[0, t_{\mathrm{r}})$. For practical applications, it is reasonable to consider a discrete approximation, because measurements from practical detectors are quantized. In the following, we provide a simple method to construct the approximated transition probability matrix; a more rigorous treatment can be found in [19].

The same formulation applies to all (i,j), so we proceed without the unnecessary pixel index notation. Let $\{T_\ell\}_{\ell\geq1}$ with $T_\ell \in [0,\infty)$ denote the set of *absolute* detection times and let $\{X_\ell\}_{\ell\geq1}$ denote detection times defined as $X_\ell := T_\ell \mod t_{\mathrm{r}} \in [0, t_{\mathrm{r}})$. It is well-known that given the most recent absolute detection time T_ℓ, the next absolute detection time $T_{\ell+1}$ takes values in $[T_\ell + t_{\mathrm{d}}, \infty)$ with conditional PDF [24]

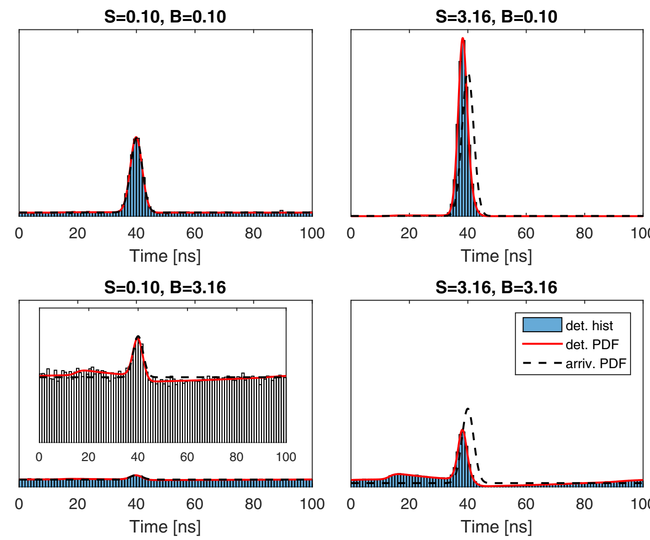

Fig. 1: A close match between the histogram of detection times and the computed PDF $\mathbf{f}^{\mathrm{D}}$ shows that our Markov chain method accurately predicts the effect of dead time and the resulting limiting distribution of detection times. The plot further demonstrates how the arrival PDF $\mathbf{f}^{\mathrm{A}}$, a discretization of f^{A} (3), is more distorted as the overall flux increases, with the dead time causing ripples in the waveform and a shift in the peak towards earlier times. The simulation was generated with $\sigma_{\mathrm{p}} = 2$ ns, $n_{\mathrm{r}} = 50000$ illuminations, $t_{\mathrm{r}} = 100$ ns, and $t_{\mathrm{d}} = 75$ ns. The vertical scale is constant for each subplot, with a magnified inset to show the otherwise unseen ripple when $S = 0.1$ and $B = 3.16$.

$$f_{T_{\ell+1}|T_\ell}(t_{\ell+1}|t_\ell) = \lambda(t_{\ell+1}) \exp\left(-\int_{t_\ell+t_{\mathrm{d}}}^{t_{\ell+1}} \lambda(\tau)d\tau\right). \tag{4}$$

Notice that $T_{\ell+1} > kt_{\mathrm{r}} + T_\ell + t_{\mathrm{d}}$ implies there are no photon arrivals for successive k periods, which has probability $\exp(-k\Lambda)$ and this probability decays exponentially in k. Therefore, we can approximate $f_{T_{\ell+1}|T_\ell}$ as being supported on a finite interval $[T_\ell + t_{\mathrm{d}}, Kt_{\mathrm{r}}] \subset [T_\ell + t_{\mathrm{d}}, \infty)$ for some positive integer K. Note that K need not be large for the approximation to be accurate, especially when Λ is large.

