ARL-TR-9780 ● SEP 2023

Demystifying Event-Based Camera Latency: Sensor Speed Dependence on Pixel Biasing, Light, and Spatial Activity

by Jonah P Sengupta

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-9780 ● SEP 2023

Demystifying Event-Based Camera Latency: Sensor Speed Dependence on Pixel Biasing, Light, and Spatial Activity

Jonah P Sengupta
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
September 2023	Technical Report	April 2023	August 2023

4. TITLE AND SUBTITLE
Demystifying Event-Based Camera Latency: Sensor Speed Dependence on Pixel Biasing, Light, and Spatial Activity

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Jonah P Sengupta

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
DEVCOM Army Research Laboratory ATTN: FCDD-RLA-TD Aberdeen Proving Ground, MD 21005	ARL-TR-9780

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES
ORCID ID: Jonah P Sengupta, 0000-0003-3555-7207

14. ABSTRACT
This report explores how various mechanisms affect the response time of event-based cameras (EBCs). EBCs are based on unconventional electro-optical IR vision sensors, which are only sensitive to changing light. Because their operation is essentially "frameless," their response time is not dependent to a frame rate or readout time, but rather the number of activated pixels, the magnitude of background light, local fabrication defects, and analog configuration of the pixel. A test apparatus was devised using a commercial off-the-shelf EBC to extract the sensor latency's dependence on each parameter. Under various illumination levels, results show a mean latency and temporal jitter can increase by a factor of 10 via configuring bias parameters. Furthermore, worst-case latency can exceed 1–2 ms even when 0.005% of the array is activated simultaneously. These and many other findings in the report hope to inform use of event-based sensing technology when latency is a critical component of successful application.

15. SUBJECT TERMS
Electromagnetic Spectrum Sciences; Photonics, Electronics, and Quantum Sciences; Terminal Effects; EO-IR sensing; event-based camera; optical characterization; sensor latency

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	40

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Jonah P Sengupta	(410) 278-5183

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

1. Introduction

Effective battlefield situational awareness (SA) requires the use of sensors that provide critical data to software that discriminate and localize threats, which in turn inform higher-level Warfighter-facing systems to convey tactical guidance. In many battlefield scenarios, time is one of the most critical dimensions to the SA task; thus, uniform and fast-sensor response time and low latency are paramount. Latency is a metric that can have multiple technical definitions depending on context: network, communication, signal processing. However, the overarching concept pertains to the temporal difference in data inception and arrival. Sensor latency can be broken down into four main components:

1) Sensory registration latency – comprises the amount of time between presentation of physical stimuli to the sensor, transducing into analog value, and encoding into a digital representation.

2) Intermediate hardware delays – from digital electronics and interface devices used to packetize, filter, and process this information.

3) Communication delays – from hardware used to propagate and receive the data via some universal protocol (GigE Ethernet, USB, Wi-Fi, long range [LoRA]). This is heavily dependent on protocol, connection, network traffic, and end-point distances.

4) Software delays – kernel-level processes that buffer communication data and algorithms that ultimately process data for a task.

Specific manifestation and magnitude of each of these delays is highly dependent on sensor modality and system configuration. Another aspect of sensor latency in the context of SA is jitter. As with most metrics, temporal response is not uniform but rather exists within a distribution. Jitter, otherwise known as timing accuracy, describes the variance of this distribution, which inherently places a bound of sensor response speed. Thus, analysis of a sensory solution needs to capture multiple aspects of the latency profile to fully characterize it for a specific application.

Sensing approaches such as lidar and radar offer high-response speed, selectivity, and sensitivity by actively illuminating the environment and leveraging a high signal-to-noise ratio (SNR). However, active sensing methods increase the probability of detection and exploitation; thereby, inadvertently endangering the platforms they are integrated within. In contrast, passive electro-optical infrared (EO-IR) sensing systems do not have an active emission signature and are an

improvement from a size, weight, power, and cost (SWaP-C) perspective but suffer from lower speeds and SNR.

Typical passive EO-IR approaches time integrate radiation incident on a pixel array, utilize scanning-based sampling methodologies to encode these values, and reset the pixels. This process is cyclic at a frequency indicative of the temporal resolution and latency of the sensor. Global shutter schemes expose the entire array simultaneously before reading out the entire array. Frame rates (the time per cycle) are therefore a function of integration time and the line time (time to readout a row) multiplied by the number of rows. Thus, worse-case sensor latency is a function of pixel location as longest delay will be seen for the last row. Rolling shutter schemes stagger the exposure by this line time and starting readout once integration has completed. This increases overall frame rate and decreases latency, but the occurrence of motion blur artifacts increases as a function of array resolution and line time. Furthermore, on a row level, there remains a pixel position dependent latency as the last row will have the longest delay between exposure and readout. Integrate-while-read (IWR), in contrast to the typical integrate-then-read (ITR), techniques seek to lower this latency further at the expense of longer exposure times. Nonetheless, state-of-the-art, high-speed camera technology have been able to produce cameras whose frame rates exceed hundreds of millions of frames per second (fps). However, these cameras utilize many high-speed serial ports and onboard dynamic-random-access memory (DRAM), whose operation is optimized for fast-recording operations and post-event, time-magnified playback.

A potential route for the advancement of EO-IR-aided SA capability is event-based imaging. The technology addresses the dynamic range and temporal sampling constraints present in EO-IR arrays by incorporating analog signal processing at the pixel level. In contrast to the frame-based approaches using traditional machine vision, event-based cameras (EBCs) comprise pixels and asynchronous interfaces that only transmit data upon temporal changes in incident light, otherwise known as temporal contrast. The analog front end logarithmically compresses the signal thus allowing for six decades of dynamic range and simultaneous high-/low-light imaging. In tandem, change detection circuitry greatly reduces bandwidth by an order of magnitude and simultaneously increases temporal resolution (tens of microseconds).[1–3] The waveform in Fig. 1 outlines the basic operation of an event-based sensor (EBS).

When there is a positive change in intensity (green curve in Fig. 1), the pixel transmits “ON” events while “OFF” events are propagated upon negative changes. The frequency of spike transmission is directly proportional to two factors: the requisite threshold to illicit an event and the gradient of input change. Output from the EBC is formatted in an address event representation (AER) where each event

datum describes spatial origination (x,y), polarity (p = ON/1 or OFF/0), and camera-relative time of capture (time stamp). However, since the camera is inherently "frameless" in its operation, the matter of latency is more complex. Due to the involved intra-pixel processing and asynchronous nature, response time is no longer a factor of global scan rates and deterministic pipeline delays, but rather of nonuniform transistor characteristics and stochastic interface arbitration.

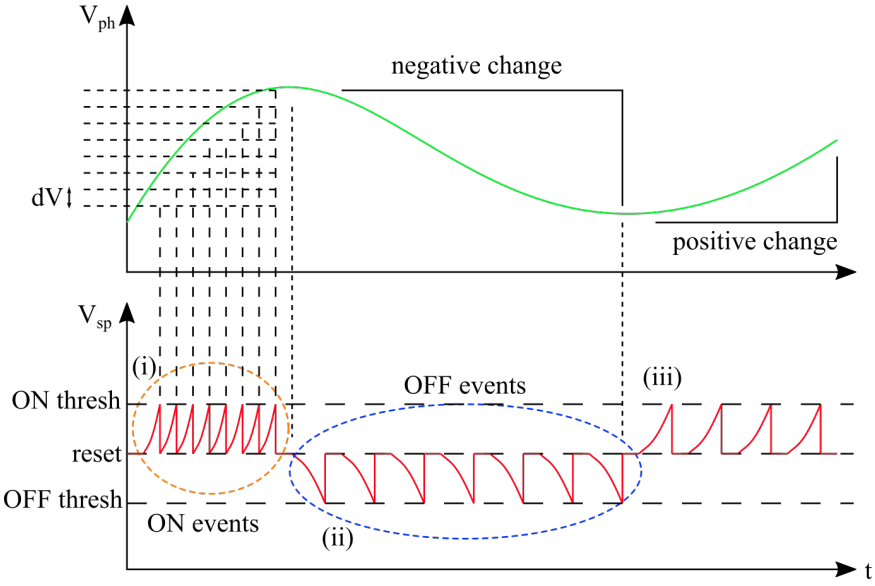

Fig. 1 Event-based operation: (top) logarithmic encoding of input intensity and (bottom) event activity with respect to ON and OFF threshold (dotted lines). Examples of (i) high-positive contrast, (ii) low-negative contrast, and (iii) low-positive contrast responses are shown.

This report outlines the mechanisms and methodical extraction of EBS latency. Prior work has briefly highlighted general methods for computing the metric, but none has presented a larger, in-depth study that provided response distributions while investigating the various causal mechanisms. Thus, multiple factors will be exploited to explore the full design space of variables that influence the EBS' response time. Section 2 will detail the specific mechanisms and theoretical underpinnings behind EBS latency. Section 3 outlines the experimental method utilized in this report to extract sensor and system latency. Section 4 describes the data analysis and reports the results of the study. Section 5 presents overall conclusions from the report and outlines next steps in the study of EBS latency.

2. Event-Based Sensor Latency

Other than jitter from the clocking networks on the chip, frame-based sensors have response times that are deterministic while being largely dependent on spatial location and readout scheme. In contrast, the asynchronous nature of EBSs introduces latencies that are dependent on bias conditions, chip irradiance, spatial location, and array activity.

2.1 Pixel Latency

The complex pixel circuitry needed to realize the event-based behavior incurs a nonuniform response time, which varies across the array. To understand this latter mechanism, a brief overview of the typical event-based pixel circuitry is needed. Most commercially available sensors utilize similar variants of a low-power analog signal processing chain.[1–3] A block diagram representation of this circuitry is shown in Fig. 2. This pipeline comprises three stages:

1) A transimpedance amplifier (TIA) logarithmically compresses the current from the light-sensitive photodiode into an analog voltage. This behavior is the mechanism behind the EBSs' ability to operate over six decades of light. Typical implementations utilize an adaptive, high-gain negative-feedback stage to extend bandwidth across a large range of input intensities. The output of this stage is represented in the top plot of Fig. 1 as V_{ph}.

