

A direct time-of-flight image sensor with in-pixel surface detection and dynamic vision

Istvan Gyongy, Ahmet T. Erdogan, Neale A.W. Dutton, *Member, IEEE* Germán Mora Martín, Alistair Gorman, Hanning Mai, Francesco Mattioli Della Rocca, *Member, IEEE* and Robert K. Henderson, *Fellow, IEEE*

Abstract—3D flash LIDAR is an alternative to the traditional scanning LIDAR systems, promising precise depth imaging in a compact form factor, and free of moving parts, for applications such as self-driving cars, robotics and augmented reality (AR). Typically implemented using single-photon, direct time-of-flight (dToF) receivers in image sensor format, the operation of the devices can be hindered by the large number of photon events needing to be processed and compressed in outdoor scenarios, limiting frame rates and scalability to larger arrays. We here present a 64×32 pixel (256×128 SPAD) dToF imager that overcomes these limitations by using pixels with embedded histogramming, which lock onto and track the return signal. This reduces the size of output data frames considerably, enabling maximum frame rates in the 10 kFPS range or 100 kFPS for direct depth readings. The sensor offers selective readout of pixels detecting surfaces, or those sensing motion, akin to dynamic vision sensors that report changes in light intensity, leading to reduced power consumption and off-chip processing requirements. We demonstrate the application of the sensor in mid-range LIDAR.

I. INTRODUCTION

DIRECT time-of-flight (dToF) sensors measure depth by illuminating the scene of interest with a pulsed laser source, and timing the back-scattered return signal. The sensors are capable of centimetre precision over hundreds of meters and beyond, and as such are a key technology for LIDAR-based 3D vision systems providing situational awareness in self-driving cars and other autonomous systems [1], [2].

With the use of single-photon avalanche diodes (SPADs) in complementary metal-oxide semiconductor (CMOS) technology, detector arrays with integrated processing can be realised, for a single-chip, all-digital receiver unit. Combining a SPAD receiver, implemented in image sensor format, with flood illumination provides a solid-state 3D imaging solution, with increased robustness to vibrations and reduced calibration needs compared with approaches involving mechanical scanning [3]. Furthermore, by capturing the entire field of view simultaneously, motion artefacts resulting from dynamic scenes are minimised.

The first two authors contributed equally to this work

I. Gyongy, A.T. Erdogan, G. Mora Martin, A. Gorman, H. Mai and R.K. Henderson are with the School of Engineering, Institute for Integrated Micro and Nano Systems, The University of Edinburgh, Edinburgh, EH9 3FF, UK (istvan.gyongy@ed.ac.uk, ahmet.erdogan@ed.ac.uk, german.mora@ed.ac.uk, h.mai@ed.ac.uk, agorman@ed.ac.uk, robert.henderson@ed.ac.uk).

N.A.W. Dutton is with STMicroelectronics Imaging Division, Tanfield, Edinburgh, EH3 5DA, UK (neale.dutton@st.com)

F. Mattioli Della Rocca was with The University of Edinburgh and is now with Sony Europe Technology Development Centre, 38123 Trento, Italy (francescopaolo.mattiolidellarocca@sony.com)

The premise of SPADs in 3D flash LIDAR has led to significant research efforts of late, aiming to overcome the key challenges in developing viable, array format sensors for mid- to long range outdoor applications [4], [5], [6], [7], [8], [9], [10], [11], [12]. These challenges include equipping the sensor with a sufficient photon throughput to accommodate the maximal data rate of several tens of millions of events per second that can be generated by each SPAD. A second, related problem is ensuring an appropriate level of on-chip data compression to overcome readout bottleneck issues (as well as potential high memory requirements) and enable depth imaging at video rates and beyond. The overall sensor architecture should ideally be scalable to large array sizes without incurring excessive power consumption.

Recent reviews of SPAD architectures for dToF/LIDAR include [13], [14], [15]. In view of the high photon throughput and data compression requirements, one of the key trends has been embedded processing in the form of histogramming of the time-of-arrival of photons, with some architectures achieving higher still compression by processing these histograms to compute depth maps on-chip. The histogram processing can be carried out in-pixel, or outside the pixel array. In the latter case, there is more area available for histogram memory, but transferring data out of the pixel array can become challenging for large arrays. In either case, the accumulation of full histograms, covering the entire distance range of interest, can be impractical, not just because of the large memory requirements, but due to the increased power consumption and potential bottlenecks associated with the transfer of large amounts of data in and out of memory.

It has therefore become common to apply multi-step, partial histogramming approaches instead, which typically use just 16 or 32 bins. The approaches, which can be categorised into “zooming” and “sliding” (Fig. 1), are contrasted in detail in [16]. The zooming method starts by histogramming the whole time range (using a limited number of bins), and then in each step the peak bin is identified, and the histogramming logic zooms in on the corresponding time range, homing in on the signal peak. In the sliding approach, the partial histogram is swept through the full time range, covering the entire time range over multiple steps.

As demonstrated in [16], the sliding scheme is more effective than zooming over longer ranges and under high ambient levels, due to a large build up of background counts in the initial steps of zooming (resulting from the wide bin width), making it difficult to detect the signal peak reliably, unless very high laser powers are used (in combination with a suitably high

histogram bit depth).

The present chip uses the sliding approach, but rather than continually scanning the whole time range, we designed pixels that lock onto peaks and track them. As a result, there is an increase in the effective temporal aperture for signal photons, and once a pixel has homed in a peak, the remaining time range is ignored, precluding the possibility of false peak detection in the excluded range. Furthermore, the approach also provides a mechanism for detecting motion in the observed scene. Read out options include a modality where only pixels detecting a change in depth are reported, demonstrating a pathway towards 3D imaging sensors that not only produce point clouds, but identify spatio-temporal features in the data to reduce data rates, saving power and facilitating low latency vision.




Fig. 1. Partial histogramming techniques in single-photon direct time-of-flight imaging. In each step of the zooming approach (top plot), the peak bin is identified, and the histogramming logic zooms into the corresponding time range. The sliding approach (middle plot) sweeps through the entire time range in multiple steps. Unlike the zooming approach, the histogram bin width remains fixed. The peak tracking scheme (bottom plot) further develops the sliding concept by estimating the background (baseline photon) level, and detecting signal peaks (marked in red), which are subsequently tracked. The dashed yellow curves correspond to the firing of the laser pulse. Adapted from [15].

II. CHIP ARCHITECTURE

The imager is implemented in a standard 40 nm process and features a 64×32 array of pixels (Fig. 2). Each pixel is composed of 4×4 SPADs plus a processing unit, which occupy matching areas (of approximately $57 \mu\text{m} \times 54 \mu\text{m}$), to enable a potential future 3D stacked implementation [17], where the processing units are placed underneath the detector array for increased pixel density and fill factor. The SPADs are passively quenched, and are based on a design, including front-end, similar to [18]; thus similar dead-time, jitter and a photon detection probability (PDP) of around 4% at 850nm is expected.

When depth imaging, each pixel generates an 8-bin photon timing histogram that is automatically shifted in time to locate and track peaks, via an electronic time gate that is controlled by the pixel itself. There is additional, and optional, processing outside the array for computing depth values from the histograms with sub-bin resolution using a centre-of-mass (CMM) technique.

