Year: 2023

Shining light on the DVS pixel: A tutorial and discussion about biasing and optimization

Graca, Rui ; McReynolds, Brian ; Delbruck, Tobi

Abstract: The operation of the Dynamic Vision Sensor (DVS) event camera is controlled by the user through adjusting different bias parameters. These biases affect the response of the camera by controlling - among other parameters - the bandwidth, sensitivity, and maximum firing rate of the pixels. Besides determining the response of the camera to input signals, biases significantly impact its noise performance. Bias optimization is a multivariate process depending on the task and the scene, to which the user's knowledge about pixel design and non-idealities can be of great importance.In this paper, we go step-by-step along the signal pathway of the DVS pixel, shining light on its low-level operation and non-idealities, comparing pixel level measurements with array level measurements, and discussing how biasing and illumination affect the pixel's behavior. With the results and discussion presented, we aim to help DVS users achieve more hardware-aware camera utilization and modelling.

Shining light on the DVS pixel:
A tutorial and discussion about biasing and optimization

Rui Graça, Brian McReynolds, Tobi Delbruck

Inst. of Neuroinformatics UZH/ETH Zurich,Switzerland

Corresponding author: rpgraca@ini.ethz.ch

Abstract

*The operation of the Dynamic Vision Sensor (**DVS**) event camera is controlled by the user through adjusting different bias parameters. These biases affect the response of the camera by controlling - among other parameters - the bandwidth, sensitivity, and maximum firing rate of the pixels. Besides determining the response of the camera to input signals, biases significantly impact its noise performance. Bias optimization is a multivariate process depending on the task and the scene, to which the user's knowledge about pixel design and non-idealities can be of great importance.*

In this paper, we go step-by-step along the signal pathway of the DVS pixel, shining light on its low-level operation and non-idealities, comparing pixel level measurements with array level measurements, and discussing how biasing and illumination affect the pixel's behavior. With the results and discussion presented, we aim to help DVS users achieve more hardware-aware camera utilization and modelling.

1. Introduction

Over the past decade, neuromorphic or event-based DVSs [1]–[4] gained significant attention as a disruptive sensing technology, demonstrating key performance advantages over conventional frame-based sensors, including wide intra-scene dynamic range, low latency and power consumption, and a data-sparse output capturing only the dynamic information in a visual scene [5]. These benefits come at the cost of increased circuit complexity compared to frame-based sensors, which is both a blessing and a curse. On one hand, DVSs are highly adaptable, offering many degrees of freedom in performance through tunable biases; however, optimizing biases for a particular application is extremely difficult, even for experts in the technology. The goal if this paper is to present the challenging topic of bias optimization in an easy to understand way and provide specific guidelines for setting event camera biases based on task requirements and scene limitations.

General biasing guidance is discussed in [1], [6], [7], but descriptions are mostly qualitative and do not capture the trade-offs encountered when adjusting biases. Key performance metrics for DVS are measured and reported in [8]–[10] for varied illumination levels; however, full characterization across varied bias settings is not reported.

In [11], [12], we show how DVS noise performance and bandwidth depend on illumination and photoreceptor bias, and [13] presents a discussion regarding the effect of different biases on the output noise rate. In the same work, an algorithm based on feedback control that dynamically tunes the bias settings with the goal of keeping the output event rate within a programmable target window is presented. A discussion about how biasing affects the sharpness of the output of a event camera is presented in [14]. In [15], we propose a DVS characterization method to infer Temporal Contrast (**TC**) event threshold, pixel bandwidth, and refractory period, and show and discuss how the inferred parameters depend on the biases.

A good understanding of the pixel operation and the effect of the biases is also important for the development of realistic DVS models and simulators [16]–[18]. Accurate models are important to predict the response of a DVS to an arbitrary scene, which can both help users optimize their set-up and generate of simulated event datasets from pre-existing frame-based datasets. In particular, *v2e* [17] explains and models DVS motion blur and noise.

Most of the knowledge about DVS operation existing in the literature, including that regarding the effect of each bias, is based on the empirical observation of the DVS output, often supported by a simplified model of the pixel operation. Therefore, many second-order effects and non-idealities are often overlooked. In this paper, we go a step further, and present and discuss how visual input signal is converted and encoded into a stream of ON and OFF events at the DVS output, and how the different biases can be optimized according to the scene and the application requirements. Combined with looking at the relation between input and output, we present and discuss what is going on inside

the DVS pixel, supported by measurements of a test-pixel isolated from a DAVIS346 array.

During extensive measurements on internal pixel signals, we realized certain phenomena are far from intuitive and many aspects have generally been overlooked or over-simplified, especially regarding sensitivity and bandwidth. This oversimplification often leads to incorrect assumptions about how to bias or to model a DVS. In this work, we present a thorough description of DVS operation, including non-ideal second order effects, and how biasing affects operation. We present this description using terms that are simple enough for DVS users with minimal circuit expertise yet still accurately describe the pixel's operation.