2) A change-detection (CD) circuit filters out the DC component of the output from the TIA and amplifies the AC component. CD output is shown as the bottom plot in Fig. 1 as V_{sp}. The switch-capacitor amplifier, nominally used in most EBSs, has the AC gain set by a capacitive ratio.

3) A pair of comparators initiate output event transfers when the output from the CD exceeds one of two globally defined thresholds (horizontal dotted lines in Fig. 1). When either comparator is flipped from its resting state, an intra-pixel circuit is triggered and initiates the peripheral interface. This interface registers the spatial coordinates of the pixel and the direction of intensity temporal gradient as ON or OFF event polarity. Simultaneously, the internal circuitry resets the CD circuit. Upon release of the reset, the pixel can be triggered once again when a requisite positive/negative change is detected by the comparators. As shown in the profile of V_{sp}, this series of triggering and reset is dependent on the steepness of the intensity gradient (along with the other delays experienced from handshaking with the peripheral interface and auxiliary intra-pixel circuits).

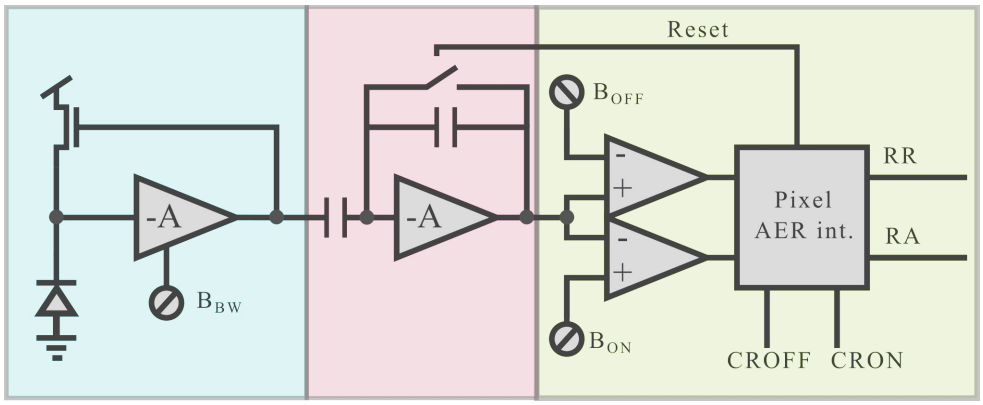

Fig. 2 Block diagram of pixel circuitry broken down into three subsections: 1) blue – TIA circuit, 2) pink – change detection circuit, and 3) yellow – comparators and intra-pixel AER circuitry. Small gray circles represent the analog biases used to tune pixel performance. Note: CROFF/ON = column request off/on, RR = row request, RA = row acknowledge.

Latency of the event-based pixel is largely governed by two mechanisms whose behavior is dependent on analog circuit biasing and incident illumination. These are the bandwidth of the TIA and the thresholds of the comparator. The CD amplifier and other pre-amplification circuits also play a role; however, because the formerly mentioned circuits are common throughout all EBSs and are programmable, those will be the only ones closely examined.

As mentioned, the TIA utilizes a high-gain, negative-feedback arrangement to pin the input photodiode cathode. In turn, this increases the input impedance by a factor proportional to the gain of the amplifier, which extends the bandwidth. Latency of an analog low-pass filter is directly related to this bandwidth as the corner frequency of the filter is correlated to the temporal delay. Amplifier gain can be directly controlled through a series of user-programmable analog biases. However, in most EBSs, there is a distinct trade-off between dynamic range and pixel latency. Moreover, the input impedance is also proportional to the photocurrent sunk by the photodiode thus the bandwidth has a light-dependent component. Overall, the temporal delay for the TIA can be approximated by time constant, τ_{TIA}:

$$\tau_{TIA}\ \alpha = R_{in}(I_{ph},\theta_{TIA})C_{in} \quad (1)$$

where R_{in} is the light- and bias-dependent input impedance, I_{ph} is the input photocurrent, and θ_{TIA} is the analog bias that controls the amplifier gain. The pixel-level latency is also directly influenced by the temporal contrast thresholds. As can be gleaned from Fig. 1, the threshold setting is directly linked to the frequency of "ON" and "OFF" events. Thus, temporal period is proportional to the ON and OFF comparator biases such that higher thresholds correspond to longer delays. In an EBS pixel, an ON event is triggered under the following condition:

$$\Delta V + V_{rst}(\theta_{CD}) = \theta_{on} \quad (2)$$

$V_{rst}(\theta_{CD})$ represents the the voltage after pixel reset. The analog level of this signal is dependent on the user-defined, CD amplifier bias. Thus, the requisite change in V_{sp} is

$$\Delta V = \theta_{on} - V_{rst}(\theta_{CD}) \, . \quad (3)$$

Dividing this quantity by the slew of the output, $S_{sp} = \Delta V / \Delta t_{sp}$ with units (V/s), the comparator temporal delay, τ_{CMP}, is quantified as

$$\tau_{CMP} = (\theta_{on} - V_{rst}(\theta_{CD}))/S_{sp} \, , \quad (4)$$

where S_{sp} is proportional to the gain of the CD amplifier and the photocurrent's rate of change such that a higher gradient will result in smaller delays. Equation 4 shows how the difference between the potential after pixel reset and the threshold directly impacts the temporal delay.

In total, pixel latency can be approximated as the sum of the TIA and comparator delays,

$$\tau_{PIX} = \tau_{TIA} + \tau_{CMP} = R_{in}(I_{ph},\theta_{TIA})C_{in} + (\theta_{on} - V_{rst}(\theta_{CD}))/S_{sp} \, , \quad (5)$$

which succinctly outlines the bias and light-dependent latency of the EBS pixel.

2.2 Interface Latency

As alluded to in Section 2.1, event transmission is solely dictated by a peripheral interface once the comparators have been triggered. This propagates the occurrence of the event activation through a channel that serializes transmission by arbitrating access and encoding each event with its spatial address and polarity. Such a channel is known as an AER interface.[4]

Fundamentally, event activation is asynchronous as local transients drive data transmission rather than the global clocking schemes seen in frame-based sensors. Therefore, a common situation arises where multiple pixels are activated simultaneously. The probability of this occurrence increases proportionally to sensor resolution as sub-sampling the same scene with higher spatial fidelity leads to more simultaneous event activations and a higher likelihood of spurious transmissions. Pixels are organized such that requests to the interface are shared among columns and rows. When a pixel is triggered, it first asserts the common row request. The AER interface then must arbitrate between the various rows that have requested access to the output channel. Once a row has been selected via acknowledgement from the AER, all pixels in that row that have triggered raise a

column request. At this time, the activated pixels are reset, and column requests are latched. The row request and subsequent column requests are encoded into digital words that describe each activated pixel address and event polarity. As the data from the selected row is sent off the chip, the AER interface acknowledges another asserted row and continues reading out events.

Much like the EBS pixel, AER interfaces too have evolved over time. Initially, column requests were acknowledged on a pixel-wise basis, which greatly limited the throughput of the channel. Including the mentioned capability to latch column requests greatly increased this metric at the expense of including extra circuitry to hold prior states. Latency of this type of interface can be analyzed under two circumstances: high- and low-array activity. In the former case, the mean event latency is proportional to the number of activated rows, mean number of activated pixels in each row, and the time needed by the interface to cycle through each event in a burst. In the light-load scenario, mean event latency is dictated by time needed to acknowledge the row request by the arbiter. An in-depth analysis of these scenarios founded in queuing theory and experimental chip results can be found elsewhere.[5,6]

As alluded to, newer, high-resolution sensors have a greater probability to yield array activity that quickly converges to the high-load scenario seen for burst-mode AER interfaces. Thus, burst-mode interfaces can only expedite cycle times to accommodate larger activities, which quickly becomes unscalable for megapixel arrays. However, modern EBSs have adapted to utilize spatial compression and column-wise scanning schemes to accommodate.[7,8] Because of this approach, mean latency under high load is synonymous to numbers seen from frame-based counterparts.

2.3 Field-Programmable Gate Array Latency

Beyond the AER interface, an external programmable logic device (PLD) or field-programmable gate array (FPGA) is typically utilized to coordinate data transfer from the EBS, append a local time stamp, and buffer the data prior to communication to a remote or local device. Data transfer is coordinated between the EBS and FPGA using a four-phase handshake protocol. After encoding the pixel address as described previously, the EBS asserts a request to the FPGA. After ensuring the data is valid, the FPGA asserts an acknowledgement to the EBS. This prompts the EBS to lower the request, which in turn causes the FPGA to lower the acknowledgement and complete the protocol. FPGA latency is set by the time needed to write a valid word to the internal buffer. Along with the speed of the clock, this is determined by the number of cycles needed by the device to

acknowledge the initial row transaction from the EBS, detect a valid column address, and register the arrival time using a local counter. This second operation is greatly dependent on AER protocol as delay-insensitive or compressive schemes (i.e., the ones adopted by modern EBS) may require a large amount of circuitry to decode addresses.[6–8]

3. Experimental Setup

An experimental apparatus and experimental procedure were designed to explore the EBC latency mechanism described in Section 2. The intention was to isolate each tested parameter, extract response time, and map out correlative behavior. The means to achieve this goal was to present a uniform, repeatable stimuli to the EBC whose operation can be programmatically altered in a controlled, well-characterized environment. Each trial consisted of utilizing a unique pixel bias profile, background illumination level, sampling scheme, and region of interest (ROI). In total, 100 points were collected per trial such that each point comprised 6–10 response times recorded on a 1-Hz frequency per pixel. By collecting data in this fashion, subsequent analysis could yield a mean pixel latency (average across all latency measurements per point), mean trial latency (average over all mean pixel latencies), intra-pixel jitter (variance across the latency measurements per point), and inter-pixel jitter (variance across the trial over all points).