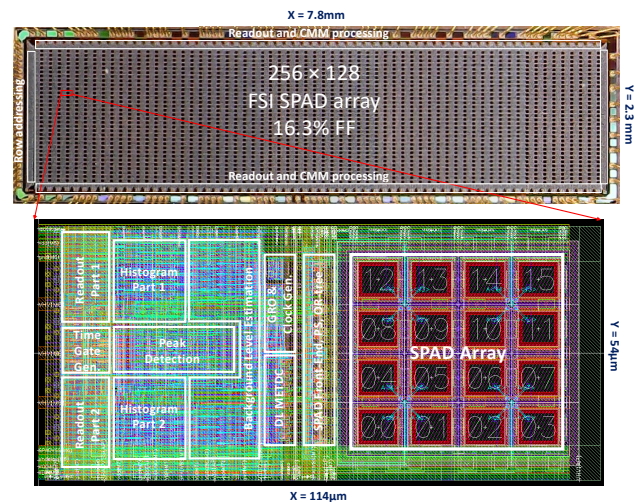


Fig. 2. Micrograph of the sensor and a close-up of the layout of an individual pixel, showing the 4×4 array of SPADs, and the processing unit next to it. The main constituent processing blocks are framed in white. The chip was implemented in a front side illuminated (FSI) architecture and has an overall fill factor (FF) of 16.3%.

The sensor can also be operated in an intensity imaging modality, where SPADs are grouped in pairs, with each pair reporting a 12-bit photon count for an overall image resolution of 128×128 .

Data is read out from the sensor over 64, 100 MHz output lines, resulting in maximal frame rates in the 10 kFPS or 100 kFPS range, depending on the mode of operation.

Pixel architecture

A block diagram of the pixel is shown in Fig. 3. The outputs of the 16 SPADs are fed to pulse shorteners, which produce digital pulses of width ≈ 300 ps for each photon event. In the photon timing modes, the SPAD outputs are then combined using an OR tree and sent to a multi-event time-to-digital converter (TDC) [19]. The TDC is able to register multiple photons per laser cycle, even if they fall within the same timing bin (provided that the events are sufficient far apart in time to form distinct digital pulses, in other words more than around 300 ps apart according to extracted simulations). The output of the TDC goes to 8, 12-bit counters representing the histogram bins. At the end of every exposure, the contents of these counters are processed to estimate the background photon level, to detect statistically significant peaks and to move the time gate position accordingly. If a peak is detected in bin 1 or 2, the time gate position is decremented, if the peak is in bin 7 or 8, or no peak is detected, the time gate is incremented. Otherwise (if the peak is in the middle, so in bins 3 to 6), the time gate position is maintained.

Fig. 1 illustrates the peak tracking functionality using histograms captured over four consecutive exposures. The pixel is initially at time gate position 1, and the ambient level is estimated at this position. Based on this ambient level, a peak detection threshold is computed according to Poisson statistics. Next, the time gate is shifted to position 2. In position 2, the last bin is found to exceed the peak threshold. As this is an




Fig. 3. Block diagram of the pixel. When photon timing, the multi-event time-to-digital converter (METDC) uses either a delay line (DL) or a series of delay flip flops (DFF), clocked by an internal ring oscillator (GRO) or external reference clock. For time gate generation, either the GRO or the external clock is used. The “background estimation” module also identifies the histogram bin with the maximum count (“peakbinID”), and the “peak detection” module calculates the peak threshold and applies it to the histogram. If a peak is detected then the “peakBinFlag” and the time gate adjusted according to the location of the peak (“peakbinID”). In case one of the histogram bin counters reaches full capacity (4095), all the bins are frozen, and the “overFlag” is raised.

outlying bin, the time gate is moved again to bring the peak into the middle. The peak is now contained in a middle bin, the time gate position is unchanged for exposure 4. Should the peak disappear in a subsequent exposure, the time gate position is first decremented, before scanning in the normal direction is resumed, if a peak is still not detected. In addition to the ambient level being estimated in exposure 1, the estimate is updated every time there is a shift in the time gate position.

When peak tracking, sensor is typically configured so that there is a 50% overlap between consecutive time gate positions, which with the signal peak occupying 2-3 histogram bins (identified as optimal for peak extraction assuming a Gaussian peak profile [20]) guarantees that there is time gate position containing the peak in its entirety. Assuming a sufficiently large return signal, and neglecting false peak detections, the maximum time to converge to a peak can be expressed as $N \times t_{exp}$ where N is the number of time gate positions and t_{exp} is the exposure time.

Instead of peak tracking, it is also possible to scan through time gate positions continually, as well as to assign specific time gates to different pixels. Up to 128 different time gate positions are available, leading to an effective overall histogram size of 1024 bins per pixel.

The coarse timing or time gate position can be generated either using an internal ring oscillator (GRO), or there is the option of using an external clock. For the fine timing or histogram bin width, either a ring oscillator, a delay line (DL), or an external clock is used. In the case of GRO and DL timing the resolution (bin width) is adjustable via external voltages, and can be made as small as 250 ps for DL. In practice, higher photon levels can lead to distortion in the GRO timing (potentially due to noise on the GRO voltage supply), making the delay line and external clock the preferred timing options.

We note that thanks to the adjustable bin size, and the ability to prescribe time gate positions, the sensor could also be conceivably operated in a zooming mode (Fig. 1) with suitable external control.

In addition to the histogram modality, there is also a photon counting mode (as indicated in Fig. 3), where SPADs are combined in pairs, and each histogram bin is repurposed to serve as a photon counter for a pair of SPADs. It is possible to switch between the histogram and photon counting modes on

the fly (with time gate positions being preserved), for instance for the purpose of off-chip intensity-guided depth upscaling [21].

In-pixel peak detection

The existence of a significant peak within the 8-bin histogram captured by a pixel is determined by comparing the bin with the maximum count h_{max} with the threshold given by the equation [22]:

$$h_{thresh} = B + 1.75\alpha\sqrt{B}, \quad (1)$$

where B is the estimated background level and α can be set to 1 or 2. Thus, under the assumption of Poisson noise on the bin counts, the peak has to be more than 1.75 or 3.5 standard deviations away from the baseline ambient level for it to be considered a “true” peak, with α setting the desired level of detection sensitivity.

The bin with the highest count (or “peak bin”) is identified using the network of comparators shown in Fig. 4, with each comparator returning the higher of its two inputs. To save circuit space, the network is reused in the estimation of the ambient level B , which is taken as the lower of the two inputs of the final comparator COMP7.

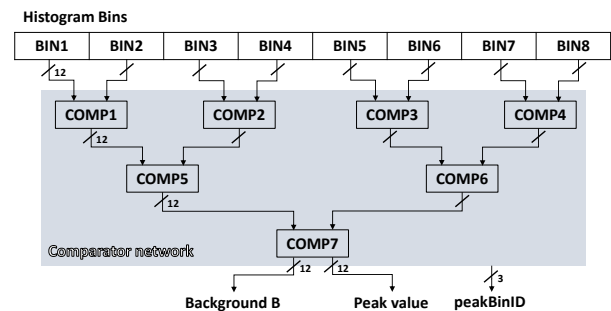


Fig. 4. Comparator tree used within the pixel to identify the histogram bin with the maximum count h_{max} , and to estimate the background level B .

To calculate the peak detection threshold in eqn. 1, a piecewise linear approximation is used for the factor $1.75\sqrt{B}$, as indicated in Fig. 5, which is implemented using comparator,

bit shifting, and summation logic. Each linear segment in the approximation has a gradient in the form of $1/2^n$, $n = 1, 2, \dots$ to allow computation using simple bit shift logic. The start and end points of the segments (which are hardwired into the logic) were determined using least-squares optimisation.

Fig. 6 shows the accuracy of the approximation of $1.75\sqrt{B}$ over the full range of possible background levels B . The maximum error is seen to be < 3 photons.

Output modes for depth imaging

In addition to the default histogram output, where all the histogram bins are read out, there are two modes providing further data compression in the form of direct depth readings (Table I).

In the bin-resolution depth mode, each pixel reports the position of the bin with the maximum count (and whether this is a statistically significant peak). An almost tenfold reduction is thereby achieved in the amount of data read out per frame.

