Count-Free Histograms with Race Logic for Single-Photon LiDAR

Atul Ingle David Maier
Portland State University

***Abstract*—Low-power 3D perception is useful in a wide range of computer-vision applications. Thanks to the increasing availability of high-resolution single-photon avalanche diode (SPAD) arrays, single-photon LiDARs (SPLs) have emerged as a promising technology for 3D sensing. The conventional image formation model for an SPL involves capturing the time-varying light intensity—which we call the *transient distribution*—of a reflected laser pulse in the form of an *equi-width* (EW) histogram. Unfortunately, this approach leads to unmanageable data rates (∼gigabytes/second) with high-resolution arrays, severely limiting the applicability of SPLs in power- and bandwidth-constrained scenarios (e.g., mobile devices). We propose a radically different approach based on race logic processing to construct equi-depth histograms with variable bin widths. This method avoids storing high-resolution histogram counts, thereby reducing the bandwidth requirement while maintaining similar ranging accuracy. We show simulation results with bandwidth reduction of over $100\times$.**

***Index Terms*—LiDAR, 3D imaging, SPAD, compression, race logic**

I. INTRODUCTION

Single-photon sensing is a promising technology for high-resolution 3D imaging. Low-power 3D perception is useful in a range of computer-vision applications, including industrial robotics, autonomous driving and augmented reality. Image sensors that can capture single photons, such as single-photon avalanche diodes (SPADs), are popular as detectors for such applications. High (kilo-to-megapixel) resolution SPAD arrays with additional data-processing embedded in the same hardware are increasingly available. However, their high sensitivity and speed is a double-edged sword: the data generated by such arrays greatly exceeds what can be reasonably processed or transferred in real time, limiting their applicability, especially where there are power and bandwidth constraints. A single-photon camera (SPC) captures depth information using the time-of-flight principle [4]. A laser illuminates the scene with a short light pulse. The corresponding camera pixel captures a stream of return events as photons arrive at different delays with respect to the pulse time. The aggregation of these delays over many laser cycles forms a transient distribution that we wish to capture. Traditional methods construct histograms of 1000's of bins by storing event counts at different delays over multiple laser cycles. The peak of a histogram gives an estimate of the true distance of the scene point (by the relationship that speed of light × delay = twice the distance to the scene point).

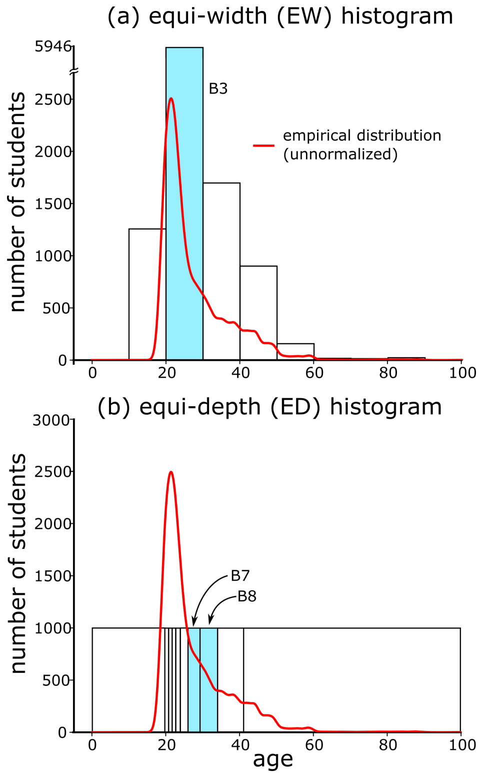

Fig. 1. **Advantage of equi-depth (ED) histogram over conventional equi-width (EW) histogram for peaky distributions.** This figures shows two types of histograms for student-age data from Dale et al. [2]: (a) A 10-bin EW histogram has a single bin B3 near the peak. Many bins are close to zero and provide no useful information about the peak location. (b) A 10-bin ED histogram with approximately 1000 students per bin reliably captures the shape of the peak. In this work we apply this intuition to SPL data: ED histograms adaptively cluster around the transient-distribution peak, providing accurate distance information with only a few ED bins.