3.1 Equipment

In the event-based regime, there is little-to-no array activity when the EBS is presented a static stimulus. Since event generation only occurs upon presentation of a dynamic scene, the latency can be extracted by computing the difference between the times of stimulus presentation and the arrival of the first event from the pixel under test. Figure 3 provides a diagrammatic overview of the experimental setup used to produce the well-calibrated and accurate stimuli for extraction of this metric in a synchronized manner. A host PC initiates the experiment by configuring and priming the EBC before enabling the trigger controller. The EBC's analog biases and ROI are parameterized with respect to the desired trial in the design of experiment. The trigger controller is used to provide a simultaneous pulse to the optical source and the EBC. Optical stimulation is provided by a pair of 12-V LEDs whose anode voltages are both provided by a common power supply. One LED constitutes a stable background illumination while the other has its cathode shunted to ground through a power transistor whose base is controlled by the trigger controller (as shown in Fig. 3). By illuminating the EBC in this manner, a quasi-uniform, fast-rising stimulus with 100% temporal contrast is achieved. Moreover, magnitude can be scaled electrically through the power source or optically through

neutral density (ND) filters in front of the EBC. To note, the LEDs were oriented to ensure full coverage of the array by both sources, but uniformity illumination is not guaranteed. Future experiments will utilize optical apparatus (e.g., an integrating sphere) to improve on this issue.

The same signal that powers the foreground LED is also propagated to the EBC. This is captured through the camera's trigger in port, which is used to synchronize the EBC with external sources. Once the camera hardware detects an edge on this input, it injects a special event into the output stream that contains the arrival time and polarity (1 encodes a rising edge, while 0 encodes a falling edge). Hence, the time of stimulus trigger is now recorded with respect to the same clock used to record the arrival time of each event. Thus, latency can be directly computed by taking the difference between the camera-synchronized event and trigger times. To note, both ON and OFF latency can be directly observed using this approach as rising and falling edges of the trigger signal are made distinct by the encoded polarity.

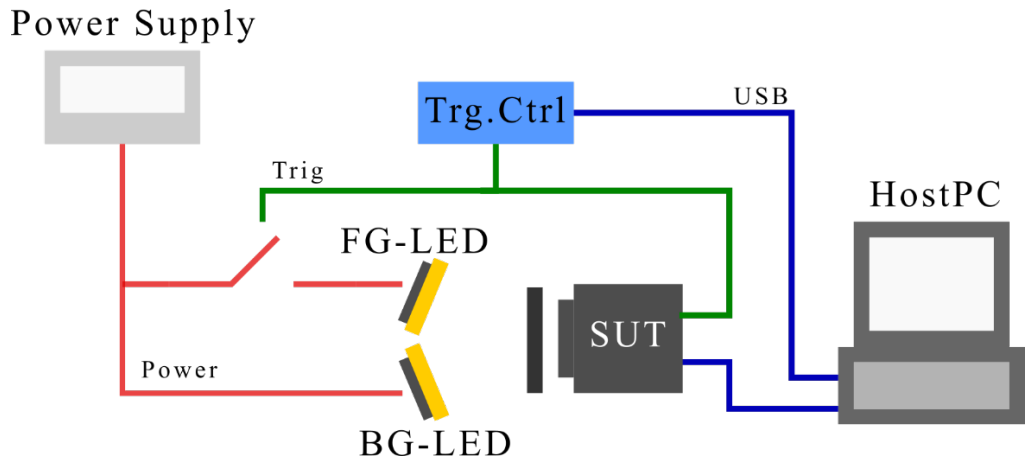

Fig. 3 Block diagram of experimental setup: (Red lines) 12-V LED supply, (Green lines) trigger signal from the µC, and (Blue lines) USB communication. The "SUT" represents the event-based Sensor Under Test. Controlled background (BG) and foreground (FG) illumination is realized using the BG- and FG-LED devices, respectively. This signal is further attenuated by the optical density, represented with the black shape in front of the SUT.

3.1.1 Commercial Off-the-Shelf Event-Based Camera

The EBC used for the experiment was the SilkyEvCam-VGA, which is equipped with the PPS3MVCD sensor from Prophesee.[9] As the name suggests, the sensor has a 640 × 480 resolution, comprising event-based pixels whose operation was described in Section 3.1. The sensor output is shuttled to the host PC via intermediary FPGA and microcontroller. An FPGA first converts the data from an AER to Mobile Industry Processor Interface (MIPI) format before being reformatted into USB packets by the microcontroller. The AER to MIPI conversion

incurs a deterministic latency of 2 ms while the USB transactions could experience wait times ranging from 0.2 to 1.0 ms depending on host PC activity.

Camera configuration and data streaming were accomplished using the Metavision SDK 3.1.2 on Microsoft Windows 11. A custom software stack was designed in Python on top of the provided application programmer interface (API) to rapidly configure hardware signals, camera modes, ROIs, and bias settings, thereby extending the capability of the existing commercial off-the-shelf (COTS) software, Metavision Studio. Trigger and event data are logged simultaneously through the API. Upon detection of a trigger event, event data was recorded for 1 ms and stored in a first in/first out (FIFO) buffer. Subsequent triggers initiate further recordings that are appended to this buffer. This capture routine is ultimately halted after a requisite number of triggers have been detected. Output buffer data was then saved in a comma-separated value (CSV) format to the PC disk.

3.1.2 Trigger Controller

An external microcontroller was used to trigger the optical stimulus in synchrony with the hardware trigger to the EBC. A BluePill development board, which houses the STM32F103 microcontroller, was chosen to serve as the trigger controller. A USB to UART FTDI chip communicated between the host PC and the board. The interface was leveraged to implement a serial console on the microcontroller. This allowed for execution-time initiation of the stimulus in addition to configuration of pulse width and period. The microcontroller used these provided parameters to program the clock rate and threshold of a high-resolution timer peripheral whose output flag was used to control the base of the power transistor as previously described.

3.2 Experimental Procedure

The main thrust of this report was to explore the various hardware-level mechanisms behind event-based sensor latencies. Thus, there were five parameters that were explored in tandem: background illumination, pixel biasing, number of active pixels, location of active pixels, and inter-sensor variances. Background illumination was controlled by modulating the LED power supply and attenuating the output signal using ND filters. Pixel biasing was configured through the API using one of four unique profiles per trial. These included the nominal (nom), a high-sensitivity (minimum-threshold), a low-sensitivity (maximum-threshold), and wide bandwidth (high-gain) parameterization. Nom bias setting represents the default settings for the camera provided by Prophesee. Max and min settings configure the camera biases to their lowest and highest threshold settings. Wide bias settings configured the TIA's front end to detect higher-temporal frequencies.

By adjusting contrast sensitivity and bandwidth independently, the predominant mechanism for low-latency operation can be explored. Exact settings of the pixel biases can be seen in the Appendix, Table A-1.

The API was also used to configure the locations of ROIs. Each ROI comprises four elements that describe its origin and dimensions. Pixels outside of this region are disabled and cannot transmit events through the AER. This can be exploited to observe singular pixel characteristics while also exploring the effect of relative location, active rows, and active columns on average latency. Finally, multiple sensors were tested to characterize the role of inter-chip variation.

4. Results

Five trial parameters were exploited to explore their respective relationship to EBC latency. These include background illumination, pixel biases, sensor under test (SUT), ROI origin, and ROI dimensions. Testing was further segmented into the trials, which involved capturing data from individual pixels per trial to those that analyzed larger ROIs per sample. Analysis leveraging data from the former tests helped illustrate an absolute performance bound on latency with respect to each design parameter (i.e., illumination, bias, and SUT). On the contrary, stimulating regions of the sensor array simultaneously constructs a relation between array activity and average response time. This analysis provides insight into how row and column orientation factor into EBC latency.

4.1 Single-Pixel Testing

Spatially local and array-level trends in response time and jitter were extracted by varying individual pixel biases, illuminations, and the SUT. This required the use of a sampling scheme to select which pixel to test. Two such schemes were deployed. An *ordered* approach was first taken to sample pixels within a 10×10 ROI in the center of the sensor array. This 100-point data set allowed for pair-wise testing of extracted latency metrics across multiple experimental variables, thus enabling statistically powerful hypothesis testing. The second scheme randomly sampled 100 points from the 307,200-member array to test per trial. In the presence of potentially spatial-correlated results, this random method constructed a better estimate of a sample mean for the population. Performing the same trials for each sampling method also enabled the observation of this local spatial correlation. In total, there were two sampling methods, two SUTs, four biases, and three background-illumination settings tested. The SUTs used for the characterization were both enclosed in two SilkyEvCam modules whose only difference was the sensor. LED power and ND filter OD factors were modulated to achieve

background-irradiance levels of 6.0, 0.6, and 0.06 W/m^2. These numbers were measured and verified using a Newport Optical Power Meter paired with a calibrated visible photodetector (SL-81).

4.1.1 Ordered Spatial Sampling

Figure 4 depicts the results of the trials for a subset of points extracted using the ordered sampling scheme. Ten points sampled along the diagonal of the ROI were overlaid on the trial average to provide a coarse visual representation of the inter-pixel jitter. Error bars on each individual pixel sample represent the intra-pixel jitter. These two variances outline the timing accuracy of a singular point and the variance seen for all points in a single trial. Each pixel is stimulated four times, yielding a temporal-latency computation on the rising and falling edge of the intensity pulse. The mean of these measurements constitutes the mean-pixel latency while the deviation represents the aforementioned intra-pixel jitter. These results can then be averaged to comprise the mean-trial latency for a specific combination of a bias and irradiance sample.