1.1. Contributions and Outline

Our goal in this paper is to compile the most comprehensive description of the DVS pixel available by explaining precisely how a signal produces events, and how pixel behavior is influenced by the multivariate combination of user-defined biases and physical non-idealities along the way. To accomplish this, Sec. 2 reviews the the pixel response in general terms by presenting pixel-level node voltage measurements that illustrate a time varying signal propagating through each stage of a DVS pixel. Secs. 3 and 4 describe the pixel stages and corresponding biases in detail and report key performance metrics and noise rates across a range of bias values. We include never-before-published node voltage measurements from a DAVIS346 pixel and array measurements to graphically communicate low-level circuit behaviors in an easy-to-digest format. In Sec. 5, we aggregate this information to present a first-of-its-kind bias optimization spreadsheet tool which recommends specific bias adjustments or "tweaks" across a comprehensive range of task requirements and scene limitations. This deep dive into pixel performance leads to novel interpretations of threshold levels drifting in time as a function of brightness, and that every event encodes information about signal, noise, and junction leakage, which have important implications for improving future vision algorithms.

2. Basic DVS Operation

The basic operation of the DVS pixel has been extensively covered in [1], [5], [13], [17]. Fig. 1A shows how each pixel converts an illumination signal changing in time to an output stream of events representing changes in relative light intensity. Fig. 1B shows a simplified pixel model. Input light converted to a voltage by a logarithmic photoreceptor, and changes in this voltage are encoded in events. Events can be of *ON* polarity, encoding an increase in light intensity, or of *OFF* polarity, encoding a decrease in light intensity. An ON event is associated with ON threshold ($\Theta_{\mathbf{ON}}$) – each ON event encodes that the input light intensity changed by a factor of $\exp(\Theta_{\mathbf{ON}})$ since the last event

A: Dynamic Vision Sensor (DVS) principle of detecting brightness changes
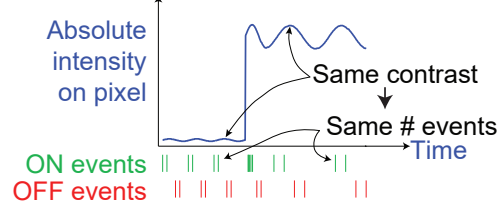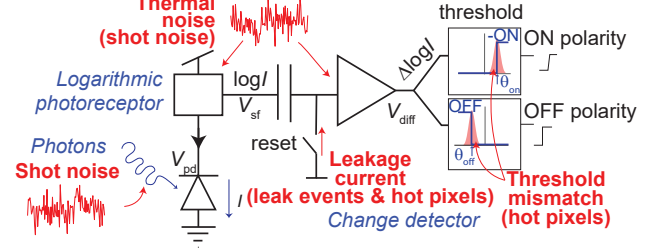

B: DVS pixel and noise sources

Figure 1. **A:** idealized output from Dynamic Vision Sensor (**DVS**) pixel. **B:** Simplified DVS pixel circuit with main noise sources (adapted from [19].)

(of either polarity). Similarly, each OFF event, associated with OFF threshold ($\Theta_{\mathbf{OFF}}$), encodes a change by a factor of $\exp(-\Theta_{\mathrm{OFF}})$ since the last event. Noise and pixel-to-pixel mismatch affect the circuit through the sources illustrated.

Fig. 2 shows the detailed DVS pixel circuit. The photoreceptor, **A**, consists of a photodiode **PD** that transduces light into an electric current I_{p}, and a feedback loop that ensures that the output voltage signal V_{pr} is proportional to the logarithm of I_{p}. Stage **B** cleans the signal by removing fast changes in V_{pr}, which is good when the signal of interest moves slower than the noise components. The maximum frequency of the signal allowed through this stage is called the pixel "bandwidth," and depends on two biases: I_{pr} and I_{sf}, as well as the input illuminance. One important remark here is that the pixel bandwidth is a totally different concept from the camera's readout bandwidth, i.e., the total rate of events we can read out from the camera, which is limited by the readout circuitry and not by the pixel.