II. KEY IDEAS

We want to estimate scene distances using an SPC with minimal power and data bandwidth. This goal rules out forming 1000-bin histograms on the image sensor, which requires large storage on chip and high bandwidth during readout of bin counts. Conventional histogram-based methods are especially inefficient with strong background illumination, as

---

Author emails: {`ingle2`, `maier`}`@pdx.edu`. Mailing address: 1900 SW 4th Ave Ste 120, Portland OR 97201. Phone: +1-503-725-3234. Research supported in part by NSF ERI Award 2138471.

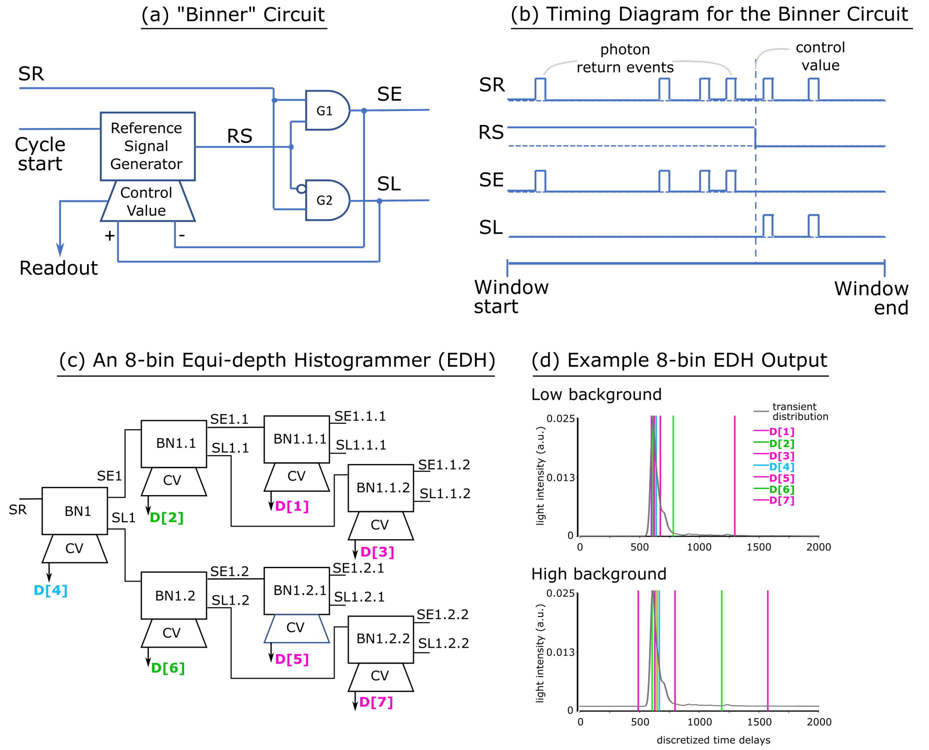

Fig. 2. **The proposed binner circuit and equi-depth histogrammer (EDH).** (a) The binner circuit splits the incoming photon stream (SR) into an early stream (SE) and a late stream (SL) depending on transition point of a reference signal (RS) generated from a control value. (b) In this example, there are more photons in the early stream than the late stream, so the control value will decrease for the subsequent laser cycle, thus moving the transition point of RS earlier. The control value eventually settles close to the overall median. (c) An 8-bin ED histogram can be captured using a collection of 7 binners arranged in a 3-level binary tree. A binner at one stage feeds streams of early and late photon events to two binners at the next stage in the tree. (d) This example shows a transient distribution and the simulated results of an 8-bin EDH for low and high background levels. Notice that a majority of the bins cluster around the true peak location. The location of the narrowest bin provides a reliable estimate of scene distance.

most bins count background photons. We offer a radically different approach for time-of-flight imaging compatible with multiple detector and illumination schemes. It has two key elements: constructing equi-depth (ED) histograms rather than the traditional equi-width (EW) histograms, thereby using many fewer bins to approximate the transient distribution and using race logic to process information in the “delay domain”, avoiding conversion of return events to digital timestamps.

Equi-Depth Histograms: ED histograms have *variable-width* bins each with approximately the same number of items. Fig. 1 shows a conceptual example of 10-bin EW and ED histograms for a dataset of 10,000 Canadian college students [2]. The majority of this population is in the bin labeled “B3” of the EW histogram, while some other bins are empty, contributing little information. In contrast, a 10-bin ED histogram for the same distribution does a better job of capturing the shape of the peak. In case of an SPL, ED histograms can provide reliable estimates of scene distance with very few bins that cluster around the peak of the distribution.