With the above observation, we can now construct our approximated transition probability matrix of the Markov chain $\{X_\ell\}_{\ell\geq1}$ that models detection times. We first partition the state space $[0, t_{\mathrm{r}})$ into n_{b} equally-spaced time bins and the bin size Δ satisfies that $n_{\mathrm{b}} := t_{\mathrm{r}}/\Delta$ and $n_{\mathrm{d}} := t_{\mathrm{d}}/\Delta$ are positive integers. Let $\{\Delta_n\}_{1\leq n\leq n_{\mathrm{b}}}$ denote the bins and the transition probability matrix is then an $n_{\mathrm{b}} \times n_{\mathrm{b}}$ matrix P with $P_{m,n}$ an approximation to the conditional probability $Pr(X_{\ell+1} \in \Delta_n | X_\ell \in \Delta_m)$. We further partition the truncated time interval $[0, Kt_{\mathrm{r}}]$ into Kn_{b} time bins with bin size Δ and bin centers $\{b_n\}_{1\leq n\leq Kn_{\mathrm{b}}}$. Entries of P are then defined as

$$P_{m,n} = \sum_{k=\min\{k:kn_{\mathrm{b}}+n>m+n_{\mathrm{d}}\}}^{K-1} \lambda(b_n) \exp\left(-\sum_{u=m+n_{\mathrm{d}}}^{kn_{\mathrm{b}}+n} \lambda(b_u)\right)$$

for $m,n = 1,\ldots,n_{\mathrm{b}}$. Each summand is a discrete approximation of

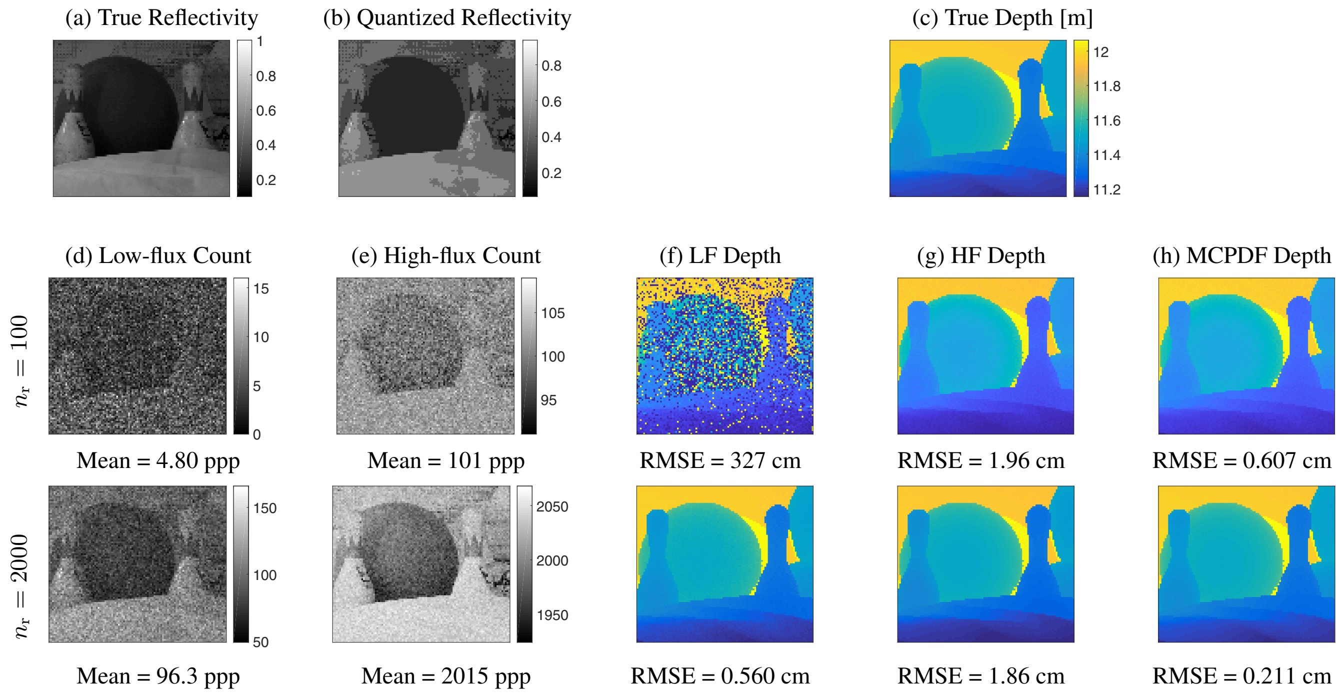

Fig. 2: Results for simulated detections from a true 3D scene illustrate the effectiveness of using the Markov chain modeling for high-flux acquisition. Using a 3-bit grayscale reflectivity approximation, our MCPDF method outperforms both low-and high-flux depth estimate for the number of illuminations $n_{\mathrm{r}} = 100$ and 2000. Our MCPDF method approximately matches the LF performance with 20× fewer illuminations, greatly speeding up acquisition.