Stage **C** amplifies voltage changes on node V_{sf} by the ratio of the capacitors ($\frac{-C_1}{C_2}$), which is around 20 for the DAVIS346. The resulting amplified signal V_{diff} is then compared to an ON threshold proportional to $\log(I_{\mathrm{on}}/I_{\mathrm{d}})$ and to an OFF threshold proportional to $\log(I_{\mathrm{d}}/I_{\mathrm{off}})$ [20]. When either of the thresholds is met, either an ON or an OFF event is generated and transmitted to the outside of the array. At the same time, a reset signal is generated in **E**, that sets V_{diff} to a reset state value independent of the input and controlled uniquely by I_{d}. V_{diff} is held at the reset state for a time period know as "refractory period". The duration of the refractory period is inversely proportional to its control bias

Table 1. Bias names used in this paper (Fig. 2) and their equivalent name for DAVIS346 and other iniVation cameras

Short name	DAVIS346	Description
I_{pr}	PrBp	Photoreceptor bias
I_{sf}	PrSFBp	Source Follower (Buffer) bias
I_d	DiffBn	Change Amplifier bias
I_{on}	OnBn	ON Threshold bias
I_{off}	OffBn	OFF Threshold bias
I_{refr}	RefrBp	Refractory bias

I_{refr}, and during the refractory period the input signal is ignored by the pixel. This way, the pixel operation switches between a tracking mode, in which V_{diff} amplifies the input signal, and a reset mode after each event, where V_{diff} is reset to a known state and ignores input changes.

The equivalence between the bias names used in this paper (and generally adopted in the literature), their descriptive name, and the name used in iniVation documentation and software is presented in Table 1.

To illustrate pixel response at each stage, Fig. 3 presents direct measurements of node voltages V_{pr} (**A**) and V_{diff} (**B**), and a computer-reconstructed V_{sf} signal (**C**). The recording captures 1.25 periods of a $5\,\text{Hz}$ sinusoidal input light signal with a contrast of 0.62 log-e units (meaning that the ratio between maximum and the minimum brightness of the sine wave is $e^{0.62}$, or 1.86). This signal results in one OFF event (at t ≈ 0.12s) and two ON events (t ≈ 0.17s ≈ 0.22s). The measurements clearly demonstrate how noisy V_{pr} is in practice, and how effectively bandwidth control with I_{sf} removes high frequency noise – V_{sf} is considerably cleaner, indicating a much more favorable Signal-to-Noise Ratio (**SNR**). From V_{sf} to V_{diff}, the signal is inverted and amplified, resulting in a much larger peak-to-peak signal amplitude. Additionally, we can clearly see the effects of the reset logic after each threshold crossing, as V_{diff} returns to a fixed level immediately after each event. On this timescale (and bias configuration), the refractory period is too small to see.

3. The Logarithmic Photoreceptor

3.1. Phototransduction

In the journey from incident photons to DVSs events, the first and crucial step is phototransduction. Ideally, every photon the lens captures is focused onto the photodiode and transfers its energy to break free an electron. The photoelectric effect dictates one photon generates a single "photoelectron," provided it has enough energy. Silicon is well suited to detect visible light, as only ≈ 1.1 eV of energy is needed to generate a photoelectron, and visible photons have 1.8 (red) - 3.1 (violet) eV. These photoelectrons move freely when a voltage is applied, comprising "photocurrent".

In reality, several non-ideal behaviors influence pho-

tocurrent. Each photon has a finite probability of interacting with an atom in the silicon, and some photoelectrons will recombine with an atom when they encounter a vacancy. These factors decrease quantum efficiency – essentially, the fraction of photons that become usable current. Additionally, even when there is no light on the pixel, the photodiode still conducts a small amount of *dark current*, and photocurrent smaller than this value is indiscernible.

Even when illumination is constant, photons and electrons do not arrive at a fixed rate. Instead, there is some statistical variability to the number of photons arriving on the pixel during a given time window, as well as in the number of electrons arriving at each node in the circuit. This fluctuation results in "shot noise", which is the dominant noise source in DVS operation [11], [12] and the reason why the measured photoreceptor output in Fig. 3**A** looks far noisier than the dashed orange signal obtained by averaging the photoreceptor output over many cycles.

3.2. Feedback logarithmic photoreceptor

One of the key features of the DVS is the logarithmic photoreceptor (**A** in Fig. 2) [22]. The logarithmic relationship between input light and output voltage enables two of the main advantages of the DVS: first, it compresses the representation of the output signal, allowing a much higher dynamic range than conventional frame based cameras (which encode light intensity linearly); and second, it makes small changes around a background value proportional to the background value. This means that it directly encodes reflectance — to first-order, the same object observed in different light settings will always result in the same output signal (Fig. 1**A**). Fig. 4 shows how the photoreceptor voltages V_{pr} and V_{pd} change with input illuminance. V_{pd} is nearly constant with illuminance at a value determined by I_{pr}, while V_{pr} varies with the logarithm of input illuminance with a gain (slope) mostly independent of I_{pr}. The measurements also clearly show the lower limit of dynamic range, set by the dark current.