In the sub-bin resolution depth mode, peak extraction is performed outside the array (in column parallel logic) using a centre-of-mass (CMM) approach, which again results in a significant ($\times 7$) data compression compared with reading out histograms.

In addition to the modes in the table, there are two smart readout options, where only pixels with a peak or those with a peak that is moving are read out, and the rest are replaced by zeros. A moving peak is defined here as one requiring a shift in time gate position. These selective readout options aim to reduce toggling on the output lines as well as downstream processing requirements. The latter option where only pixels with changing depth are read out has parallels with dynamic vision sensors [23], although the readout is frame-based here rather than asynchronous, event-based.

Primarily intended for low-ambient environments, there is the further option of applying multiple subexposures between frame readouts to pixels that are peak-searching, to allow for faster convergence to peaks. Furthermore, during readout, rows can be selectively skipped, via external control, for example if none of the pixels in a row have detected a peak.

TABLE I

Comparison of the different output formats offered by the sensor for depth imaging, in terms of the type of output data, the number of bits reported per pixel, and the maximum achievable frame rate. Note that the frame rates achievable in practice depend on the exposure time required to register histograms with suitable signal-to-noise ratio, and thus, to a large part, on the optical power of the laser source. For the laser source used here, frame rates of > 1 kFPS can be obtained.

Mode	Output format*	# bits per pixel	Max. frame rate (FPS)
Histogram	hist. bins and hist. peak data	108	29k
Bin res. depth	hist. peak data	12	260k
Sub-bin res. depth	centre of mass of hist bins**	15	208k

*in addition to 7-bit time gate position; the histogram peak data consists of peak bin flag, peak bin ID and overflow flag

**after compensation for background counts

On-chip depth computation using CMM

For computing a sub-bin resolution estimate for the depth d , we use a variant of the centre-of-mass scheme [21], described by the following equation:

$$\hat{d} = \frac{\sum_{t=1}^8 t (h_t - B)}{\sum_{t=1}^8 (h_t - B)}, \quad (2)$$

where h_t ($t = 1 \dots 8$) are the histogram bins at a given pixel, and B is the estimated background level. To simplify the hardware implementation, B is taken here as the minimum bin count amongst the 8 bins.

The implementation of the CMM module is shown in Fig. 7, with the division being carried out by a parallel divider to a three decimal point accuracy, allowing the position of the peak to be determined with a resolution of $1/8$ of the bin size. The bin counts are loaded into the module two at a time. Due to a design bug affecting the current version of the chip, there is an error in the order in which the bins are presented to the module. As a consequence, the histogram peak has to be contained in the middle two bins for \hat{d} to be correctly calculated. We note that in a future implementation, the design could be modified to read out a photon count alongside the peak position, to facilitate off-chip, reflectivity-based range walk error correction.

III. RESULTS

To create a camera system, the sensor is attached to a custom PCB board alongside an Opal Kelly XEM7310 FPGA integration module, which generates the required control signals and streams the SPAD data over a USB 3.0 link to a PC running a Matlab software interface. A C-mount, 25 mm/f1.4 objective (Thorlabs MVL25M23) is used in front of the sensor, giving an approximately 20×5 degree field-of-view (FOV). For depth imaging, the sensor is paired with a compact 850 nm VCSEL source with 10 ns pulse width and 60 W peak optical power triggered at 1.2 MHz. Lenses are used to shape the laser beam according to the sensor FOV, and a 10 nm ambient filter (Thorlabs FL850-10) is added to the receiver optics. The whole system is powered using standard USB ports: a USB 3.0 port supplies the camera board, and a USB-C port powers the laser. The SPAD excess bias was set to $VEB = 2 V$.

We note that although the camera is sufficiently compact and lightweight to be mounted on a small robot or drone, it would be difficult to miniaturise it into a module suitable for mobile phones, for example, due to the relatively high electrical and optical power of the laser source. The laser source, as currently configured, is eye-safe from a distance of around 0.5 m, and could potentially be made fully eye-safe, by using a source with a larger emitter area (to decrease the power density) in a suitable enclosure, and/or reducing the pulse width. The bandwidth of the ambient filter is relatively narrow so may require temperature control of the laser source to stabilise the laser wavelength regardless the environmental temperature.

Unless stated otherwise, the presented results correspond to the histogram readout mode. For short range 3D imaging, the exposure time (which in turn sets the frame rate) was chosen such that the signal peak for all pixels would be below the




Fig. 5. Logic used for peak threshold (h_{thresh}) calculation, based on the estimated background level B . The logic includes a piece-wise linear approximation of the quadratic function (see also Fig. 6). B is compared to a set of values, defining the end points of the linear segments. The linear functions within the segments are implemented via bit shifting (" \gg ") and the addition of a constant term (specific to the given segment). The detection sensitivity can be adjusted via the α input.

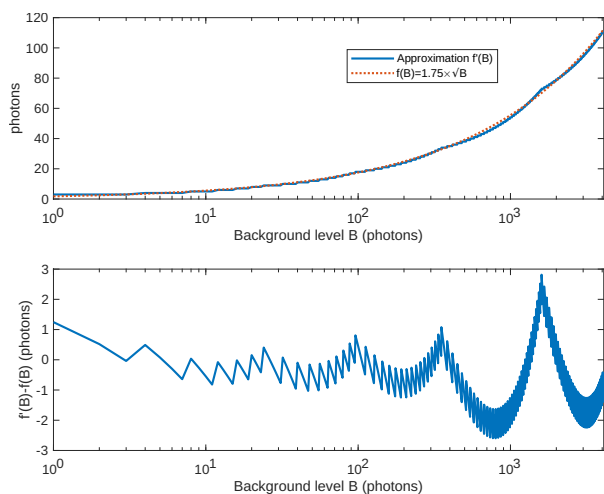


Fig. 6. Accuracy of the approximation $f'(B)$ for the quadratic function $f(B) = 1.75\sqrt{B}$ built into the pixels for peak threshold calculation. The top graph plots the quadratic function $f(B)$ and the approximation $f'(B)$ for background levels B ranging from 0 to 4095 photons (the maximum possible level that can be recorded). The bottom graph shows the error in the approximation over the same range of B .

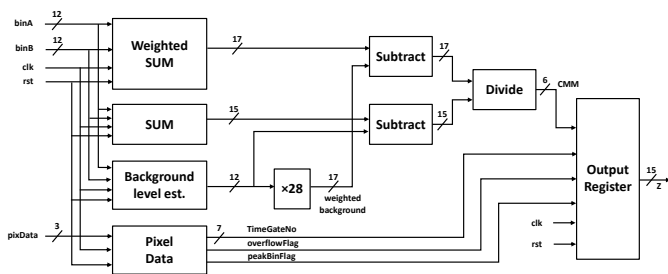


Fig. 7. Block diagram of on-chip centre-of-mass (CMM) processing module providing sub-bin resolution depth estimates. The block multiplying by the fixed constant 28 corresponds to the factor $\sum_t t$ by which B is multiplied in the numerator of 2, assuming the indexing $t = 0 \dots 7$.

maximum count of 4095; as the histogram is frozen when this value is reached, there is no benefit to exposing for longer. For longer range imaging, the exposure time was chosen to be sufficiently long to ensure reliable peak detection in the given conditions.

Intensity imaging

Fig. 8 shows an example photon counting image captured by the sensor. The raw image (top image) is seen to be pixelated horizontally, due to the uneven spacing of pixels (as a result of the gaps created by the processing unit in each pixel). By accounting for this uneven spacing, and applying interpolation, the level of pixelation is reduced (bottom image).