$f_{T_{\ell+1}|T_\ell}(kt_{\mathrm{r}} + b_n|b_m)$ defined in (4) on the finite interval $[0, Kt_{\mathrm{r}}]$. To make P a transition probability matrix, we normalize P such that it has row sum equal to 1. Finally, the approximated limiting distribution of detection times $\mathbf{f}^{\mathrm{D}}$, or equivalently the stationary distribution of the Markov chain, is the leading left eigenvector of P. Fig. 1 shows that such an approximation is sufficiently accurate, as it closely matches the simulated histogram in all tested (S, B) pairs.

3. ESTIMATION ALGORITHM

Periodic, time-quantized acquisition results in a histogram $\mathbf{h} = [h_1, \ldots, h_{n_{\mathrm{b}}}]$ of detection counts in each time bin. For a Poisson process whose intensity $\lambda(t)$ is known up to a delay, the maximum likelihood (ML) estimate of the delay is the shift that maximizes the output of the log-matched filter, i.e., the filter with impulse response $\log[\lambda(t)]$ [25]. For an acquisition made at low flux, the log-matched filter can correctly use the photon arrival intensity because the effects of dead time can be considered negligible. If used for high-flux acquisition, however, a filter matched to the arrival intensity is suboptimal because of the mismatch between the pulse shape and the actual distribution of detection times. In [19], the authors outlined a number of depth estimation strategies, several of which we will test here. For each method, the histogram of detection counts is circularly convolved with a time-reversed detection time PDF:

$$\widehat{m} = \arg\max_m \mathbf{h}_m \circledast \log(\mathbf{f}_{-m}), \tag{5}$$

where the negatively-indexed subscript denotes time-reversal. Conversion to distance units is the straightforward transformation $\widehat{z} = (\widehat{m} - \frac{1}{2})\Delta c/2$.

Two separate acquisition types are proposed: a *low-flux* acquisition resulting in histogram $\mathbf{h}^L$ first attenuates the photon flux (e.g., with a neutral density filter) so the average photon arrival rate for the scene is 0.05 and the effect of dead time is minimal; a *high-flux* acquisition applies no attenuation and yields a detection histogram $\mathbf{h}^H$. We also consider two possible filters, with $\mathbf{f}^{\mathrm{A}}$ a discretization of the arrival PDF $f^A(t)$ and the detection PDF $\mathbf{f}^{\mathrm{D}}$ as computed in Section 2. The methods combine acquisitions and filters as follows: the conventional method LF uses histogram $\mathbf{h}^{\mathrm{L}}$ and filter $\mathbf{f}^{\mathrm{A}}$, the naïve high-flux method HF uses $\mathbf{h}^{\mathrm{H}}$ and $\mathbf{f}^{\mathrm{A}}$, and our proposed Markov chain-based method MCPDF uses $\mathbf{h}^{\mathrm{H}}$ and $\mathbf{f}^{\mathrm{D}}$.

One crucial component in the accurate computation of $\mathbf{f}^{\mathrm{D}}$ is the requirement that $\alpha_{i,j}$ is known. Estimating $\alpha_{i,j}$ directly from the photon detections is difficult, as the coupled relationship between the depth and reflectivity parameters is further complicated by the dead time. Instead, since the target application of our depth imaging system is autonomous navigation, other sensing modalities are likely available that could supplement the lidar with information relevant to the reflectivity. We assume a conventional camera coaxially aligned with the lidar and spectrally filtered to accept the same wavelength can acquire a grayscale image that is a sufficient approximation of $\boldsymbol{\alpha}$. Then the camera image $\widetilde{\boldsymbol{\alpha}}$ acquired simultaneously with the lidar data can be used in the depth image reconstruction. Since recomputation of $\mathbf{f}^{\mathrm{D}}$ with each new $\widetilde{\alpha}_{i,j}$ would be a slow and inefficient process, we instead precompute $\widetilde{\mathbf{f}}^{\mathrm{D}}$ for a small set n_{q} of evenly spaced values of α over $[0, 1]$. Then for each (i, j) the $\widetilde{\mathbf{f}}^{\mathrm{D}}$ for the closest value to $\widetilde{\alpha}_{i,j}$ is used in the log-matched filter to estimate depth.