3.3. Second-order effects of illumination

Even though to a first order approximation the photoreceptor output is independent of the absolute illumination level, this is far from true for practical applications. Absolute illumination strongly affects noise and bandwidth, as discussed in [11]. The bandwidth of the photoreceptor is generally directly proportional to light intensity (unless I_{pr} is set very low), meaning that in the dark, the ability to detect fast changes in the scene is limited and the latency from signal change to pixel response is longer. Also the Background Activity (**BA**) profile (i.e. the rate of "noise" events not encoding signal change) strongly depends on light intensity, as shown in Fig. 5. The figure shows the BA rate vs. illuminance for both the array and an isolated test pixel op-
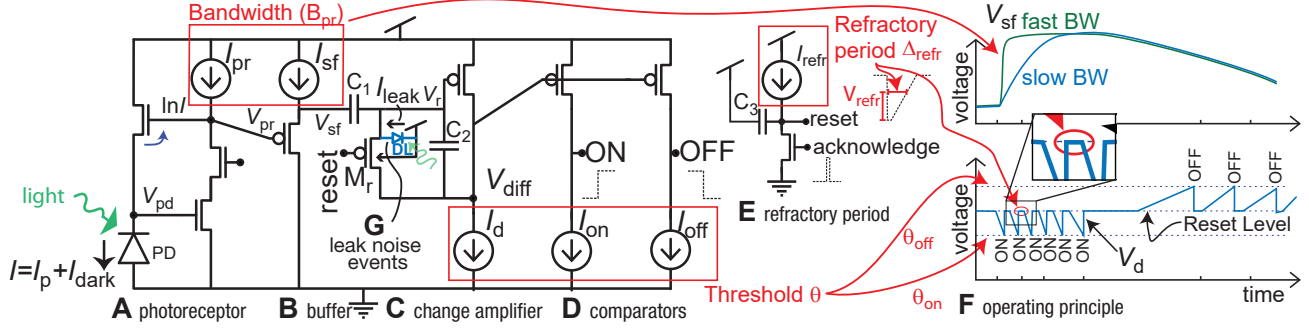
Figure 2. Typical DVS pixel circuit [21]. The active logarithmic photoreceptor (**A**) is buffered by a source-follower (**B**), which drives a cap-feedback change amplifier (**C**), which is reset on each event by a low-going *reset* pulse. A finite refractory period (**E**) holds the change amplifier in reset for the refractory period Δ_{refr}. Comparators (**D**) detect ON and OFF events as seen in **F**. Periodic leak events result from junction and parasitic photocurrent I_{leak} in diode DL (**G**).
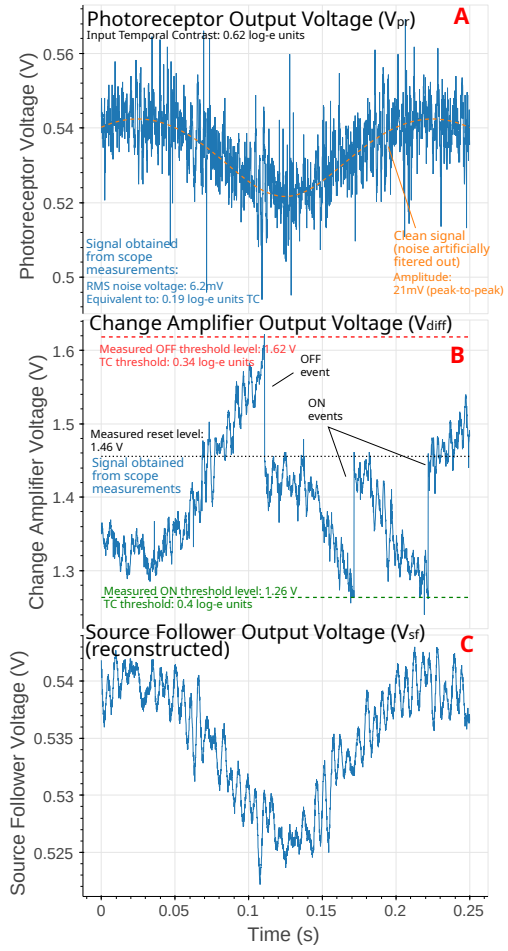

Figure 3. Test pixel recordings from internal nodes of the DVS pixel in Fig. 2. **A** and **B** show the measured voltage at V_{pr} and V_{diff} for a 5 Hz input sine wave with a TC contrast of 0.62. **C** shows a computer reconstruction of V_{sf} based in **A** and **B**, obtained by low-pass filtering **A** at a cutoff frequency of 200 Hz. I_{pr}, is biased strongly (large current), so the bandwidth is large and thus V_{pr} is noisy. The fast frequency components are then filtered out by I_{sf}
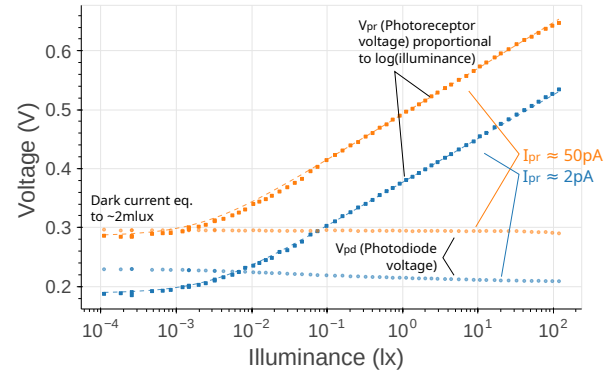