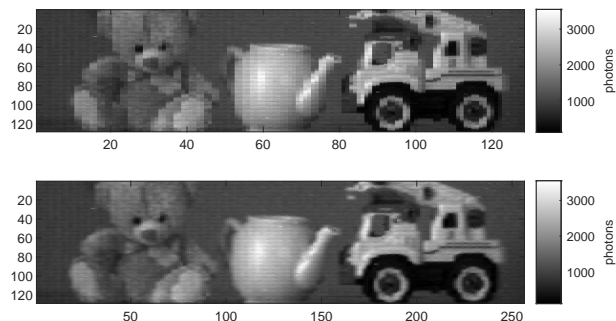


Fig. 8. Example photon counting (intensity) image. The top image depicts the raw data from the sensor, the bottom image is the result of interpolation to account for the uneven spacing of pixels horizontally. The data was captured indoors with a $200 \mu s$ exposure time, equating to 5 kFPS. The laser was switched off and the ambient filter removed.

Histogram non-linearity

The non-linearity of the sensor when photon timing was measured by capturing a large number of exposures under ambient conditions, and assessing the bin-to-bin differences in the mean photon count per bin within each macro pixel. Figs. 9a and b show example results for the delay line and external clock timing options, respectively, each based on 4000 exposures. For delay line fine timing with a 1 ns bin size the maximum DNL for a typical (median) pixel is found to be $\approx 14\%$. In the case an external clock is used for timing, a nonlinearity of $\approx 1.6\%$ was measured for a 8 ns bin size

(the minimum bin size obtainable in this configuration being subject to bandwidth limitations in feeding a clock into the chip and distributing it across the array).

Due to the superior linearity of external clocking timing, and the pulse width of the compact laser sources used here being better matched to wider bin sizes, the external clock option was adopted as a default configuration, and depth imaging results presented here were obtained with an external clock.

Specificity of peak detection

The specificity of the in-pixel peak detection logic was tested by recording sequences of frames under ambient light. The measurements were repeated for different background levels by varying the exposure time. Fig. 10 plots the resulting false positive rate for the two different sensitivity settings, in the case of both external clock and delay line timing. For external clock timing with $\alpha = 1$, the false positive rate (FPR) is around 10% at higher background levels; for $\alpha = 2$ the FPR approaches 0.2%. Similar values of 12% and 0.14% are obtained in a Monte Carlo simulation of the same experiment, with the square root in eqn. 1 calculated exactly, and the bin counts all drawn from the same Poisson process under the assumption of $B = 1000$ photons. It is noted that these values are higher than what would be expected in theory assuming a known background level B (4% and 0.023% for z scores of 1.75 and 3.5 [24]) due to the approximate nature in which B is estimated in-pixel. For delay line timing, the FPR continues to climb as the background level is raised, due to the higher non-linearity in this mode. To ensure an FPR of, say $<2\%$, the background level should be around ≈ 100 counts/bin or lower.

Short range, indoor imaging

For short range imaging, there is no need to scan across time gate positions, as a single time gate provides a 9.6 m unambiguous distance range (assuming external clock timing). The time gate is thus fixed at the first position. Fig. 11 shows example results obtained indoors with a fixed time gate. Fig. 11b depicts the depth map from a single exposure in the histogram readout mode, with depth extraction carried out off-chip [21]. As previously demonstrated in [7], it is possible to increase the lateral resolution of a macro-pixel-based imager by successively enabling individual SPADs within the pixel. A similar time-multiplexing approach can be applied to the present sensor, with multiple exposures being taken and combined off-chip to compose a higher resolution depth map (bottom image in Fig. 11c and d).

Another short range example is given in Fig. 12, which compares two depth maps of the same scene, the top image having been obtained in histogram mode (with off-chip peak extraction) and the bottom image being a result of on-chip (CMM) depth computation. Recalling the bug currently affecting CMM, an exact correspondence cannot be expected. Nevertheless, the two depth maps are visually similar. Noting that the depth in the scene varies by a distance equivalent to less than a single histogram bin, the results demonstrate the ability of CMM to output sub-bin resolution readings.

Medium range, outdoor imaging

For imaging outdoors the peak tracking functionality was activated with $N = 16$ time gate positions (giving 81.6 m of unambiguous range) and $\alpha = 2$. The accuracy and precision of the system were evaluated for flat targets of different reflectivity η up to range of 50 m. Ground truth readings were taken with a laser range finder (Bosch GLM 250 VF), and the ambient level was measured with a calibrated light meter to be 20 kLux. Fig. 13 summarises the results for $\eta=0.4,0.85$, obtained at a frame rate of 50 FPS. The results indicate a maximum non-linearity of 5 cm and a precision (standard deviation) of <13 cm. A target detection rate of 100% was measured throughout. In practice, the performance of the system is dependent on the laser source that the sensor is paired with, and given the availability of laser sources with much higher peak power than the 60 W VCSEL source used in this study, we expect longer ranges, under potentially higher ambient levels, to be achievable.

The peak tracking functionality is illustrated in Fig. 14 through a sequence of consecutive exposures after switch on. A panel is shown for each exposure, plotting the time gate position (top graph), the depth extracted from the histograms (middle graph), and the histograms corresponding to two pixels (bottom), one in the left side of the field-of-view and the other in the centre. In the first exposure (panel 1) no surfaces can be seen, and both of the highlighted pixels move to the next time gate position. After exposure 2, a surface starts emerging in the captured scene. The second pixel has detected a peak but it is at an outlying bin, so a further time shift is required. The first pixel is still peak searching. Exposure 3 reveals most of the scene, which is of a person sitting next to a tree and holding a ball, with a swing on the other side. Pixel 1, which observes the ball has now locked onto the peak. Pixel 2, looking at the background has detected a surface but still needs a time shift. After exposure 4, most of the pixels have homed in, including both of the highlighted pixels. More details have appeared in the scene, including the second leg of a swing. Subsequent exposures (not shown), are mostly unchanged, apart from a few remaining pixels also converging.

Figs. 15 and 16 present snapshots of data from dynamic scenes. In Fig. 15, a person is walking around a moving swing. Two depth maps are shown: sub-bin resolution depth, extracted off-chip from histogram data, and bin-resolution depth, as reported by the sensor. Even though the latter output is one tenth of the size compared to histogram frames, and provides a very coarse depth reading given the large bin size of 8 ns, it is still seen to capture key features in the scene. Fig. 16 gives a demonstration of the smart readout mode, where only pixels with changing depth are reported, and the rest replaced by zeros. On the top, we see the original sequence, of two people passing a ball to each other, and on the bottom we see the result of only presenting pixel data for pixels with a change in time gate position. This gives a “dynamic vision sensor”-type output [23], with the movement of the ball clearly indicated.

Videos of these examples, and of another garden scene where the camera is panned, can be found in the supplement-

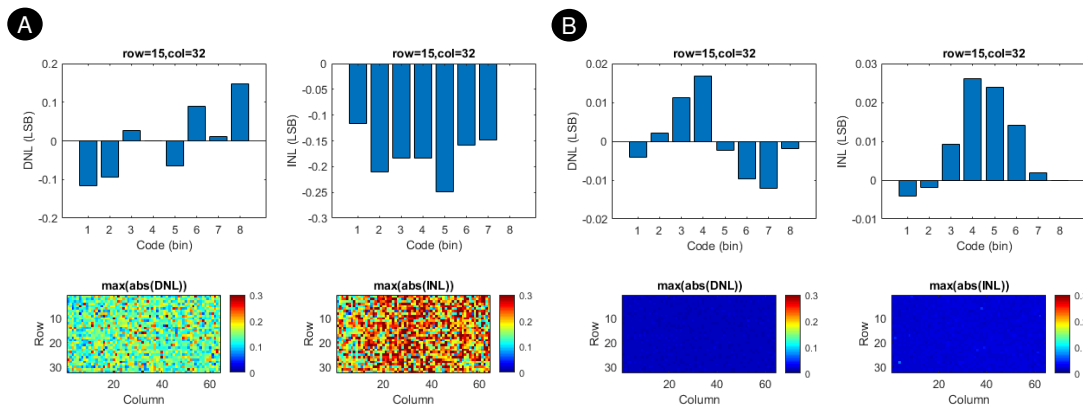


Fig. 9. Non-linearity in the photon timing histogram: a) for a nominal bin width of 1 ns, when the delay line is used for timing, b) for a nominal bin width of 8 ns, when an external clock (of frequency 125 MHz) is used for timing. The top figures plot the differential and integral nonlinearity (DNL/INL) for a selected pixel; the bottom figures show the maximum DNL and INL across the pixel array.