Race Logic: Race logic [5] is a novel approach to computation in which values are represented as time delays, rather than as analog or digital quantities.

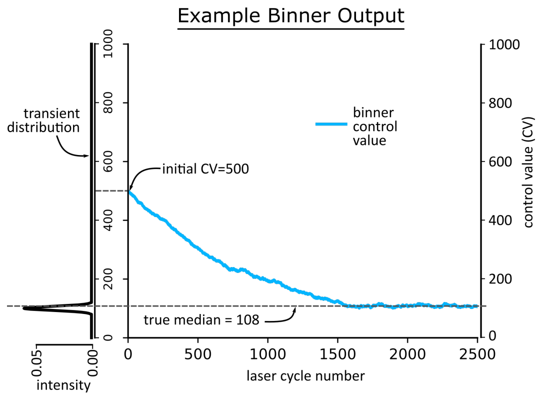

Fig. 3. **Example trajectory of the control value of a binner.** We show a binner’s CV over multiple laser cycles. The true peak is at 100, signal strength is 1.0 and signal-to-background ratio is 1.0.

Race logic is well suited to SPAD data processing, as the photon arrivals are already in the delay domain. Working with information in this form avoids the energy required to convert return events to digital timestamps, which most existing methods need.

III. METHODS

Fig. 2(a) shows our proposed binner circuit. “Early” or “late” photons relative to the current estimate (called

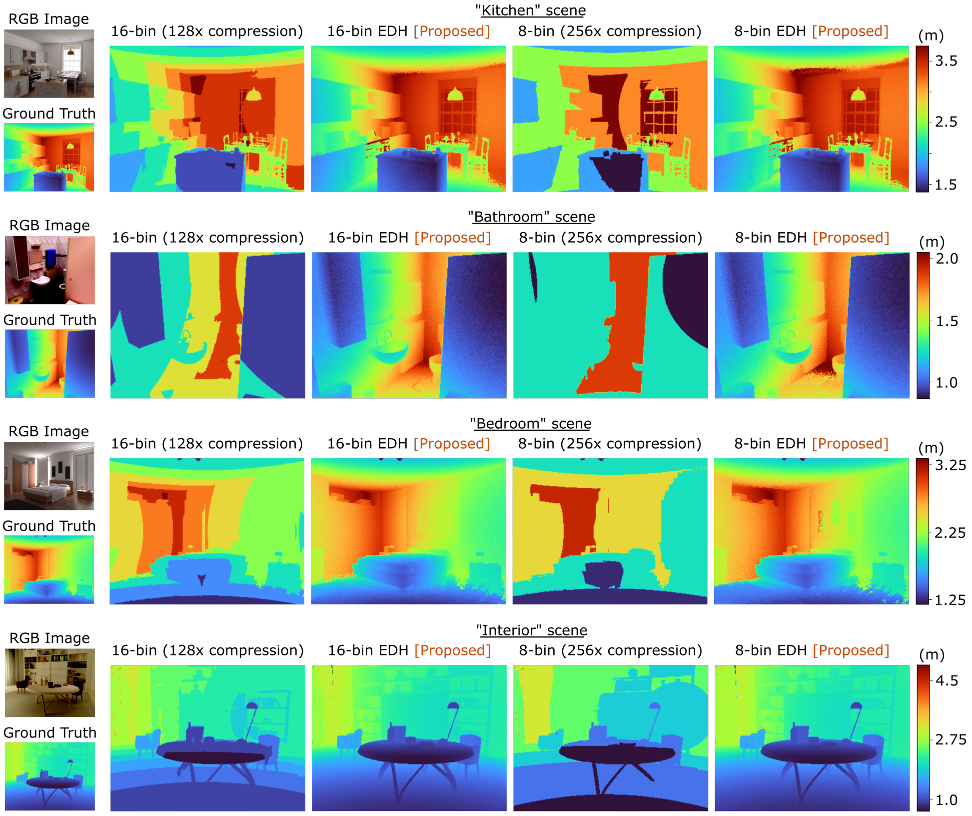

Fig. 4. **Simulated results using a transient rendering dataset [3].** This figure shows results for “Kitchen”, “Bathroom”, “Bedroom” and “Interior” scenes from a simulated transient-rendering dataset. RGB and ground-truth distance maps are shown alongside distance map reconstructions using conventional 8-bin and 16-bin coarse EW histograms. Observe the lack of details and increased quantization artifacts in distance maps computed using the compressed coarse EW histogram method. In contrast, our method reliably captures scene distances with as few as 8 ED bins while achieving over 100× compression.