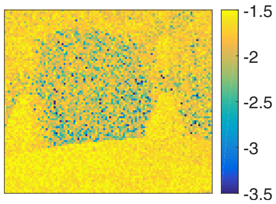

(a) HF

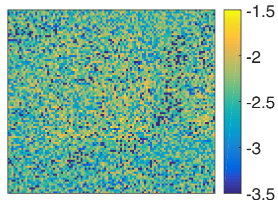

(b) MCPDF

Fig. 3: Comparison of $\log_{10}(|\hat{\boldsymbol{z}} - \boldsymbol{z}|)$ for HF and MCPDF for the high-flux acquisition with $n_{\mathrm{r}} = 100$ illuminations. The error for HF in (a) is lower for darker scene patches with detection PDFs less distorted by dead time, whereas the error for MCPDF in (b) is lower for lighter scene patches which reflect back more signal photons.

4. SIMULATION RESULTS

To validate our depth estimation algorithm, we simulate detection data using ground truth depth and reflectivity images from the Middlebury stereo dataset [26]. The color images are first converted to grayscale and then normalized so that $\alpha_{i,j} \in [0.1, 1.0]$. The disparity image is converted to a depth map using intrinsic camera properties, and the scene is arbitrarily shifted by 10 m (66.7 ns). Both images are downsampled to 93×105 pixels to reduce processing time. For all simulations, we use parameter values of $\beta = 6$, $B = 3$, $\sigma_{\mathrm{p}} = 0.2$ ns, $\Delta = 0.02$ ns, $t_{\mathrm{r}} = 100$ ns, and $t_{\mathrm{d}} = 75$ ns.

Fig. 2 shows the results of simulated acquisitions and depth estimation for the Bowling scene with the number of illuminations $n_{\mathrm{r}} = 100$ and 2000. Reflectivity measurement by a 3-bit grayscale camera is emulated. Low-flux acquisition results in 4.80 and 96.3 detected photons per pixel (ppp) over the scene for the short and long acquisitions, respectively. Without first attenuating the photon flux, the high-flux acquisition detects photons much faster, with an average rate of just more than one photon detection per illumination. The depth estimation results for the short acquisition demonstrate why the LF approach is insufficient for real-time applications—there are simply too few signal detections to reliably estimate the depth. Increasing n_{r} enables an improvement in root mean squared error (RMSE) of several orders of magnitude for the LF method, with RMSE computed as

$$\mathrm{RMSE}(\widehat{\boldsymbol{z}}) = \sqrt{\frac{1}{n_i n_j}\sum_{i}^{n_i}\sum_{j}^{n_j}(z_{i,j} - \widehat{z}_{i,j})^2}. \tag{6}$$

On the other hand, the results from the HF estimator barely improve as n_{r} increases since it does not take dead time into account, so the error is dominated by a bias. The RMSE of our proposed Markov chain-based MCPDF method continues to decrease as more data is acquired. Furthermore, MCPDF achieves nearly the same RMSE for the short acquisition as LF does for the long acquisition, enabling accurate depth imaging to be performed 20 times faster.

A comparison for the high-flux acquisition in Fig. 3 further illustrates the advantage of MCPDF. The absolute error for HF and MCPDF is shown on a logarithmic scale. The figure reveals that when the overall arrival rate is high enough, the smallest errors for HF somewhat counterintuitively occur for the darkest pixels, since their detection PDFs are least distorted by the dead time. On the other hand, by correctly anticipating the dead time distortion, errors for MCPDF occur roughly proportionally with the number of detected photons.