Figure 4. Photoreceptor voltage (V_{pr}) and Photodiode voltage (V_{pd}) versus illuminance for two different I_{pr} settings. The dashed lines show a fit of a model of the photoreceptor including dark current: $V_{\text{pr}} = \frac{U_T}{\kappa} log(E_{v,\text{p}} + E_{v,\text{dark}}) + V_{\text{p0}}$, where $E_{v,\text{p}}$ is the input illuminance, $E_{v,\text{dark}}$ is the illuminance equivalent to the dark current observed (fitted to approx. 2 mlx), U_T is the thermal voltage ($\approx$25 mV at room temperature), κ is a parameter of the transistor M_{fb} and is fitted to approx. 0.75. V_{p0} is a parameter determining the voltage level at 1 lx, and it varies with the logarithm of I_{pr}.

erating in the same conditions. As explained in [11], [20], the BA consists of a high rate of random shot noise events of both polarities in the dark, exclusively ON events with rate proportional to light intensity in bright scenes (described in Sec. 4), and a dip in BA between these two regions.

3.4. The Source-Follower Buffer

The Source-Follower (**SF**) Buffer (**B** in Fig. 2) serves two main purposes: it decouples the sensitive photoreceptor stage (**A**) from the spiking stages of the pixel (**C**, **D**, and **E**), and it can be used to limit the bandwidth, filtering out noise components faster than the signal of interest. This effect is visible in Fig. 3, where V_{sf} (**C**) and V_{diff} (**B**) exhibit significantly less noise (and of lower frequency components) than
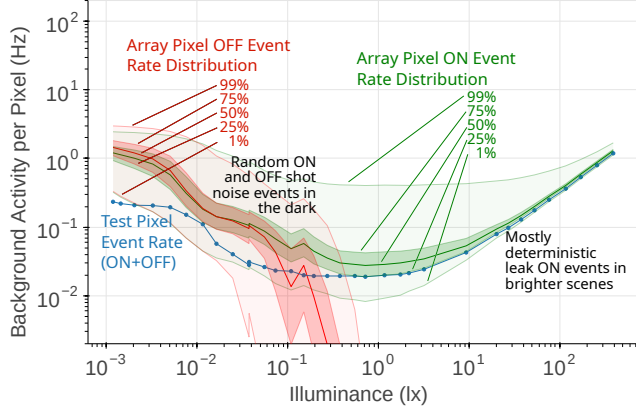
Figure 5. Distribution of the BA event rate per pixel in a DAVIS346 biased with I_{pr} of $30\,\text{pA}$ and I_{sf} of $15\,\text{pA}$ and nominal threshold and refractory period settings (tweaks of 0 in Figs. 7 and 9) for varying on-chip illuminance. The green and red lines show the quantiles of the distribution of ON and OFF event rates in pixels of the array indicated in the figure. The blue line shows the event rate for a test pixel in the same camera, but isolated from the array.

V_{pr} (**A**). The bandwidth limit imposed by the SF stage is proportional to its bias current I_{sf}.

3.5. Photoreceptor and Buffer biasing

A discussion about optimal biasing of the photoreceptor and SF buffer is presented in [12]. Both I_{pr} and I_{sf} affect bandwidth and noise, but in different ways. The main practical difference is that I_{pr} potentially introduces significantly more noise than I_{sf}, but because it is downstream in the circuitry, I_{sf} can be used to remove excess noise introduced by I_{pr}. The higher I_{pr}, the faster the noise components it introduces, and the easier we can filter them out by adjusting I_{sf}. As a result, contrary to what was proposed in [13], the DVS achieves significantly better noise performance when I_{pr} is high and I_{sf} is low enough to filter out noise introduced by I_{pr}. I_{sf} can be used to limit signal bandwidth, but also to only remove high frequency noise components introduced by I_{pr} without affecting signal bandwidth.