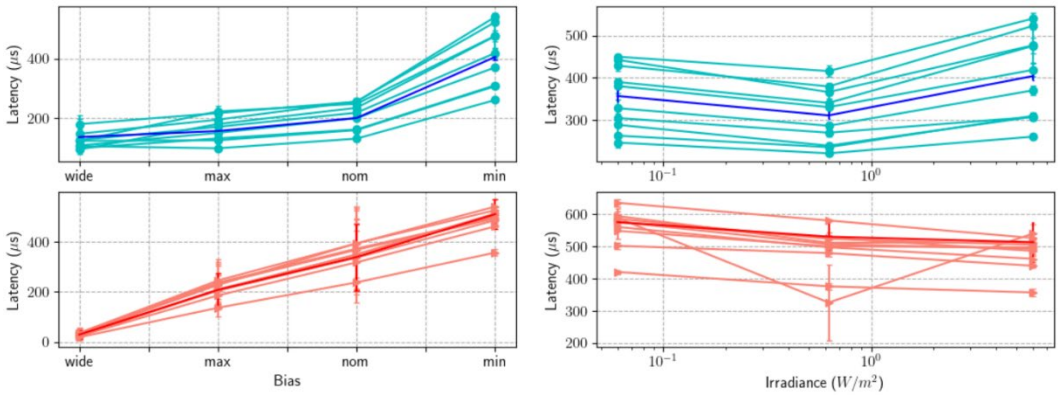

Fig. 4 Mean latency in microseconds extracted from 10 × 10 ROI in center of EBC. Left plots correspond to sampled trials at irradiance = 0.06 W/m^2. Right plots depict trial data with bias = *min*. Bottom red = OFF data; Top blue = ON data. Light colors correspond to data sampled on the diagonal of the ROI; dark colors depict the average response of the 10 × 10 ROI. Error bars denote trial jitter (or mean jitter). R^2: 0.79, 0.98, 0.94, 0.26 (counterclockwise from top left).

As expected from prior results and theoretical pixel operation, there is a clear monotonically increasing relationship between pixel sensitivity and temporal latency. Furthermore, a linearly decreasing relationship was observed in the OFF temporal response and irradiance. Again, this aligns with what was previously shown. However, the ON response did not exhibit the same response. This could be attributed to the smaller range of background-irradiance values tested. As seen from prior work, an operating regime under dim background illumination where latency is invariant exists.[1,3] This non-decreasing behavior when extracted could

also fall within that regime. Future tests will exploit a larger array of background-illumination levels to explore this result further.

Figure 5 provides a spatial map of the mean temporal response within the tested ROI for a certain trial. In the presence of large spatial correlation, one could expect to observe periodic, sub-regions of similar activity. However, the contour plot in Fig. 3 offers little such insight suggesting lack of a strong spatial correlation in temporal response. This suspicion is confirmed later when comparing the results between ordered and random sampling trials.

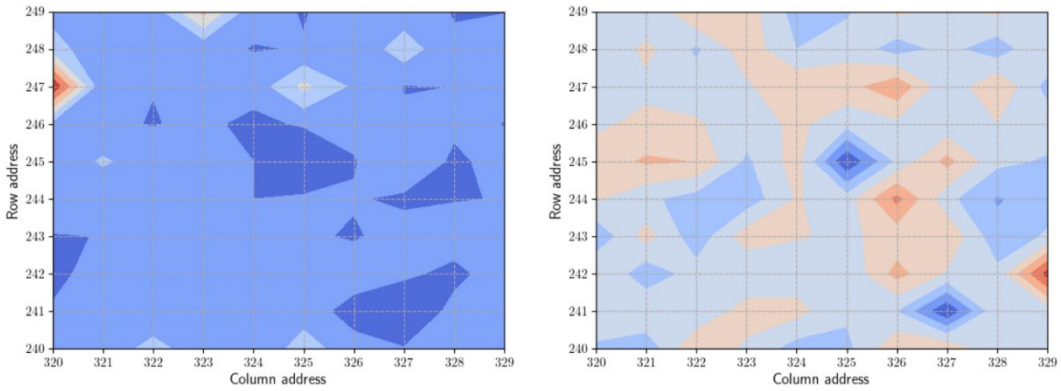

Fig. 5 Mean latency contour plot: color coding designates mean latency values in the 10 × 10 ROI. Dark blue corresponds to small values, while dark red denotes pixels with slow responses. Left plot = ON event data; right plot = OFF event data.

Figure 6 portrays the distribution of mean pixel latencies over a subset of trials with biases and irradiances overlaid over one another. Like the relation shown in Fig. 2, the inter-pixel jitter increases in tandem with the mean response when sensitivity is lowered. The mean of the ON latency distribution sees little-to-no shift when the irradiance is modulated.

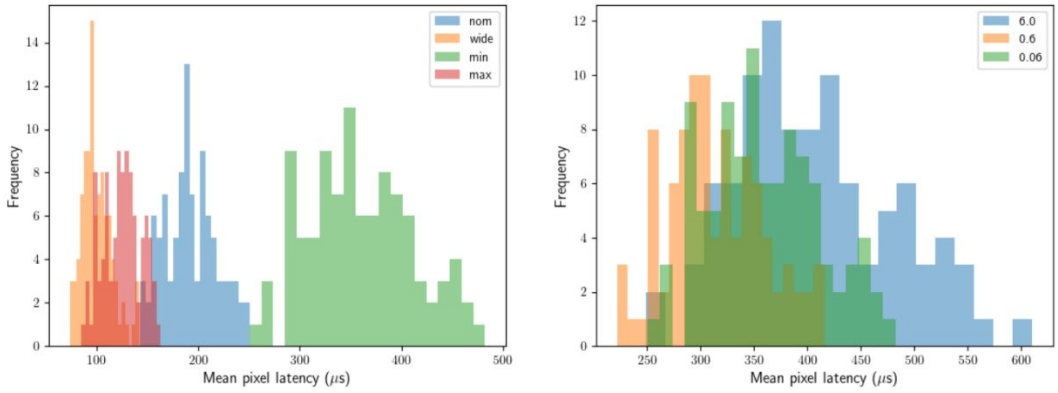

Fig. 6 Mean latency histograms depicting distribution of ON event responses within ROI for each bias (left) and irradiance (right) value. These histograms reflect the same trials depicted in Fig. 2.

A one-sided Wilcoxon ranked-signed test was used to provide proof of statistical significance. A non-parametric test was chosen because a Shaprio test confirmed that the ordered data was not Gaussian. Tables 1 and 2 show the test statistic and p-value for each paired test performed for the bias and irradiance tests. There is strong evidence to suggest that pixel biases play a large role in the resulting sensor response time. Moreover, given the R^2 values in Fig. 2 and the level of statistical significance, these paired tests also provide evidence that biasing influences temporal response more significantly than background illumination.

Table 1 Wilcoxon test statistic (T) and p-value (p) for each pair of bias samples using the ordered sampling scheme over the set of irradiance values

Irradiance (W/m^2)	Polarity	Nom,wide (T/p)	Nom,max (T/p)	Nom,min (T/p)
6	ON	4459/3.31e-14	4752/6.27e-18	0/6.09e-18
	OFF	4753/6.09e-18	4753/6.08e-18	28/1.40e-17
0.6	ON	4465/1.90e-17	4465/1.90e-17	0/1.90e-17
	OFF	4465/1.90e-17	4465/1.90e-17	1/1.96e-17
0.06	ON	4851/4.16e-18	4851/4.16e-18	0/4.16e-18
	OFF	4851/4.16e-18	4851/4,16e-18	0/4.16e-18

Table 2 Wilcoxon test statistic (T) and p-value (p) for each pair of bias samples using the ordered sampling scheme over the set of irradiance

Biases	Polarity	0.6,6.0 (T/p)	0.6,0.06 (T/p)
Nom	ON	192/0.99	4273/4.76e-17
	OFF	94/1.74e-10	148/2.29e-16
Wide	ON	256.5/0.999	632.5/6.32e-11
	OFF	4655/1.36e-14	94/4.75e-17
Max	ON	185/0.99	94/4.75e-17
	OFF	4378/2.2e-12	188/7.16e-16
Min	ON	190/0.99	94/4.76e-17
	OFF	3930/1.9e-7	631/6.12e-11

4.1.2 Random Spatial Sampling

Figure 7 provides the distribution of the sampled pixel addresses for a specific trial using the random sampling scheme. These addresses span the entire range of rows and columns with no bias for a certain region. The mean column and row addresses were 302 and 246, respectively; they are relatively close to the actual mean addresses and center of the array (320,240).