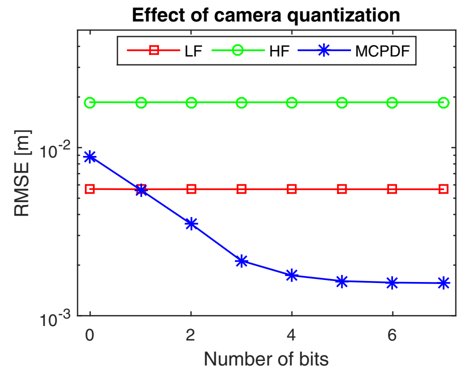

Fig. 4: The performance of MCPDF improves as more quantization levels are used for the reflectivity estimate. The plot shows the median RMSE values for 100 realizations of detection data for the Bowling scene with $n_{\mathrm{r}} = 2000$ illuminations.

We also explore the effect of the quantization of $\widetilde{\boldsymbol{\alpha}}$ on the reconstruction error. For the same experimental parameters as previously used, 100 realizations of photon detection data were generated for different numbers of bits for the Bowling scene reflectivity. Fig. 4 shows the median RMSE results over the 100 trials. The median is plotted since significant outliers occasionally occurred when the RMSE for the LF method was dominated by a small number of pixels with large depth errors. It is clear from the plot that the performance of MCPDF greatly improves with the number of quantization levels, whereas the RMSE of the LF and HF methods does not depend on information about the reflectivity. Fig. 1 helps illustrate why the methods depend differently on the quantization of $\widetilde{\boldsymbol{\alpha}}$. For the LF and HF methods, which use $\mathbf{f}^{\mathrm{A}}$, a change in the estimated value of $\alpha_{i,j}$ only changes the strength of the signal relative to the background, but the position of the peak is unaffected, so the log-matched filter depth estimates are mostly unchanged. On the other hand, the shape of $\mathbf{f}^{\mathrm{D}}$, including the position of its peak, depends strongly on the exact $\alpha_{i,j}$ value, so a closer approximation of $\alpha_{i,j}$ from finer quantization yields a more accurate approximation of $\mathbf{f}^{\mathrm{D}}$ and thus a better log-matched filter depth estimate. Moreover, Fig. 4 implies that for the Bowling scene used in our experiment, 3 bits are sufficient to achieve small performance degradation due to reflectivity quantization, though the number of bits needed is likely affected by the range of $\alpha_{i,j}$ values.

5. CONCLUSION

Real-time depth imaging is a critical component of navigation systems for autonomous vehicles. In this work, we demonstrated that precise modeling of the effects of dead time on photon-counting lidar systems enables fast and accurate depth image acquisition. A simple, discrete-time approximation to the PDF of the next photon arrival time after a detector reset leads to efficient computation of the limiting distribution for the Markov chain of photon detection times. A grayscale camera is used to capture the approximate scene reflectivity in order to choose the correct limiting distribution for each pixel. Results show our method can achieve accurate depth images 20 times faster than the conventional method. Future work could incorporate spatial information such as in [23, 27], allowing for accurate imaging with even lower photon counts and further reducing acquisition times.

6. REFERENCES

[1] D. J. Fagnant and K. Kockelman, "Preparing a nation for autonomous vehicles: Opportunities, barriers and policy recommendations," *Transp. Res. Part A Policy Pract.*, vol. 77, pp. 167–181, July 2015.

[2] A. P. Garza, "Look Ma, no hands: Wrinkles and wrecks in the age of autonomous vehicles," *New Engl. Law Rev.*, vol. 46, 2011.

[3] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Z. Kolter, D. Langer, O. Pink, V. Pratt, M. Sokolsky, G. Stanek, D. Stavens, A. Teichman, M. Werling, and S. Thrun, "Towards fully autonomous driving: Systems and algorithms," in *IEEE Intell. Veh. Symp. Proc.*, June 2011, pp. 163–168.

[4] T. Schamm, C. Von Carlowitz, and J. Marius Zöllner, "On-road vehicle detection during dusk and at night," in *IEEE Intell. Veh. Symp. Proc.*, June 2010, pp. 418–423.