Fig. 6 shows the distribution of noise event rate per pixel in a DVS camera in dark conditions (around $40\,\text{mlx}$) for varying I_{pr}. We see that for very low I_{pr}, noise increases with I_{pr} (because bandwidth increases), but for higher I_{pr} noise starts to decrease because the faster (higher frequency) noise components introduced by I_{pr} are filtered out by the SF. The results also show that from a certain point, increasing I_{pr} has little effect on the output noise. This, combined with the fact that increasing I_{pr} leads to higher power consumption, suggests that there is an optimal sweet spot for I_{pr} (dependent on the selected I_{sf}) at the lowest I_{pr} where all of its noise components are filtered out.
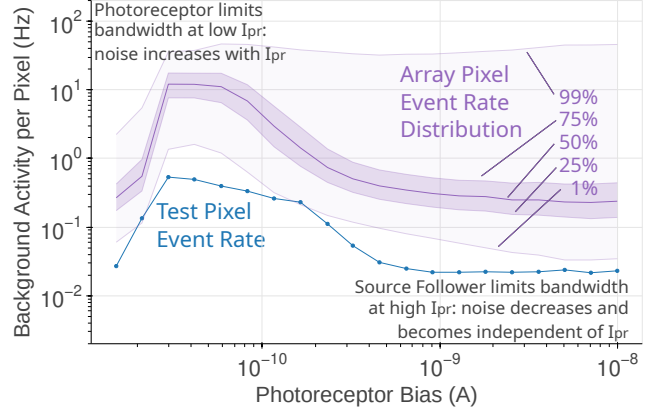

Figure 6. Distribution of the BA per pixel in a DAVIS346 array (in purple) for varying I_{pr} and for an isolated test pixel (in blue) at an on-chip illuminance of $40\,\text{mlx}$ and I_{sf} of $30\,\text{pA}$ and nominal threshold and refractory period settings

This point can be found empirically by finding the minimum I_{pr} setting beyond which any increase will result in negligible noise rate reduction. Under these circumstances, bandwidth is limited by (and proportional to) input light intensity up to an illumination level dependent on I_{sf}, and then limited by I_{sf} at higher illumination levels. In [12], we can see how bandwidth increases with I_{sf}, allowing higher frequency noise components to pass through – leading to higher noise rates and pixel bandwidth.

4. Event Generation

The last step of the DVS pixel is the conversion of a voltage signal to a stream of events, as described in Sec. 2. Understanding the process of event generation is fundamental to understanding how DVS encodes data. One key aspect of event generation is that the state of the pixel (i.e. current illumination level) is stored (and reset after each event) in V_{diff} as an analog voltage. Unlike digital memory, analog memory destructively decays in time. In the case of the DVS, this happens because of current that leaks through the reset switch (**G** in Fig. 2). The effect of leakage is that the voltage V_{diff} linearly decays in time at a rate that increases with absolute input illuminance. One way to look at this, is that V_{diff} is continually drifting toward an ON event, even in the absence of changes to the input signal. In this case, V_{diff} hits the ON threshold level at a fixed rate, generating so called "leak events," which dominate the BA in brighter scenes, as depicted in Fig. 5 [20]. An important note is that leakage is not noise – its behavior is deterministic, and signal information is not destroyed but rather encoded in the timing between events.

Even though events are often discriminated into either "signal", "noise", or "leak" for simplification [11], [17], [19], it is important to note that an event never consists
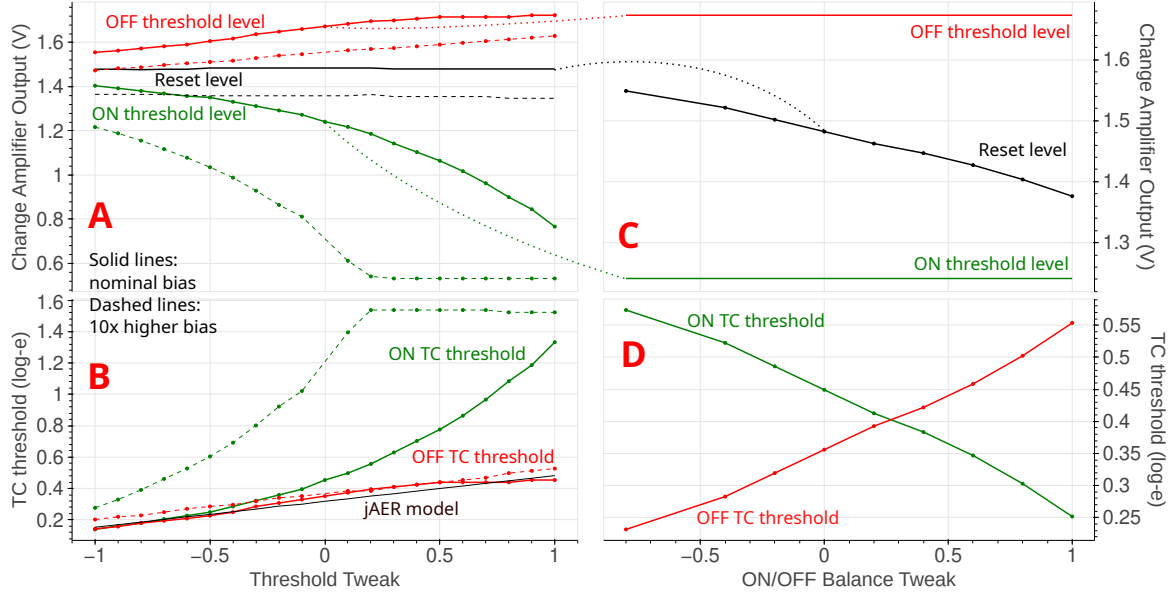
Figure 7. Threshold variation with jAER tweaks for threshold and ON/OFF balance. **A** shows the voltages measured at V_{diff} at which ON and OFF events are triggered (green and red lines) and to which the pixel is reset (in black). The solid lines were measured for nominal I_{on}, I_{off} and I_{d}, and the dashed lines were measured for 10x higher I_{on}, I_{off} and I_{d}. **B** shows the TC threshold equivalent to the same measurements. **C** shows the same measured voltages as **A**, but for varying ON/OFF balance tweak, and **D** the equivalent TC thresholds.