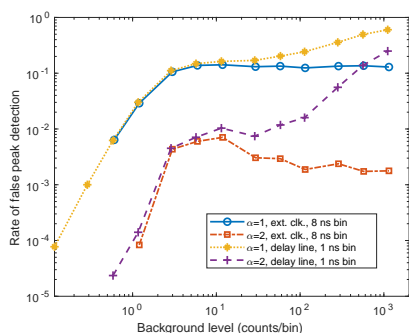


Fig. 10. False positive rate in the peak detection versus the background level for different α settings and timing modes. Each data point was generated using 500 frames, and corresponds to the mean rate of false positive detections across the array. The background level was measured as the mean bin count and varied by changing the exposure time. Assuming the camera setup used here, together with a 20 ms exposure time, a background level of $B = 1000$ photons/bin corresponds to 15 klux and 120 klux ambient for bin widths of 8 ns and 1 ns respectively.

tary material (Visualisations 1, 2, 3 and 4, which are accessible by clicking on the embedded hyperlinks). In addition, two longer range sequences are presented. In the first example, a person is walking towards the camera in an open space, with pixels observing the person successfully tracking the motion in z (Visualisations 5 and 6). The second example (Visualisation 7) compares sequences captured in the tracking and sliding (continuous scanning) modalities using the same exposure time. In the sliding mode, the camera steps through time gate positions 1 to 16, and the first peak detected by each pixel is plotted. Both sequences are played back in real time. In the tracking mode, a $\times 16$ higher frame rate is achieved, which could potentially be exploited in post-processing to obtain more accurate depth or motion estimates in z for tracked surfaces. No post-processing has been applied to any of the data, apart from centre-of-mass peak extraction (where indicated), and zero-offset correction, based on subtracting an a priori depth image of a flat calibration target.

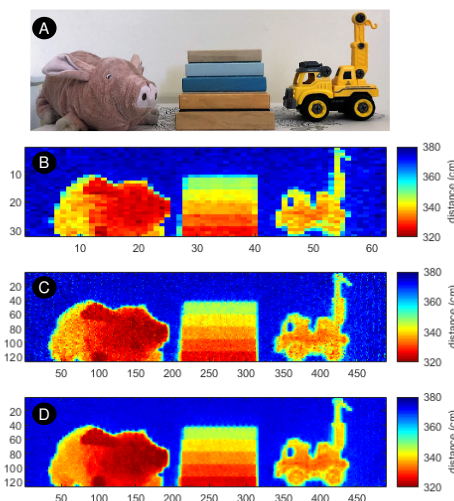


Fig. 11. 3D-imaging results for indoor scene with toys: a) RGB image of the scene; the separation between the wooden blocks is $\delta z \approx 5$ cm, b) depth map from a single 1 ms exposure (1 kFPS frame rate), c) the result of 16×3.2 ms exposures being combined off-chip for a high resolution depth map (20 FPS overall frame rate), d) the average of 20 high resolution images. To produce the higher resolution depth map, one SPAD at a time was enabled within the 4×4 array in each mpixel, the vertical dead zones between the pixel columns being interpolated over in post-processing.




Fig. 12. 3D-imaging results for indoor scene with plush pig and wooden arcs. The top image corresponds to off-chip depth extraction (histogram readout), the bottom image is for on-chip (CMM) depth computation. Both images were captured with a 1 ms exposure time (1 kFPS).

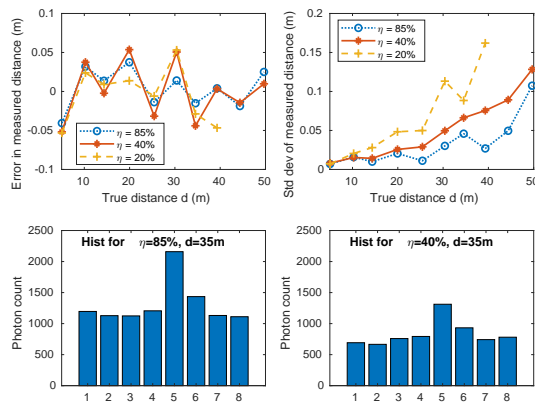


Fig. 13. Characterisation of system accuracy and precision for two different target reflectivities over a 50 m distance range. The top graphs plot the error and standard deviation in the measurement. Each data point corresponds to a single pixel observing the target, and is computed from 168 consecutive exposures. The bottom graphs show sample histograms for the two targets at a distance of $d = 35$ m. Note that at $d = 45$ m and 50 m the observed target width became less than the size of the pixel, resulting in an increased level of background photons relative to signal photons. A target of 20% reflectivity was also used in the test and was reliably detected up to $d = 40$ m, where the standard deviation was 16 cm, and the histograms showed a signal-to-background ratio of 0.04.

In-built and external time gate control

When peak searching or tracking, the background estimate used for calculating the peak threshold is based on the previous time gate position, to avoid the signal peak from affecting the background estimate. By default, once a peak is detected, the background estimate is no longer updated, unless the peak moves to the edges of the histogram (bin 1, 2, 7 or 8, such that a time gate shift is required). One downside of this strategy is that should the signal peak disappear, but the background level increase at the same time (say due to a high reflectivity target, at a different range, moving into view), the pixel may continue detecting a peak (and staying in the given time gate position), until the background is updated. This can therefore slow down the tracking of surfaces. To counter the problem, the sensor could be operated such that the laser is switched off in alternating frames, from which an up-to-date background level would be estimated.

Furthermore, by controlling the time gate positions externally, more elaborate peak tracking strategies could be implemented, such as, for each pixel (1) scanning the whole time range then hopping between the time gate positions with peaks (to enable multi-surface tracking), and (2) shifting the time gate position by 8 bins at a time (rather than the default 4 bins) when no peak is detected, for faster convergence to the peak.

As with other partial histogramming approaches, problems may be encountered when an object is moving laterally at a speed faster than the time taken for pixels to converge (or capture a full histogram). In the case of the present sensor, small objects moving at a speed such that pixels are passed in under $N \times t_{exp}$ (where N is the number of time gate positions and t_{exp} is the exposure time) may potentially be missed.

Power consumption

The chip power consumption in outdoor operation at 50 FPS under 30 klux of ambient light was measured to be 70 mW, of which 32 mW is consumed by SPAD detectors, 36 mW by the digital processing logic, and 1.32 mW by the I/O pads. The latter, I/O power consumption is strongly dependent on the frame rate, as well as the readout mode. In the default histogram mode, power consumption ranges from 1.32 mW-11.6mW (50 FPS-1 kFPS), reducing to 0.5 mW-1.85 mW in CMM mode. Further reductions can be obtained using the smart readout modes, which can lower the I/O power consumption to 0.43 mW-1.62 mW in histogram mode and 0.26 mW-0.99 mW in CMM mode. These results demonstrate the effectiveness of the reduced readout modes in saving I/O power, with additional savings expected from the row skipping feature. The chip was fabricated in STMicroelectronics' general purpose 40 nm CMOS technology; a SPAD detector power improvement may be made in future versions of this chip by using STMicroelectronics' SPAD-optimised technology which would lead to lower DCR and <100 fC per SPAD pulse [25].