[5] J. Wei, J. M. Snider, J. Kim, J. M. Dolan, R. Rajkumar, and B. Litkouhi, "Towards a viable autonomous driving research platform," in *IEEE Intell. Veh. Symp. Proc.*, June 2013, pp. 763–770.

[6] K. Bengler, K. Dietmayer, B. Farber, M. Maurer, C. Stiller, and H. Winner, "Three decades of driver assistance systems: Review and future perspectives," *IEEE Intell. Transp. Syst. Mag.*, vol. 6, no. 4, pp. 6–22, 2014.

[7] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.

[8] 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. Imag.*, vol. 1, no. 2, pp. 112–125, June 2015.

[9] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform-based analysis of depth images constructed using sparse single-photon data," *IEEE Trans. Image Process.*, vol. 25, no. 5, pp. 1935–1946, May 2016.

[10] M. A. Itzler, M. Entwistle, S. Wilton, I. Kudryashov, J. Kotelnikov, X. Jiang, B. Piccione, M. Owens, and S. Rangwala, "Geiger-mode LiDAR: From airborne platforms to driverless cars," in *Imaging Appl. Opt.*, Washington, D.C., June 2017, p. ATu3A.3, OSA.

[11] D. V. O'Connor and D. Phillips, *Time-Correlated Single Photon Counting*, Academic Press, Orlando, FL, 1984.

[12] W. Becker, *Advanced Time-Correlated Single Photon Counting Techniques*, Springer, Berlin, 2005.

[13] P. B. Coates, "The correction for photon 'pile-up' in the measurement of radiative lifetimes," *J. Phys. E: Sci. Instrum.*, vol. 1, no. 8, pp. 878–879, Aug. 1968.

[14] C. C. Davis and T. A. King, "Photon pile-up corrections in the study of time-varying light sources," *J. Phys. E: Sci. Instrum.*, vol. 5, no. 11, pp. 1072–1074, 1972.

[15] G. Verma and R. J. Drost, "Inhomogeneous Poisson process rate function inference from dead-time limited observations," *J. Opt. Soc. Amer. A.*, vol. 34, no. 5, pp. 770, May 2017.

[16] A. K. Pediredla, A. C. Sankaranarayanan, M. Buttafava, A. Tosi, and A. Veeraraghavan, "Signal processing based pile-up compensation for gated single-photon avalanche diodes," arXiv:1806.07437, June 2018.

[17] F. Heide, S. Diamond, D. B. Lindell, and G. Wetzstein, "Sub-picosecond photon-efficient 3D imaging using single-photon sensors," *Scientific Reports*, vol. 8, no. 1, pp. 17726, Dec. 2018.

[18] S. Isbaner, N. Karedla, D. Ruhlandt, S. C. Stein, A. Chizhik, I. Gregor, and J. Enderlein, "Dead-time correction of fluorescence lifetime measurements and fluorescence lifetime imaging," *Opt. Expr.*, vol. 24, no. 9, pp. 9429, 2016.

[19] J. Rapp, Y. Ma, R. M. A. Dawson, and V. K. Goyal, "Dead time compensation for high-flux ranging," arXiv:1810.11145, Oct. 2018.

[20] U. S. Kamilov and P. T. Boufounos, "Motion-adaptive depth superresolution," *IEEE Trans. Image Process.*, vol. 26, no. 4, pp. 1723–1731, Apr. 2017.

[21] H. Mahmoudabadi, T. Shoaf, and M. J. Olsen, "Superpixel clustering and planar fit segmentation of 3D LIDAR point clouds," in *Proc. 4th Int. Conf. Comput. Geospatial Res. Appl.*, July 2013, pp. 1–7.

[22] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3D object detection network for autonomous driving," *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, pp. 6526–6534, July 2017.

[23] D. B. Lindell, M. O'Toole, and G. Wetzstein, "Single-photon 3D imaging with deep sensor fusion," *ACM Trans. Graph.*, vol. 37, no. 12, 2018.

[24] D. L. Snyder and M. I. Miller, *Random Point Processes in Time and Space*, Springer Science & Business Media, 2012.

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

[26] D. Scharstein and C. Pal, "Learning conditional random fields for stereo," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2007.

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