purely of one of these phenomena. Rather, every event is a combination of all three. Noise causes not only the generation of events in the absence of signal changes, but also uncertainty (jitter) in the timing of an event in response to a signal change. Observing Fig. 3, we can predict that events will most likely occur when noise adds constructively to the signal in the direction of the nearest threshold. In the same way, a small underthreshold signal can generate an event if noise adds constructively, and an over-threshold signal may not generate an event if noise adds destructively. Additionally, leakage makes it progressively harder to generate an OFF event (and easier to generate an ON event) as time increases after each reset. This can be interpreted as the OFF threshold increasing and the ON threshold decreasing linearly in time. These aspects of event generation influence how a DVS should be biased, and should also be incorporated to design more robust event-processing algorithms.

4.1. Event Threshold

The TC event threshold quantifies how much the input signal needs to change since the last event to trigger the next event. An ON event is generated when V_{diff} drops low enough to toggle the output of the ON comparator (**D** in Fig. 2). This level is controlled by I_{on}. Likewise, I_{off} controls the level to which V_{diff} has to increase to generate an OFF event. This results in TC thresholds, Θ_{ON} and Θ_{OFF}, proportional to $\log(I_{\text{on}}/I_{\text{d}})$ and $\log(I_{\text{d}}/I_{\text{off}})$ respectively.

To study the effect of the TC threshold, we adjusted the

Threshold Tweak implemented in jAER described in [13]. The Threshold Tweak is a user defined value between -1 and 1 that controls the TC threshold by adjusting I_{on} and I_{off}. The solid green and red lines in Fig. 7**A** show the measured V_{diff} values at which ON and an OFF events occur, and the solid black line shows the voltage to which V_{diff} is reset after each event. As expected, increasing the tweak increases the distance between the reset level and the threshold levels, but with two caveats: first, the ON threshold deviates from the desired linear variation with the tweak - this results from inaccurate modeling in jAER which can be corrected, and second, the OFF threshold saturates for higher tweaks - this happens because the value for I_{off} reaches the minimum allowed for proper circuit operation. The second condition can be alleviated by increasing all I_{d}, I_{on}, and I_{off} by the same factor (at the expense of higher power consumption). The dashed lines show the corresponding levels when all currents are increased by a factor of 10. Fig. 7**B** shows TC thresholds equivalent to the same measurements, as well as the expected TC level modeled by jAER. Considering the interpretation of leakage as a gradual decrease of Θ_{ON} and increase of Θ_{OFF} following each reset, these TC thresholds represent measured values immediately after reset.

Fig. 8 shows how BA varies with the Threshold Tweak. As predicted by [11], noise rates depend exponentially on the threshold. This dependence is given by the tail of the Gaussian function, since noise is approximately normally distributed. For higher thresholds, shot noise becomes neg-
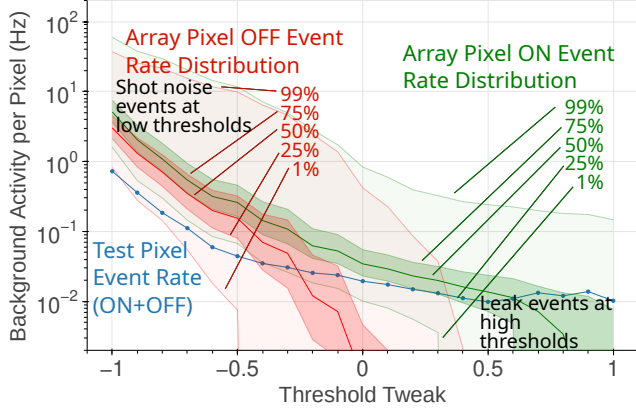
Figure 8. Distribution of ON (in green) and OFF (in red) BA per pixel in the array and BA rate in an isolated test pixel (in blue) for varying jAER Threshold Tweak at an on-chip illuminance of $40\,\mathrm{mlx}$, I_{pr} of $3\,\mathrm{nA}$, I_{sf} of $15\,\mathrm{pA}$, and nominal threshold and refractory period settings

ligible and leak events dominate the output.