Comparison with state-of-the-art sensors

A comparison of the present sensor with state-of-the-art SPAD dToF image sensors featuring on-chip histogramming is shown in Table II. The measurements for the present sensor correspond to peak tracking mode with 8 ns bin width and histogram readout; similar or higher accuracy and precision is expected at shorter bin width settings provided the laser pulse width is scaled accordingly [26]. However, the shorter the bin width, the more time gate positions are needed to cover the same range, which also correspond to the maximum number of exposures required to converge to the target.

IV. DISCUSSION AND OUTLOOK

A SPAD dToF imager has been presented that the authors believe to be the first reported imager with in-pixel ambient estimation, surface detection and tracking, resulting in a scalable architecture.

Each pixel has an independent time gate that automatically tracks the signal peak, and histogramming logic to generate photon timing histograms. Multiple events can be captured per laser cycle to enable outdoor operation without pile-up distortion. Thanks to the significant in-pixel data compression, the device can operate at high-frame rates, and on-chip depth computation is available. Experimental results are given demonstrating the viability of the device for medium range LIDAR at high frame rates.

Although not considered here, there is the potential to enhance the time gate mechanism in each pixel by switching to external control of the time gate positions, for instance, with the intention to capture multiple peaks. Such external control would also allow for more accurate background estimation and peak detection, for example, by capturing pairs of exposures, one with the laser off and the other with the laser on.

There is also scope for more optimised illumination strategies compared with the simple flood illumination at fixed repetition rate that is currently adopted. As described in [16], when




Fig. 14. Frames from sequence with person holding a ball, captured at 50 FPS. The top graphs show the time gate position, the middle graph is the extracted depth, and the bottom graphs show histograms from two selected pixels. The panels correspond to the third exposure onwards following switch on.

TABLE II
Comparison of present sensor with SPAD dToF image sensors with embedded histogramming

	This work	Zhang [4]	Kumagai [10]	Zhang [12]	Stoppa [7]	Park [27]
<i>Sensor characteristics</i>						
Technology	40 nm FSI	180 nm FSI	90/40 nm BSI	65/65 nm BSI	45/45 nm BSI	110 nm FSI
Pixel array	64 × 32	252 × 144	200 × 56	240 × 160	320 × 240	80 × 60
Pixel size	114 × 54 μm ² (4 × 4 SPADs)	28.5 × 28.5 μm ²	30 × 30 μm ² (3 × 3 SPADs)	16 × 16 μm ²	12.5 × 12.5 μm ²	75 × 75 μm ² (6 × 1 SPADs)
TDC/hist. type	In-pixel peak tracking	On-chip zoom	On-chip full hist.	In-pixel zoom ⁽¹⁾	In-pixel sliding ⁽²⁾	In-pixel zoom
bit-depth/# bins	12-bit/8	12-bit/16	12-bit/2000	10-bit/32	12-bit/32	6-bit/4
# TDCs	2048	1728	384	600	4800	4800
TDC res.	0.25-8 ns	48-195.2 ps	1 ns	97.65 ps	200-300 ps	5 ns
TDC range	129-4128 ns	50-200 ns	2000 ns	100 ns	51.2-76.8 ns	320 ns
Power consumption	70 mW	2540 mW	1192 mW	306 mW	<1500 mW ⁽³⁾	132 mW
<i>Time-of-flight measurements</i>						
Type of illumination	Flash	Flash	MEMS scanning	Flash	Flash	Flash
Field of view (H×V)	20° × 5°	40° × 20°	25.2° × 9.45°	63° × 41°	74.5° diagonal	8.1° × 10.8°
Max. distance	50 m	50 m ⁽⁴⁾	200 m	9.5 m	2-5.8 m ⁽⁵⁾	45 m ⁽⁶⁾
Background	20 klux	dark	117 klux	500 lux	50 klux	30 klux
Target reflectivity	40%	60%	10%, 90%	50%	92%	-
Frame rate	50 FPS	30 FPS	20 FPS	20 FPS	5-30 FPS	30 FPS
Accuracy	5 cm	8.8 cm	< 15 cm	1 cm	< 5%	2.5 cm
Precision	< 13 cm	1.4 mm	-	0.9 cm	< 0.5%	1.5 cm
Laser wavelength	850 nm	637 nm	905 nm	940 nm	940 nm	905 nm
Rep. rate/pulse width	1.2 MHz/10 ns	40 MHz/40 ps	-/4.5 ns	10 MHz/-	-/360 ps	100 kHz/5 ns
Optical power (peak, avg.)	60 W (peak)	0.5 W, 2 mW	45 W (peak)	15 W, 90 mW	35 W, 93 mW	3 mW (avg.)
Camera objective	25mm, f/1.4	25mm, f/1.5	-	2.3mm, f/1.4	f/1.4	-

⁽¹⁾each TDC is shared by a subgroup of 8 × 8 pixels, consisting of 4 macropixels; one macropixel is activated at a time

⁽²⁾each TDC is shared by a subgroup of 4 × 4 pixels; one pixel is activated at a time

⁽³⁾total system power consumption including laser source and depth map processing

⁽⁴⁾the unambiguous range is 3.5 m; the distance, accuracy and precision figures relate to single point measurement

⁽⁵⁾5.8 m range and 30 FPS is obtained at 80 × 60 resolution (125 mW avg. laser power), reducing to 4.4 m range for a 18% reflectivity target indoors

⁽⁶⁾measured with a reduced frame rate of 1.5 FPS; the accuracy and precision were measured in indoor conditions over a 3-4.5 m range

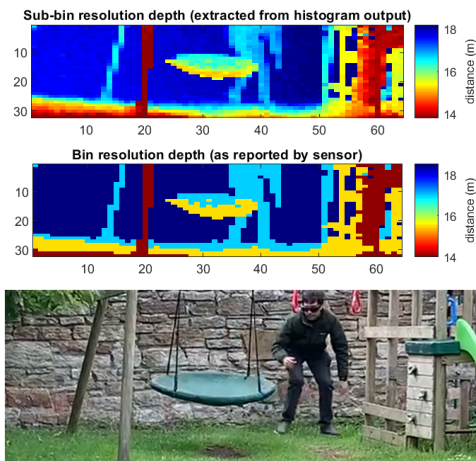


Fig. 15. Snapshot of sequence with swing, captured at 50 FPS. The top graph shows the depth extracted from the histograms; the middle graph gives the bin resolution depth (calculated on-chip based on the histogram bins with the highest photon count, with no interpolation applied, so the accuracy is commensurate with the bin width). In the bottom an RGB image is shown of the scene. See also Visualisations 1 and 2.

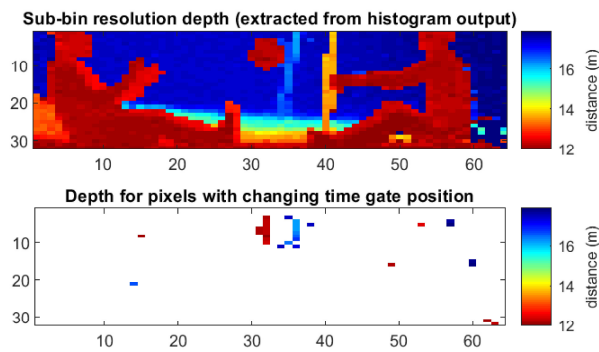


Fig. 16. Snapshot of ball passing sequence, captured at 100 FPS. The top graph shows the full depth map (obtained in histogram mode), the bottom graph illustrates selective readout according to changes in time gate position. See also Visualisation 3.

stepping through different time gate positions, it is beneficial to adjust the laser power accordingly. In particular, a higher laser repetition rate (or higher pulse energy) is advocated for farther time window positions (to ensure a sufficient SNR for the histogram peaks from far away targets to be reliably detected), and lower repetition rate (or lower pulse energy) for closer time window positions to conserve laser power. In the context of the present sensor where each pixel has its own time gate position, this would require an illumination source with individually addressable elements, allowing the number of laser cycles to be adjusted on a per-pixel basis.