the “control value” (CV)) move that estimate earlier or later, respectively. Photon-timing histograms are never explicitly stored in memory (unlike equi-width-histogram approaches that remember the timestamp of every photon received). Intuitively, the control value must settle at a location where equal number of photons arrive in the early and late streams. By definition, this point is the median of the transient distribution. Note, however, that the binner does not converge deterministically to the median because photon arrivals are random and the binner does not maintain the full history of photon arrivals. Fig. 3 shows the CV of a simulated binner over multiple cycles, for a synthetic transient distribution with a single, narrow peak, along with the true median of the distribution.

The median estimate will often lie close to the peak in the transient distribution, but can deviate significantly, if there is significant background light or the peak is asymmetric. To deal with such divergence, we cascade multiple binners stages to form an equi-depth histogrammer (EDH). The collective CVs of these binners represent the ED histogram bin boundaries. A binner at one stage feeds streams of early and late return events to two binners at the subsequent stage, each which in turn subdivides the corresponding bin from the previous stage. For example, a three-stage EDH shown in Fig. 2(c) will compute the boundaries of a 8-bin ED histogram. With multiple bins, the bin boundaries cluster around the transient-distribution peaks. Fig. 2(d) shows an example simulated output for a transient distribution with a single peak and some background illumination. Having adaptive bin boundaries helps an EDH cope with background light. The lower part of Fig. 2(d) uses the same simulated transient distribution as the upper part, but with more background light. We see a couple more bins away from the peak, which essentially “absorb” the background photons. However, there are still sufficient bins remaining to capture laser peak well. Since EDH bin boundaries cluster around the peak of the transient distribution, scene distances can be estimated by locating the narrowest bin. We are evaluating two methods for estimating the transient peak from an ED histogram. The simpler method just returns the midpoint of the narrowest bin. That estimate can be inaccurate when two bins split the peak. Thus, we also have a more sophisticated method that fits a curve to the bins in the neighborhood of the narrowest bin, then uses the argmax of that curve as the peak estimate.

IV. RESULTS

We evaluate the performance of an EDH using a transient rendering dataset [3]. We simulate photon-

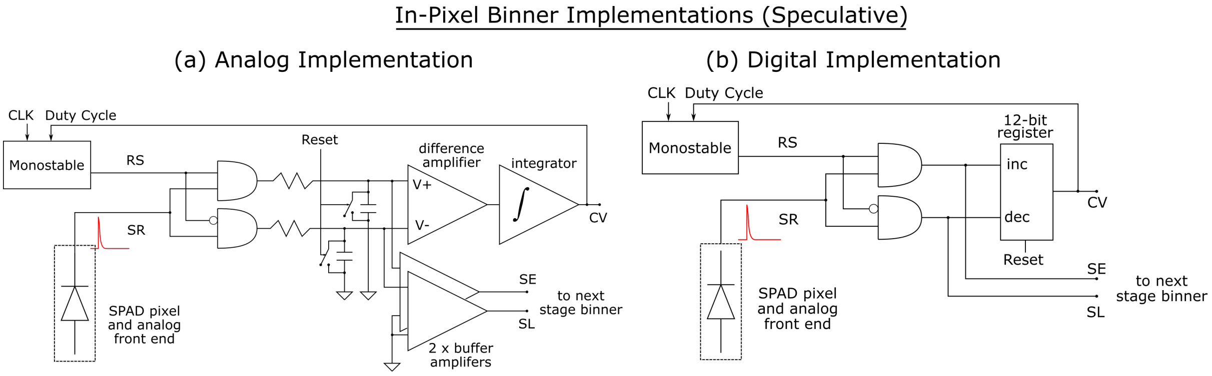

Fig. 5. **Speculative pixel architectures for an in-pixel binner implementation.** (a) An analog implementation accumulates the early and late stream signals as charges on two capacitors that are compared to increase/decrease the control value. (b) In the digital implementation, an up/down register holds the control value.

return streams using the pixelwise ground-truth transient distributions from this dataset. Our simulation accounts for the effect of signal and ambient light strength, different scene albedoes, and sensor-noise sources including shot noise and dark counts.