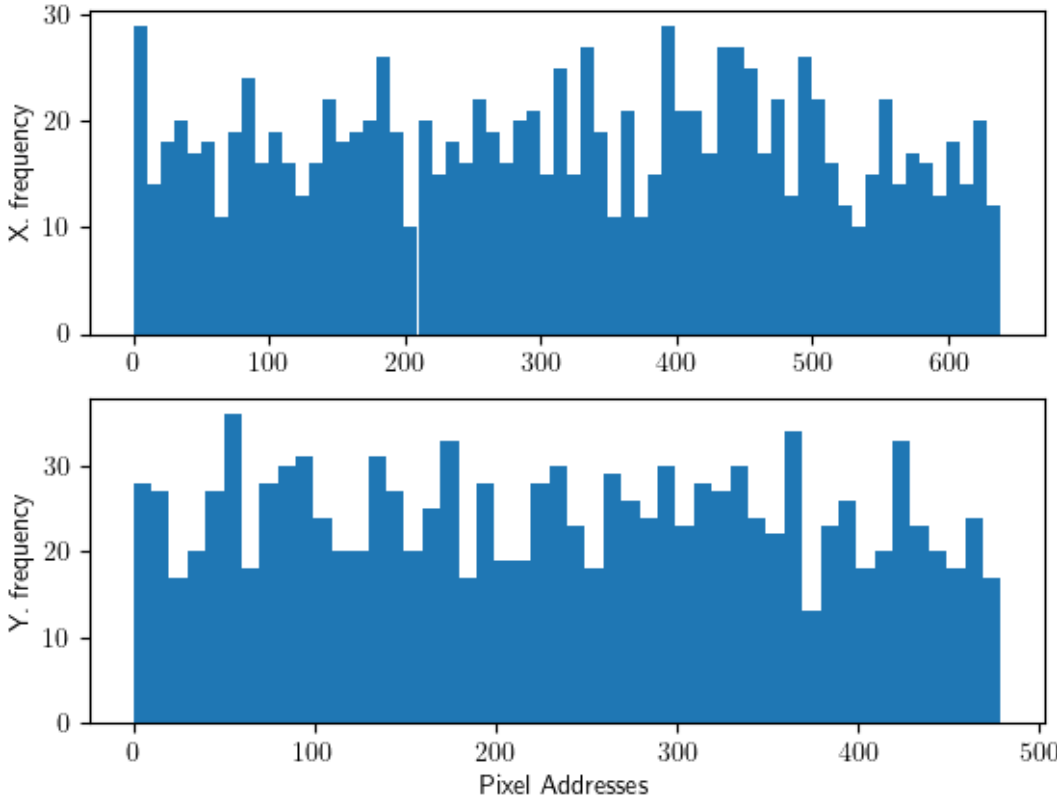

Fig. 7 Example distribution (irradiance = 0.06 W/m^2, bias = min) of addresses utilized using the random pixel sampling scheme for global latency extraction. Top: Sampled x-address frequencies; Bottom: sampled y-address frequencies.

Figures 8 and 9 provide a succinct pictorial overview of the mean trial latency and inter-pixel jitter for the random trials. Mean trial latency is extracted by averaging the 100 average pixel latency results over the distinct trial. Error bars are computed using the standard deviation of these 100 latencies. This inter-pixel latency is shown in the right sub-figures within Figs. 6 and 7. Many of the conclusions derived from the ordered sampling trials are validated when comparing the random sampling results. Again, there is a consistent relationship between pixel sensitivity and temporal response. Moreover, jitter hypothesis posited after viewing the histograms in Fig. 4 is confirmed: inter-pixel jitter also monotonically increases with respect to a decreased pixel sensitivity. Once again, an expected, monotonically decreasing relationship between OFF trial latency and irradiance levels is seen across all bias configurations. However, inter-pixel jitter does not similarly decrease. This unforeseen trend reinforces the need to repeat the experiment using a larger swathe of background-irradiance levels. The four metrics and their confidence intervals are summarized for the 12 random trials in Table 3. Confidence intervals were extracted by bootstrapping the metric of interest (mean or standard deviation) from the 100 average latencies and intra-pixel jitter values per trial.

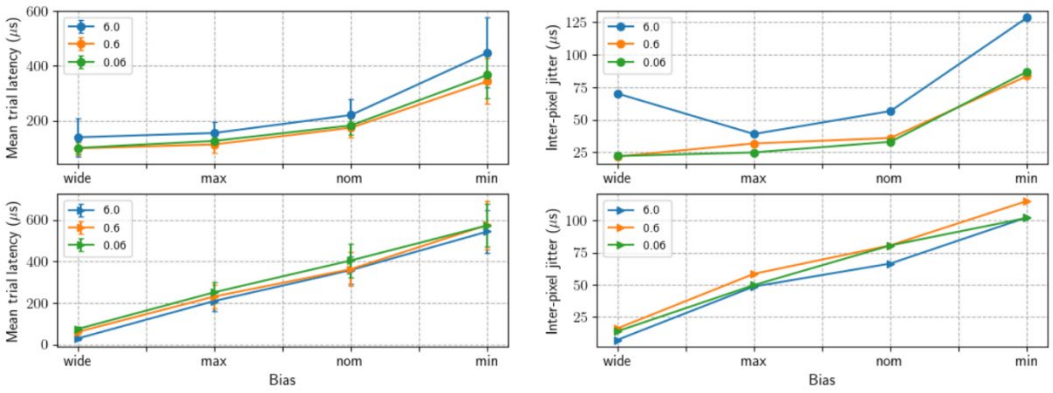

Fig. 8 Average mean latency and average inter-pixel jitter for each trial with respect to configured bias. Top: ON event response; Bottom: OFF event response. Trial irradiance denoted in inset legend. Error bars denote average jitter.

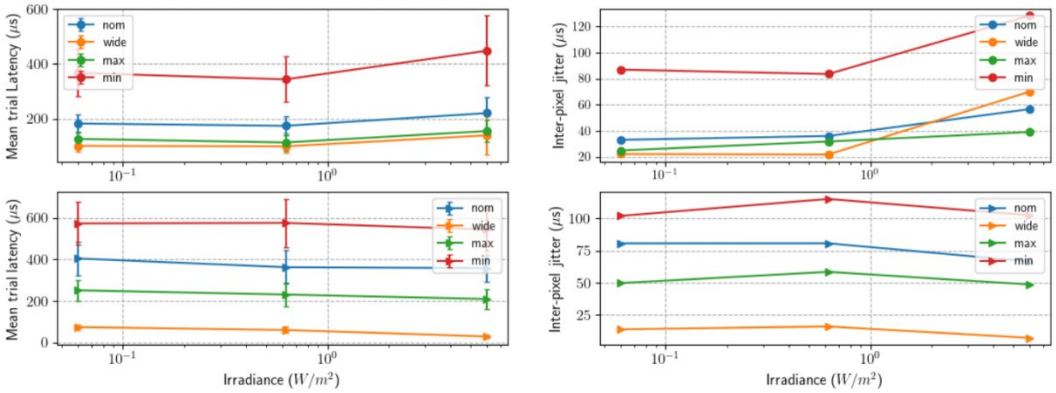

Fig. 9 Average mean latency and average jitter for each trial with respect to BG irradiance. Top: ON event response; Bottom: OFF event response. Trial biases denoted in inset legend. Error bars denote average jitter.

Table 3 Mean trial latency, intra-pixel jitter, inter-pixel jitter, and intra-pixel jitter standard deviation for ON and OFF events per trial. All confidence intervals were bootstrapped from the 100 samples per trial. All metrics are in units of microseconds.

Bias	Irradiance (W/m^2)	Mean average ON latency	Mean ON intra-pixel jitter	Inter-pixel ON jitter	Std. ON intra-pixel jitter
nom	6.04	209.39,231.20	4.27,5.45	48.41,69.47	2.60,3.66
wide	6.04	129.16,161.21	13.45,44.14	38.71,137.14	18.40,127.02
max	6.04	147.16,162.59	3.95,4.88	34.24,45.73	2.11,2.77
min	6.04	425.16,476.60	9.94,12.89	108.09,158.00	6.30,9.72
nom	0.62	167.14,181.45	4.96,18.83	30.13,48.93	2.83,61.70
wide	0.62	94.62,103.19	2.58,3.51	18.20,27.13	1.90,2.90
max	0.62	108.13,121.38	3.70,4.67	20.98,59.69	2.01,3.04
min	0.62	328.57,361.91	8.00,10.27	68.78,106.71	4.68,7.73
nom	0.06	175.89,188.85	6.28,8.10	29.33,38.27	3.96,5.78
wide	0.06	96.11,104.82	4.91,6.71	19.39,25.78	3.78,5.66
max	0.06	120.88,130.67	5.73,7.15	22.25,28.00	3.02,4.31
min	0.06	351.13,385.15	13.35,16.89	75.23,102.69	7.75,10.70
Bias	**Irradiance (W/m^2)**	**Mean average OFF latency**	**Mean OFF intra-pixel jitter**	**Inter-pixel OFF jitter**	**Std. OFF intra-pixel jitter**
nom	6.04	344.35,370.34	136.24,151.90	56.40,89.34	34.11,49.43
wide	6.04	27.51,30.46	17.94,20.57	5.71,8.22	5.05,7.60
max	6.04	200.45,219.31	65.33,73.71	41.45,57.82	17.44,26.18
min	6.04	523.16,563.73	60.03,124.05	88.59,124.22	131.04,191.57
nom	0.62	346.31,378.35	127.76,147.84	70.25,100.75	44.27,61.83
wide	0.62	58.10,65.06	24.17,27.83	9.25,30.59	7.17,12.05
max	0.62	220.54,243.59	63.86,72.39	46.59,86.58	18.98,25.46
min	0.62	553.66,599.74	52.61,117.33	97.26,142.41	125.16,198.37
nom	0.06	389.59,421.42	153.26,178.67	66.23,103.53	48.01,94.34
wide	0.06	71.66,76.91	26.18,31.22	11.75,15.84	10.09,18.17
max	0.06	241.64,261.16	75.29,84.69	43.38,58.20	20.59,27.48
min	0.06	552.17,592.59	37.51,92.19	88.48,122.65	95.19,175.25

Figure 10 furthers the discussion on the dependence of latency on pixel location. The distribution of ON and OFF average pixel latency is shown using nominal bias and background irradiance of 6 W/m^2. The results of the trial for the ordered and random sampling schemes are overlaid. From visual inspection, it can be suspected that each distribution is sampled from the same population thus indicating there is little relationship between local address and temporal response. A non-parametric, Kruskal test supports this hypothesis: p-values of 0.96 and 0.86 provide further proof that the ordered and random samples belong to the same underlying distribution.