State-of-the-art SPADs in 3D-stacked, backside illuminated (BSI) technology [28], [29], [25] offer significantly higher photon detection efficiency (PDP) compared with the front-side illuminated (FSI) technology used in the present chip. By combining the latest SPADs with the processing architecture presented here, the same signal (and background) counts, and

therefore depth precision, could be attained at much lower exposure times, and consequently, at higher frame rates.

REFERENCES

- [1] 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 Processing Magazine*, vol. 37, no. 4, pp. 62–71, 2020.
- [2] R. Roriz, J. Cabral, and T. Gomes, "Automotive LiDAR technology: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [3] D. Wang, C. Watkins, and H. Xie, "MEMS mirrors for LiDAR: A review," *Micromachines*, vol. 11, no. 5, 2020. [Online]. Available: <https://www.mdpi.com/2072-666X/11/5/456>
- [4] C. Zhang, S. Lindner, I. M. Antolović, J. M. Pavia, M. Wolf, and E. Charbon, "A 30-frames/s, 252×144 SPAD flash LiDAR with 1728 dual-clock 48.8-ps TDCs, and pixel-wise integrated histogramming," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 4, pp. 1137–1151, 2018.
- [5] S. W. Hutchings, N. Johnston, I. Gyongy, T. Al Abbas, N. A. Dutton, M. Tyler, S. Chan, J. Leach, and R. K. Henderson, "A reconfigurable 3-d-stacked SPAD imager with in-pixel histogramming for flash LiDAR or high-speed time-of-flight imaging," *IEEE Journal of Solid-State Circuits*, vol. 54, no. 11, pp. 2947–2956, 2019.
- [6] B. Kim, S. Park, J.-H. Chun, J. Choi, and S.-J. Kim, "7.2 a 48×40 13.5mm depth resolution flash LiDAR sensor with in-pixel zoom histogramming time-to-digital converter," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, 2021, pp. 108–110.
- [7] D. Stoppa, S. Abovyan, D. Furrer, R. Gancarz, T. Jessenig, R. Kappel, M. Lueger, C. Mautner, I. Mills, D. Perenzoni *et al.*, "A reconfigurable QVGA/Q3VGA direct time-of-flight 3D imaging system with on-chip depth-map computation in 45/40 nm 3d-stacked BSI SPAD CMOS," in *Proc. Int. Image Sensor Workshop*, 2021.
- [8] S. M. Patanwala, I. Gyongy, H. Mai, A. Abmann, N. A. Dutton, B. R. Rae, and R. K. Henderson, "A high-throughput photon processing technique for range extension of SPAD-based LiDAR receivers," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 2, pp. 12–25, 2021.
- [9] H. Seo, H. Yoon, D. Kim, J. Kim, S.-J. Kim, J.-H. Chun, and J. Choi, "Direct TOF scanning LiDAR sensor with two-step multievent histogramming TDC and embedded interference filter," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 4, pp. 1022–1035, 2021.
- [10] O. Kumagai, J. Ohmachi, M. Matsumura, S. Yagi, K. Tayu, K. Amagawa, T. Matsukawa, O. Ozawa, D. Hirono, Y. Shinozuka, R. Homma, K. Mahara, T. Ohyama, Y. Morita, S. Shimada, T. Ueno, A. Matsumoto, Y. Otake, T. Wakano, and T. Izawa, "7.3 a 189×600 back-illuminated stacked SPAD direct time-of-flight depth sensor for automotive LiDAR systems," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, 2021, pp. 110–112.
- [11] P. Padmanabhan, C. Zhang, M. Cazzaniga, B. Efe, A. R. Ximenes, M.-J. Lee, and E. Charbon, "7.4 a 256×128 3D-stacked (45nm) SPAD FLASH LiDAR with 7-level coincidence detection and progressive gating for 100m range and 10klux background light," in *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64. IEEE, 2021, pp. 111–113.
- [12] C. Zhang, N. Zhang, Z. Ma, L. Wang, Y. Qin, J. Jia, and K. Zang, "A 240×160 3D-stacked SPAD dToF image sensor with rolling shutter and in-pixel histogram for mobile devices," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 2, pp. 3–11, 2022.
- [13] F. Villa, F. Severini, F. Madonini, and F. Zappa, "SPADs and SiPMs arrays for long-range high-speed light detection and ranging (lidar)," *Sensors*, vol. 21, no. 11, 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/11/3839>
- [14] F. Piron, D. Morrison, M. R. Yuce, and J.-M. Redouté, "A review of single-photon avalanche diode time-of-flight imaging sensor arrays," *IEEE Sensors Journal*, vol. 21, no. 11, pp. 12 654–12 666, 2021.
- [15] I. Gyongy, N. A. W. Dutton, and R. K. Henderson, "Direct time-of-flight single-photon imaging," *IEEE Transactions on Electron Devices*, vol. 69, no. 6, pp. 2794–2805, 2022.
- [16] F. Taneski, T. A. Abbas, and R. K. Henderson, "Laser power efficiency of partial histogram direct time-of-flight lidar sensors," *Journal of Lightwave Technology*, pp. 1–10, 2022.
- [17] T. Al Abbas, N. A. W. Dutton, O. Almer, S. Pellegrini, Y. Henrion, and R. K. Henderson, "Backside illuminated SPAD image sensor with 7.83μm pitch in 3D-stacked CMOS technology," in *2016 IEEE International Electron Devices Meeting (IEDM)*, 2016, pp. 8.1.1–8.1.4.