Fig. 4 shows simulated transient-rendering results using 8- and 16-bin EDHs for four different scenes. Observe that in the "kitchen" scene (first row), our method preserves distance gradients of flat surfaces (such as the walls of the room) that appear quantized with a coarse 8- or 16-bin EWH. In the "bedroom" scene, an EDH reliably captures fine distance details such as the bed-frame with as few as 8 bins. In the "interior" scene, a coarse EWH loses details such as small objects on the table and in the shelf in the background. Our method preserves these distance details with as few as 8-bins. Unlike coarse EW histograms, our method avoids strong quantization artifacts and preserves details in the distance maps while reducing the data rate by $> 100\times$

V. CONCLUSION AND FUTURE DIRECTIONS

We have presented a method for distance imaging that is a marked departure from current approaches. Instead of capturing the transient distribution by storing photon count histograms, we instead use a binner element that maintains and adjusts the distribution-median estimate at every laser-pulse cycle. Using a cascade of binners, we can produce equi-depth histogrammers that robustly capture peaks in the distribution. Our approach requires many fewer bins than with the equi-width approach, thus reducing memory and bandwidth requirements. Moreover, part or all of a binner element can operate in the delay domain using race logic, avoiding time-to-digital conversion, thus lowering circuitry and energy requirements.

Current single-photon cameras are severely bandwidth-constrained due to the requirement of reading out individual photon timestamps at extremely high (mega-to-gigahertz) rates. ED histograms captured using race logic processing have the potential to reduce this bandwidth requirement by orders of magnitude, enabling future SPAD arrays to be scaled to higher spatial resolution. Lower bandwidth requirements can also provide power savings by requiring fewer bits to be moved off sensor, and may simplify in-pixel circuitry by avoiding the need to construct and store high-bit-depth histograms on-sensor.

Our current work is proceeding along two tracks, method refinement and hardware prototyping.

Method Refinement: We are exploring variations and extensions of the EDH approach that improve convergence rates and accuracy, as well as conducting further simulations and analysis to understand the effects of signal strength, background light and peak position on binner and EDH behavior.

Hardware prototyping: We are currently prototyping a binner circuit using a SPAD array and associated FPGA [1], and plan to extend that to a full EDH. Based on what we learn from that effort, we will explore designs that integrate the binner onto the acutal detector device. Fig 5 shows two (speculative) designs with a binner implemented in-pixel.

REFERENCES

[1] Claudio Bruschini, Samuel Burri, Ermanno Bernasconi, Tommaso Milanese, Arin C Ulku, Harald Homulle, and Edoardo Charbon. Linospad2: a 512x1 linear spad camera with system-level 135-ps sptr and a reconfigurable computational engine for time-resolved single-photon imaging. In *Quantum Sensing and Nano Electronics and Photonics XIX*, volume 12430, pages 126–135. SPIE, 2023. 4

[2] Meghan Dale. Trends in the Age Composition of College and University Students and Graduates (Archived). *Education Matters: Insights on Education, Learning and Training in Canada*, (5), Dec 2010. Online (accessed July 21, 2022) https://www150.statcan.gc.ca/n1/en/pub/81-004-x/2010005/article/11386-eng.pdf. 1, 2

[3] Felipe Gutierrez-Barragan, Huaijin Chen, Mohit Gupta, Andreas Velten, and Jinwei Gu. itof2dtof: A robust and flexible representation for data-driven time-of-flight imaging. *IEEE Transactions on Computational Imaging*, 7:1205–1214, 2021. 3

[4] Istvan Gyongy, Neale AW Dutton, and Robert K Henderson. Direct time-of-flight single-photon imaging. *IEEE Transactions on Electron Devices*, 69(6):2794–2805, 2021. 1

[5] Advait Madhavan, Timothy Sherwood, and Dmitri Strukov. Race logic: A hardware acceleration for dynamic programming algorithms. *ACM SIGARCH Computer Architecture News*, 42(3):517–528, 2014. 2