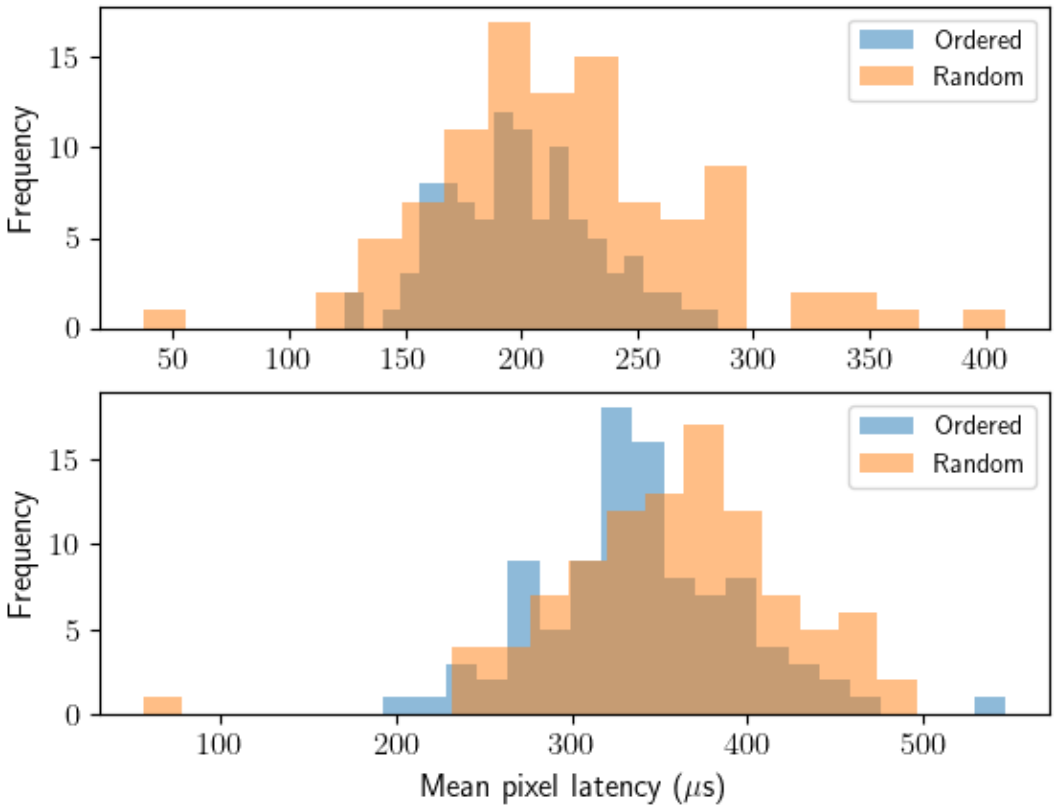

Fig. 10 Distribution of ON (top) and OFF (bottom) mean latencies for nominal bias, irradiance = 6 W/m^2 with respect to each utilized sampling scheme (legend inset). Areas common to both the ordered and random sample distributions appear brown-colored in the plots.

Figure 11 presents the results of collecting data from two different EBC using the ordered and random sampling schemes. After performing the Wilcoxon and Kruskal tests on these data sets, there seems to be sufficient evidence to suggest that there is a significant variation in local performance based on intra-sensor variation (p-value = 8.32e-18), but little effect on global latency metrics (p-value = 0.19). Future endeavors will include sampling data from other EBCs with more design parameters (biases, irradiance) to further solidify this preliminary finding.

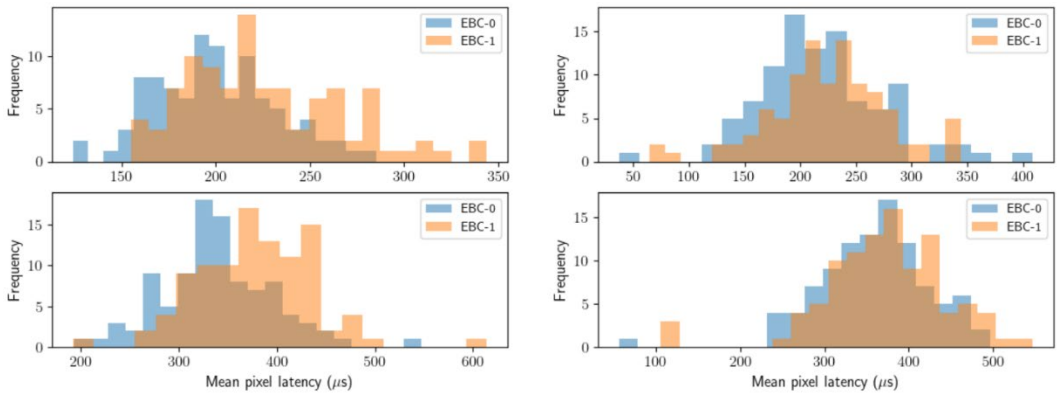

Fig. 11 Distribution of ON (top) and OFF (bottom) mean latencies for nominal bias, irradiance = 6 W/m^2 and utilized sampling scheme (ordered = left, random = right). Responses were plotted with respect to the EBC under test (legend inset).

4.2 Multi-Pixel Testing

Event-based cameras utilize asynchronous, AER interfaces to communicate pixel data on demand. Under static conditions that illicit sparse array activity, latency is bounded by the individual pixel response. However, as scene stimuli affects a large region within the array, the delay introduced by serializing all the data through the common interface starts to become prominent. A series of trials were conducted using the same apparatus as the single-pixel experiments to observe this effect. Three ROI parameters were explored: dimension, size, and location. Row and column dimensions of 1, 4, 64, 256, 480, 640 and 4, 64, 256, 480 were used to parameterize the trial ROI. In addition, each ROI was situated at the bottom right, bottom left, top right, and top left of the array. Larger ROI could only be accommodated in certain quadrants; thus, the total number of trials was 99. A nominal bias configuration and background irradiance of 6 W/m^2 were used for each trial. See Fig. 12.

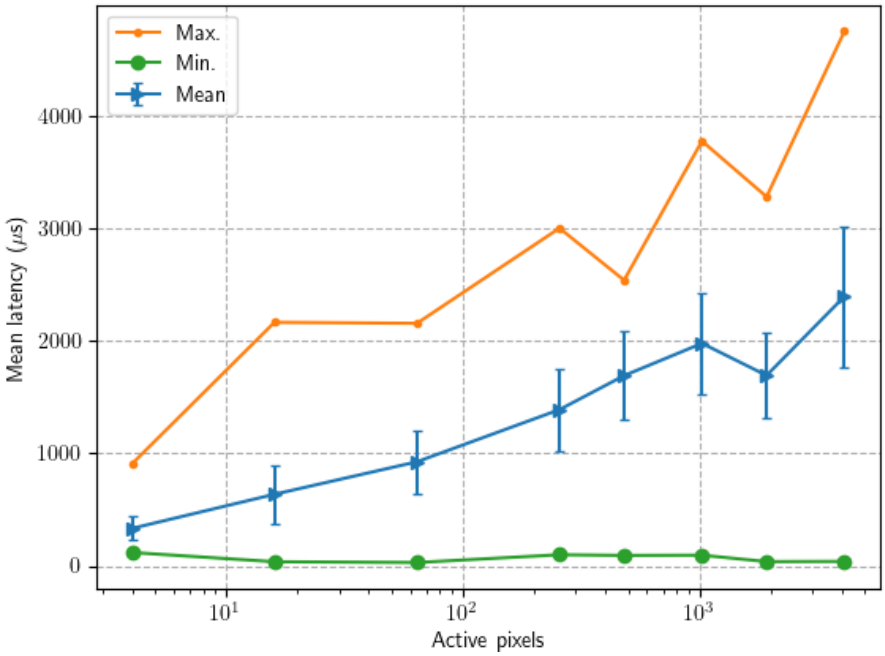

Fig. 12 Mean latency of analyzed samples with respect to total number of pixels enabled per trial. Error bars represent the variation of latency across trials with the same amount of active pixels.

After configuring the ROI, the array was stimulated 10 times for each trial. The temporal activation of the trigger signal and resulting events were correlated to compute the latency of each event per signal edge. Despite the presence of a bright transient stimuli, not every enabled pixel was triggered due to the non-uniformity in temporal contrast discrimination and latency.[1–3] Thus, trial metrics were extracted using a weighted-average of statistics gathered after each triggered activation. Weights were computed by finding the maximum events triggered after a stimulus edge, resolving other results whose total amount of activated events are close to this maximum, and normalizing the mean and standard deviation of each result by the maximum. Using this method, a larger sample can be cultivated to draw conclusions about spatial relationships; however, low-activity edges, whose results could correlate to outlier behavior, are discarded. A comparison of three different sampling methods is shown in Fig. 13.

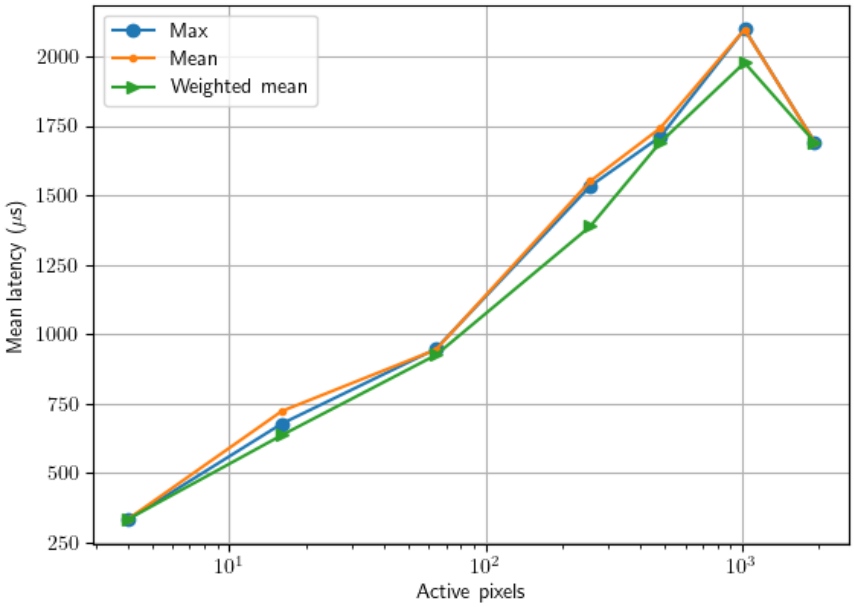

Fig. 13 Average mean latency plotted with respect to total pixels enabled per trial and sampling scheme used to extract spatial trial statistics

The results in Fig. 12 were compiled using the described technique. Figure 12 shows the relationship between the total number of active pixels in the ROI and the mean latency of trials with the same number of enabled pixels. Error bars denote the standard deviation between those trials. At low amounts of active pixels (ROI size = 4), the mean latency is similar to that seen in the single-pixel trials for the same parameterization. However, average latency grows as the ROI size increases and more pixels attempt to request access. This is where the AER interface transitions from the light-to-heavy load scenario introduced in Section 2 and described in further detail elsewhere.[2,3] This is also reflected in the maximum latency seen across all triggers for a specific trial. As seen in the orange curve, worst-case latency can exceed 2 ms even under smaller loads (ROI dimension = 16). Best-case response time is maintained across all trials as it is fundamentally still bounded by the fastest pixel latency in the ROI rather than the waiting time for a pixel to be serviced by the interface.