- [18] S. Pellegrini, B. Rae, A. Pingault, D. Golanski, S. Jouan, C. Lapeyre, and B. Mamdy, "Industrialised spad in 40 nm technology," in *2017 IEEE International Electron Devices Meeting (IEDM)*. IEEE, 2017, pp. 16–5.
- [19] N. A. W. Dutton, S. Gnecci, L. Parmesan, A. J. Holmes, B. Rae, L. A. Grant, and R. K. Henderson, "11.5 a time-correlated single-photon-counting sensor with 14gs/s histogramming time-to-digital converter," in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, 2015, pp. 1–3.
- [20] N. Hagen, M. Kupinski, and E. L. Dereniak, "Gaussian profile estimation in one dimension," *Applied Optics*, vol. 46, no. 22, pp. 5374–5383, 2007.
- [21] I. Gyongy, S. W. Hutchings, A. Halimi, M. Tyler, S. Chan, F. Zhu, S. McLaughlin, R. K. Henderson, and J. Leach, "High-speed 3D sensing via hybrid-mode imaging and guided upsampling," *Optica*, vol. 7, no. 10, pp. 1253–1260, Oct 2020. [Online]. Available: <http://opg.optica.org/optica/abstract.cfm?URI=optica-7-10-1253>
- [22] S. Gnecci and C. Jackson, "A 1×16 SIPM array for automotive 3d imaging lidar systems," in *Proceedings of the 2017 International Image Sensor Workshop (IISW)*, Hiroshima, Japan, 2017, pp. 133–136.
- [23] T. Delbrück, B. Linares-Barranco, E. Culurciello, and C. Posch, "Activity-driven, event-based vision sensors," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. IEEE, 2010, pp. 2426–2429.
- [24] R. J. Larsen and M. L. Marx, *An introduction to mathematical statistics*. Prentice Hall, 2005.
- [25] S. Pellegrini and B. Rae, "A low-power high-dynamic-range SPAD for sustainable sensing," in *Proceedings of the 2022 International SPAD Sensor Workshop (ISSW)*, 2022, p. 24.
- [26] L. J. Koerner, "Models of direct time-of-flight sensor precision that enable optimal design and dynamic configuration," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–9, 2021.
- [27] S. Park, B. Kim, S.-H. Han, J.-H. Cho, J.-H. Chun, J. Choi, and S.-J. Kim, "An 80×60 flash lidar sensor with in-pixel delta-intensity quaternary search histogramming tdc," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 11, pp. 3200–3211, 2022.
- [28] K. Ito, Y. Otake, Y. Kitano, A. Matsumoto, J. Yamamoto, T. Ogashara, H. Hiyama, R. Naito, K. Takeuchi, T. Tada, K. Takabayashi, H. Nakayama, K. Tatani, T. Hirano, and T. Wakano, "A back illuminated $10\mu\text{m}$ SPAD pixel array comprising full trench isolation and cu-cu bonding with over 14% PDE at 940nm," in *2020 IEEE International Electron Devices Meeting (IEDM)*, 2020, pp. 16.6.1–16.6.4.
- [29] K. Morimoto, J. Iwata, M. Shinohara, H. Sekine, A. Abdelghafar, H. Tsuchiya, Y. Kuroda, K. Tojima, W. Endo, Y. Maehashi, Y. Ota, T. Sasago, S. Maekawa, S. Hikosaka, T. Kanou, A. Kato, T. Tezuka, S. Yoshizaki, T. Ogawa, K. Uehira, A. Ehara, F. Inui, Y. Matsuno, K. Sakurai, and T. Ichikawa, "3.2 megapixel 3D-stacked charge focusing SPAD for low-light imaging and depth sensing," in *2021 IEEE International Electron Devices Meeting (IEDM)*, 2021, pp. 20.2.1–20.2.4.
- [30] I. Gyongy, A. Erdogan, N. A. Dutton, H. Mai, F. M. Della Rocca, and R. K. Henderson, "A 200kFPS, 256×128 SPAD dToF sensor with peak tracking and smart readout," in *Proc. Int. Image Sensor Workshop*, 2021.

ACKNOWLEDGEMENTS

The support of the Engineering and Physical Science Research Council (EPSRC) via grants EP/M01326X/1 and EP/S001638/1 is acknowledged. The authors are grateful to STMicroelectronics for chip fabrication and in particular to Sara Pellegrini and Thierry Lachaud for their support. The authors would also like to thank Filip Taneski for his assistance in experiments. Portions of this work were presented at the International Image Sensor Workshop in 2021 [30].



Istvan Gyongy received the M.Eng. and Ph.D. degrees from Oxford University, Oxford, U.K., in 2003 and 2008, respectively. Following a period in industry, where he worked on processors for smartphones and a cloud-connected activity tracking system for dairy farms, he joined The University of Edinburgh, Edinburgh, U.K. He is currently a Lecturer and is developing single-photon avalanche diode (SPAD) cameras and exploring applications in 3-D capture and the life sciences.



Ahmet T. Erdogan received the B.Sc. degree in electronics engineering from Dokuz Eylul University, İzmir, Turkey, in 1990, and the M.Sc. and Ph.D. degrees in electronics engineering from Cardiff University, Cardiff, U.K., in 1995 and 1999, respectively.

Since 1999, he has been working on several research projects at The University of Edinburgh, Edinburgh, U.K., where he is currently a Research Associate with the CMOS Sensors and Systems Group, School of Engineering. His research interests

include low-power VLSI design, reconfigurable computing, and CMOS image sensors.



Neale A.W. Dutton (Senior Member, IEEE) received the Master of Engineering degree in electronics and electrical engineering from The University of Edinburgh, Edinburgh, U.K., in 2011, and the Ph.D. degree from The University of Edinburgh in 2016, with a focus on single-photon avalanche diode (SPAD) image sensors for photon counting and time-of-flight (TOF) imaging funded by the Imaging Division, STMicroelectronics, Edinburgh.

He is currently a Principal Engineer and a Senior Member of the Technical Staff with STMicroelectronics with research and development interests in SPADs and TOF. He is supervising one Ph.D. student and is an author of 38 papers and one book chapter. He is an inventor of 39 patents and trade secrets. Dr. Dutton is a Reviewer of IEEE Journal of Solid-State Circuits, IEEE Solid-State Circuits Letters, and Sensors (MDPI). He serves on the Technical Program Committee of the IEEE VLSI Symposium.



Germán Mora Martín received the B.Sc. degree in nanoscience and nanotechnology from the Universitat Autònoma de Barcelona (UAB), Cerdanyola del Vallès, Spain, in 2017, the M.Sc. degree in physique, spécialité optique, photonique, and signal et image from Aix-Marseille Université, Marseille, France, and the M.Sc. degree in optics and photonics from Karlsruher Institut für Technologie (KIT), Karlsruhe, Germany, in 2019, through Europhotonics, an Erasmus Mundus Program. He is currently pursuing a Ph.D. degree at The University of Edinburgh, working

on high-speed 3D imaging with SPAD sensors. He develops algorithms for object detection and super-resolution using the data produced by SPAD sensors.



Alistair Gorman obtained a BSc in astrophysics and MSc in remote sensing and image processing from the University of Edinburgh, UK followed by a PhD in spectral imaging from Heriot-Watt University, UK. He was subsequently a Research Associate at Heriot-Watt University, and then at the University of Glasgow, UK, in the field of multi-spectral retinal imaging. A period in industry followed, first at retinal imaging company Optos, where he was lead research engineer, and more recently at LIDAR company Sense Photonics.

Dr Alistair Gorman is currently a Research Fellow at the University of Edinburgh and is part of the InLightenUs project, where he is undertaking research with time-resolving detector arrays for deep tissue imaging.



Hanning Mai received the M.Sc. and Ph.D. degrees in electronic and electrical engineering from the University of Edinburgh, U.K., where he is currently a Postdoctoral Research Associate and is developing single photon sensor for close range ToF imaging and biomedical applications. His research interests include CMOS SPAD sensors, CMOS circuit design, and biomedical imaging applications.



Francesco Mattioli Della Rocca (Member, IEEE) received the M.Eng. degree in electronic and electrical engineering from The University of Edinburgh, Edinburgh, U.K. in 2015, where he went on to complete a Ph.D. degree on single-photon avalanche diode (SPAD) CMOS image sensors, funded and jointly supervised by STMicroelectronics, Edinburgh.

Dr. Mattioli Della Rocca was a recipient of the IEEE ISSCC 2017 Student Research Preview Award in 2018, and is currently a Senior Digital Design Engineer at Sony European Technology Development Center.



Robert K. Henderson (Fellow, IEEE) received the Ph.D. degree in electronics and electrical engineering from the University of Glasgow, Glasgow, U.K., in 1990.

From 1991 to 1996, he was a Research Engineer with the Swiss Centre for Microelectronics, Neuchâtel, Switzerland. In 1996, he was appointed as a Senior VLSI Engineer at VLSI Vision Ltd., Edinburgh, U.K., where he worked on world's first single-chip video camera. From 2000 to 2005, as a Principal VLSI Engineer with the Imaging Division, STMicroelectronics, Edinburgh, he developed image sensors for mobile phone applications. In 2005, he joined The University of Edinburgh, Edinburgh, designing the first single-photon avalanche diode image sensors in nanometer CMOS technologies in MegaFrame and SPADnet EU projects. He is currently a Professor with the School of Engineering, The University of Edinburgh.

Dr. Henderson was awarded the prestigious ERC Advanced Fellowship in 2014.