Figure 14 outlines the dependence of ROI dimension on mean pixel latency. Two subsets of the collected data were observed: ROI with one row enabled and varied column dimension, ROI with varied row dimensions and four columns enabled. These were then compared by computing the size of each ROI and plotting the mean latency per trial. For light load, the dimensions of the ROI do not affect the mean response time. However, there seems to be a noted effect when the number of active pixels exceed 256. For this point, the interface needed to arbitrate requests from 64 unique rows, which is slower than reading out the same number of pixels in a singular row.

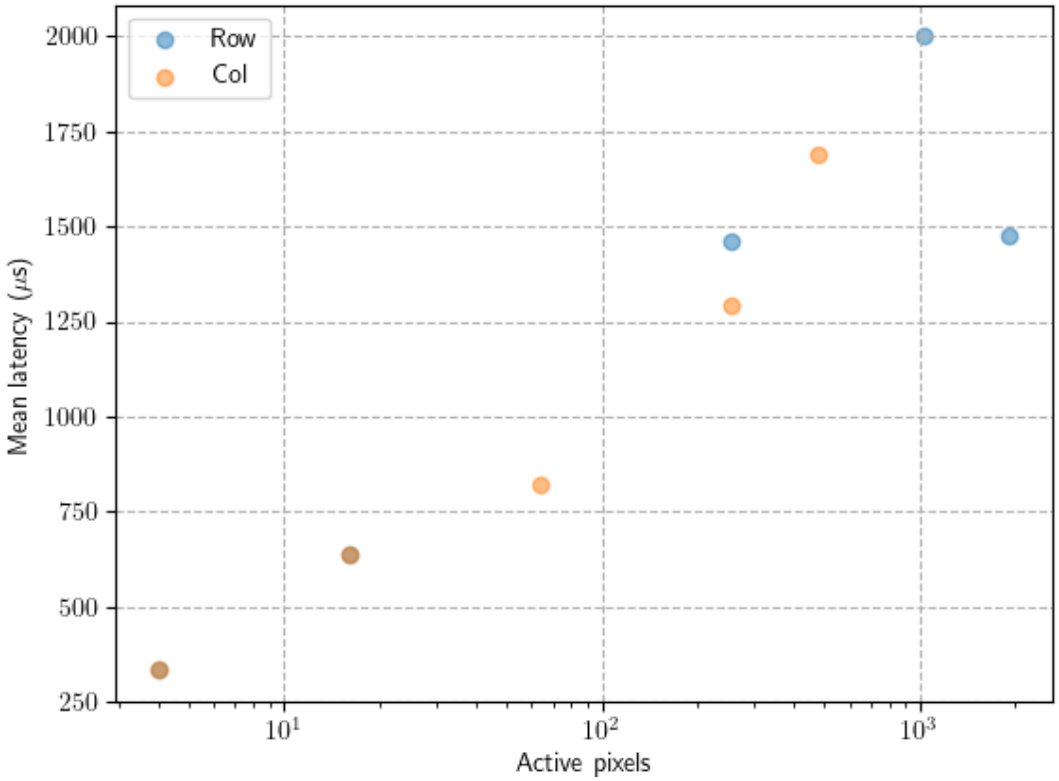

Fig. 14 Mean latency of analyzed samples with respect to total number of active pixels. A subset of trials where only columns and rows were varied are overlaid. For column analysis, each sampled trial had one row enabled. For row analysis, each sampled trial had four columns enabled.

Dependence on the ROI dimensionality was examined further in Fig. 15. Each curve corresponds to trials whose total number of active pixels are equivalent, but the number of rows vary. Trial points whose ROI comprise more rows contain less columns and vice versa. Latency increases as the ROI size exceeds 64 pixels and encompasses more rows than columns. These results reflect the nature of the EBS' AER interface: under medium-to-heavy load the AER latency is bounded by burst cycle time, which is faster than arbitration cycle time.

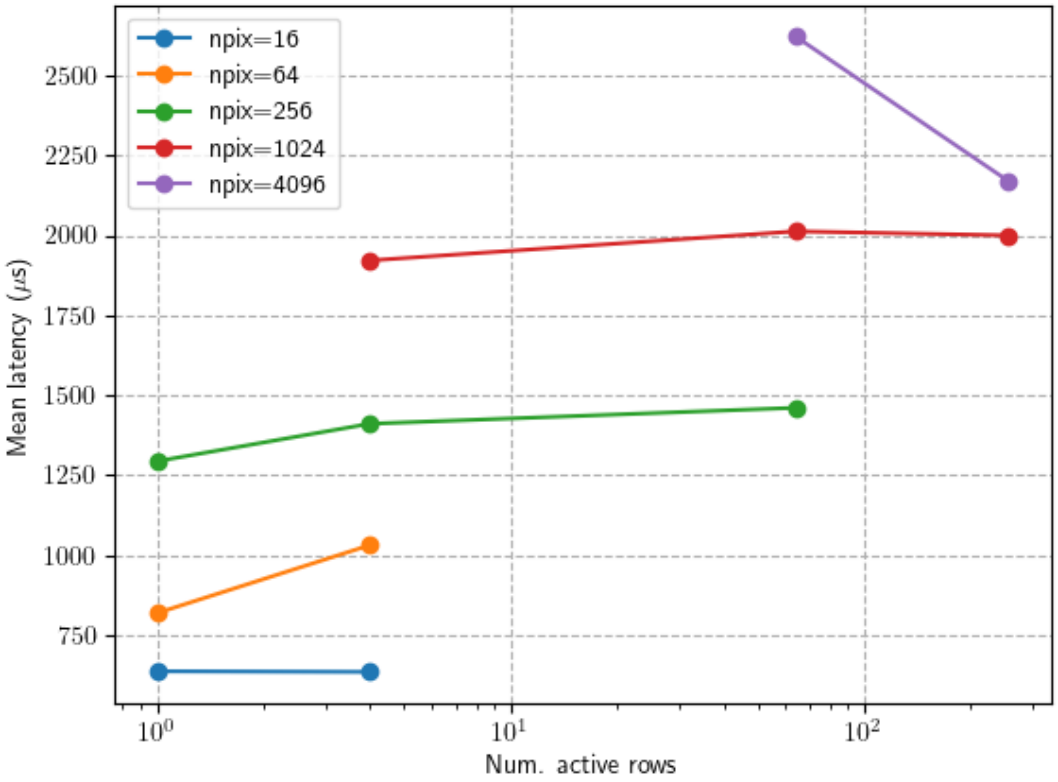

Fig. 15 Average latency plotted with respect to the number of active rows. Each color denotes trials with a common total amount of pixels enabled.

Figure 16 summarizes the analysis regarding EBC temporal latency dependence on ROI location. The top plot overlays trials that have the same number of active pixels but correspond to origin points on the left or right edge of the chip. There does not seem to be a noted effect on latency from varying the ROI horizontally. However, the data suggests that pixels at the bottom edge of the array respond slower than those near the top. The cause of this effect can be attributed to the mechanism behind the AER's handshake with the row and column. It has been reported in burst-mode AER interfaces that row handshaking is bounded by ability to reset the state of the request upon acknowledgment. This circuitry typically propagates a signal laterally across the array and makes the response of the handshake dependent on the array resolution rather than the origin of the request.[10] Conversely, the column handshake requires communication between the row interface and the column circuitry at the top of the array. This physical distance on the chip varies based on the location of the selected row; therefore, latency could be expected to vary similarly (see the bottom plot of Fig. 14).

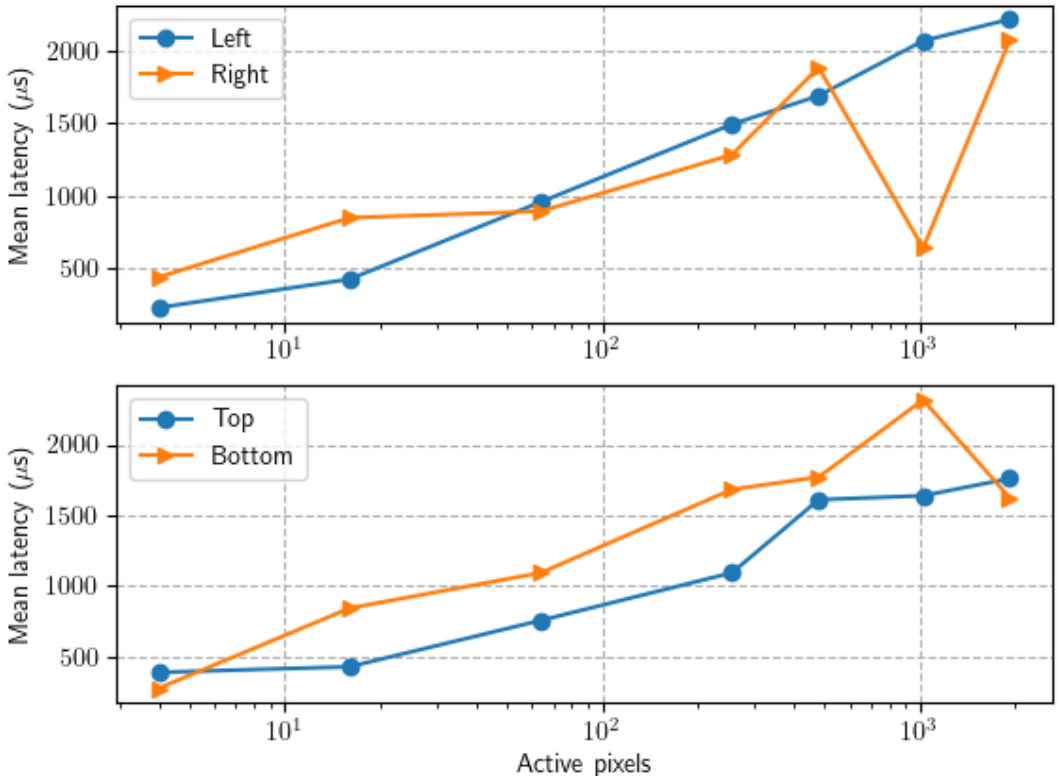

Fig. 16 Average mean latency plotted with respect to total enabled pixels and approximate trial ROI origin. Column-based and row-based shift of the ROI are depicted in the top and bottom plots, respectively.

5. Conclusion and Future Work

This contribution presented the theoretical underpinnings and experimental analysis of the temporal response for event-based cameras. Due to the relative recent inception of the field, there is not nearly as much understanding of the bounds of the technology as compared to their frame-based counterparts. This study sought to demystify the operating principles and mechanism behind the speed of event-based sensors by methodically characterizing the metric across a variety of device and experimental parameters.

Through random and ordered selection of singular pixels, a strong dependence of latency and intra-pixel jitter on analog bias configurations was discovered. Pixel latency has a direct dependence on irradiance, bandwidth configuration, and threshold sensitivity. Under the same background illumination, there is a 400% increase in latency when adjusting the pixel from its fastest to slowest configuration. The results suggest that configuring the pixel to possess a wider temporal bandwidth is the largest factor in faster response. The *wide* bias parameterization utilized a nominal sensitivity setting but was adjusted to have this

larger bandwidth that could sense faster frequencies. Thus, faster pixel response requires the ability to configure the device to exhibit such behavior. Even lower latency can be achieved by coupling this configuration with higher sensitivity (max), but the increased temporal frequency sensitivity plays the largest role. Moreover, latency is not uniform across successive samples of the same pixel or across the array. Pixel response time can vary up to 5–10 μs locally (intra-pixel jitter) and 50–150 μs globally. As event-based sensors are incorporated into real-time applications, these fundamental limits in response uniformity need to be characterized as they place a bound on closed-loop system performance.

The single-pixel tests also yielded unique insight into the variation of pixel performance across multiple cameras. As seen from analyzing data using the ordered sampling scheme, a sensitivity of the latency to specific locations existed. However, there was not sufficient evidence to suggest that average pixel response across the entire array varies considerably between the two cameras. Further experiments need to be realized to confirm this preliminary result, but they suggest that overall performance should not degrade when utilizing multiple cameras in a system, but specific regions in each sensor may see some variance. On the contrary, there was seen to be little dependence on pixel location within the same EBS. Comparing the results across ordered and random sampling schemes showed that the overall distribution did not vary—suggesting no relationship between single-pixel latency and location within the EBS.

This relationship between location and spatial activity was explored further by exploiting the EBC's ability to modify the ROI. Experimental data confirmed the theoretical operation of the AER interfaces as heavy array activity largely increased the mean latency. Worst-case waiting time under light-to-medium interface loads greatly exceeded the reported speed metrics afforded to EBS. When activating only 16 pixels, response time for certain pixels in the ROI could exceed 2 ms—much slower than the microseconds latency touted for the technology. Adding to the complexity of this claim, a dependence was shown between ROI dimension and location. An ROI with more rows (or near the bottom of the array) yielded higher response times under medium-to-heavy load. These insights into the speed dependence on spatial activity, coupled with the bias dependence shown in the single-pixel tests, are intended to inform users of the technology and empower effective utilization in future challenging endeavors.

Further study is certainly needed to cement the relationships presented and explore those that were inconclusive. These include stimulating the EBC with a wider range of background illuminations, testing additional EBC using identical sensors, and characterizing EBC from other vendors.[11] Additional work will also be devoted to studying system latency. This entails using well-characterized transient stimuli to

activate the EBC, which will then propagate data to a processing device and produce a response.[12] Such an experiment would not only capture sensor latency, but also the time needed to packetize data through the camera-level embedded devices, serial interfaces, and vendor software. Characterizing this behavior across a variety of operating conditions and current technologies will yield critical system engineering information for potential users and provide a set of realistic metrics to guide the development of future EBC technology.

6. References

1. Lichtsteiner P, Posch C, Delbruck T. A 128 × 128 120 dB 15μs latency asynchronous temporal contrast vision sensor. IEEE Journal of Solid-State Circuits. 2008;43(2):566–576.

2. Brandli C, Berner R, Yang M, Liu SC, Delbruck T. A 240 × 180 130 db 3 μs latency global shutter spatiotemporal vision sensor. IEEE Journal of Solid-State Circuits. 2014;49(10):2333–2341.

3. Posch C, Matolin D, Wohlgenannt R. A QVGA 143 dB dynamic range frame-free PWM image sensor with lossless pixel-level video compression and time-domain CDS. IEEE Journal of Solid-State Circuits. 2010;46(1):259–275.

4. Sivilotti MA. Wiring considerations in analog VLSI systems, with application to field-programmable networks. California Institute of Technology. 1991.

5. Boahen KA. A burst-mode word-serial address-event link-I: transmitter design. IEEE Transactions on Circuits and Systems I: Regular Papers. 2004;51(7):1269–1280.

6. Boahen KA. A burst-mode word-serial address-event link-III: analysis and test results. IEEE Transactions on Circuits and Systems I: Departmental Papers. 2004;51(7):1292–1300. doi: 10.1109/TCSI.2004.830701.

7. Son B, Suh Y, Kim S, Jung H, Kim JS, Shin C, Ryu H. A 640 × 480 dynamic vision sensor with a 9μm pixel and 300Meps address-event representation. Proceedings of 2017 IEEE International Solid-State Circuits Conference (ISSCC); 2017 Feb. IEEE; c2017. p. 66–67.

8. Finateu T, Niwa A, Matolin D, Tsuchimoto K, Mascheroni A, Reynaud E, Posch C. A 1280 × 720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86 μm pixels, 1.066 GEPS readout, programmable event-rate controller and compressive data-formatting pipeline. Proceedings of 2020 IEEE International Solid-State Circuits Conference-(ISSCC); 2020 Feb. IEEE; c2020. p. 112–114.

9. SilkyEvCam (VGA). CenturyArks Co, Ltd; n.d. [accessed 15 Aug 2023]. https://centuryarks.com/en/silkyevcam-vga/.

10. Lin J, Boahen K. A delay-insensitive address-event link. Proceedings of 2009 15th IEEE Symposium on Asynchronous Circuits and Systems; 2009 May. IEEE; c2009. p. 55–62.

11. DVXplorer.pdf. inivation.com; n.d. [accessed Aug 10]. https://inivation.com/wp-content/uploads/2023/03/DVXplorer.pdf.

12. Sengupta J, Linne B, Bard A. Characterization of the high-speed detection limits of a neuromorphic vision sensor. DEVCOM Army Research Laboratory (US); 2020 Nov. Report No.: ARL-TR-9117.

Appendix. Bias Values

Table A-1 outlines the bias values used to parameterize the SilkyEvCam through the Metavision SDK. Biases "bias-fo" and "bias-hpf" are used to widen the detectable temporal frequencies for the wide configuration and correlate B_{BW} in Fig. 2. Biases "bias-diff-on" and "biases-diff-off" were modified to lower and raise the thresholds and affect the pixel temporal contrast sensitivity. These correlate to B_{on} and B_{off}, respectively, in Fig. 2. More information about the biases in the context of the Metavision SDK can be found elsewhere.[1]

Table A-1 Bias values used for each configuration in the latency analysis: green and burnt orange cells reflect positive and negative deviations from the nominal (nom) bias setting. All values are unitless and do not reflect an absolute current or voltage level. Conversion is performed in software before interfacing with the EBS.

Bias name	Wide	Max	Nom	Min
Bias-diff	299	299	299	299
Bias-diff-on	200	228	200	150
Bias-diff-off	397	370	397	445
Bias-fo	1399	1477	1477	1477
Bias-hpf	1525	1448	1448	1448
Bias-pr	1250	1250	1250	1250
Bias-refr	1350	1500	1500	1500

[1] Biases — Metavision SDK Docs 4.2.1 documentation. Prophesee; n.d. [accessed 2023 Sep 15]. https://docs.prophesee.ai/stable/hw/manuals/biases.html.

List of Symbols, Abbreviations, and Acronyms

AC	alternating current
AER	address event representation
API	application programmer interface
BG	background
CD	change detection
COTS	commercial off-the-shelf
CSV	comma-separated value
DC	direct current
DRAM	dynamic-random-access memory
EBC	event-based camera
EBS	event-based sensor
EO-IR	electro-optical infrared
FIFO	first-in/first-out
FG	foreground
FPGA	field-programmable gate array
fps	frames per second
LED	light-emitting diode
ITR	integrate-then-read
IWR	integrate-while-read
LoRA	long range
MIPI	Mobile Industry Processor Interface
ND	neutral density
PC	personal computer
PLD	programmable logic device
ROI	region of interest
SA	situational awareness

SNR	signal-to-noise ratio
SUT	sensor under test
TIA	transimpedance amplifier
USB	universal serial bus

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DEVCOM ARL FCDD RLB CI TECH LIB
1 (PDF)	DEVCOM ARL FCDD RLA TD J SENGUPTA