4.2. Balance between ON and OFF thresholds

In ideal DVS operation, for most applications, it would be desirable to have the same TC threshold for ON and OFF events. However, as seen before, leakage introduces a fundamental imbalance between ON an OFF events, with a preference towards ON events. As discussed in [23], deliberately applying imbalanced thresholds towards a more sensitive Θ_{ON} can lead to better noise performance.

The balance between Θ_{ON} and Θ_{OFF} while keeping the same $\Theta_{\mathrm{ON}} + \Theta_{\mathrm{OFF}}$ can be adjusted by altering the bias I_{d}, or equivalently, the ON/OFF balance tweak implemented in jAER, which adjusts I_{d} around a nominal value. The effect of the tweak is shown in **C** and **D** in Fig. 7. In **C**, we can see that the voltage level at the output of the Change Amplifier that results in an ON or OFF event (green and red lines) is independent of the tweak, but the level to which V_{diff} resets after an event (black line) changes linearly with the tweak. In **D** we see how this translates to TC threshold.

4.3. Refractory period

The duration of the refractory period is inversely proportional to its control bias I_{refr}. Fig. 9**A** shows the measured length of the refractory period for varying values of the Max Firing Rate Tweak in jAER. A linear change in the tweak corresponds to an exponential change in I_{refr}. The refractory period was measured by directly observing the reset signal of a test pixel after an event and measuring the time the pixel is held in reset mode. Samples of pulses of the reset signal for different tweaks are shown in Fig. 9**B**. The pixel is held in reset while the reset signal is low (below $\approx 1.2\,\mathrm{V}$). As shown, the length of the refractory period is
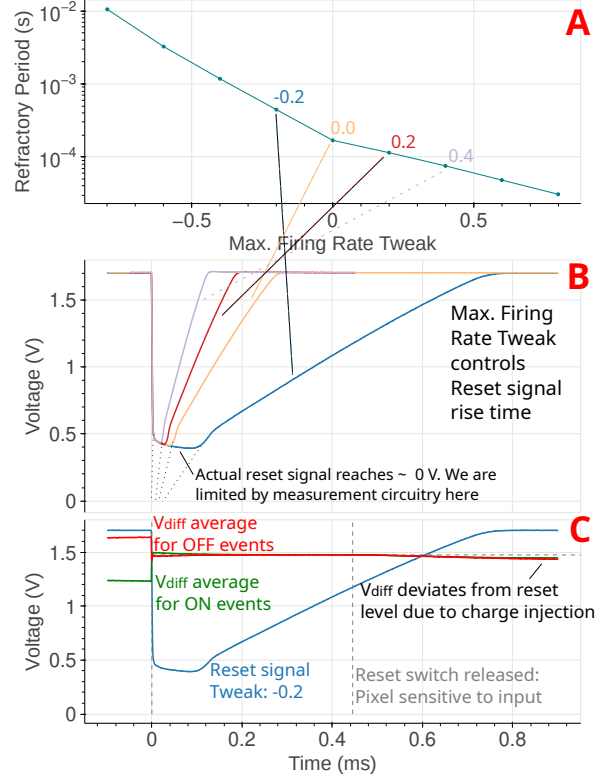

Figure 9. In-pixel measurements on the refractory period. **A** shows the length of the refractory period vs. the jAER Max Firing Rate Tweak. The length was obtained by directly measured the time during which the reset signal is active. **B** show average samples of measured pulses of the reset signal for different values of the tweak, with the correspondence with the points in **A** marked. **C** shows the same averaged reset pulse (for a tweak of -0.2) and the average of V_{diff} for ON (in green) and OFF (in red) events.

controlled by changing the speed of the rising edge of the reset pulse. We observed that below a tweak of -0.8, I_{refr} is too weak to drive the pixel out of reset, and the pixel stops operating properly. Fig. 9**C** shows the measured voltage at V_{diff} following an ON (green) OFF (red) event, and the reset signal in blue. When the reset pulse drops, V_{diff} moves to the reset level ($\approx 1.5\,\mathrm{V}$). When the reset signal nears this level, the pixel starts to respond to the input signal again. This is not depicted here since the input signal is not changing. However, right after the reset signal reaches $1.2\,\mathrm{V}$, V_{diff} drops slightly from the reset level. This effect is usually not problematic (it is equivalent to having a slightly lower reset level), but for refractory periods faster than around $100\,\mathrm{\mu s}$, the amount of this deviation increases. For extremely short refractory periods, the operation of the pixel is severely affected - setting a practical upper limit to I_{refr}.

As discussed in [13], [23], I_{refr} sets a trade-off between data fidelity and noise rate. A higher refractory period reduces noise, but of discards some signal information (e.g.
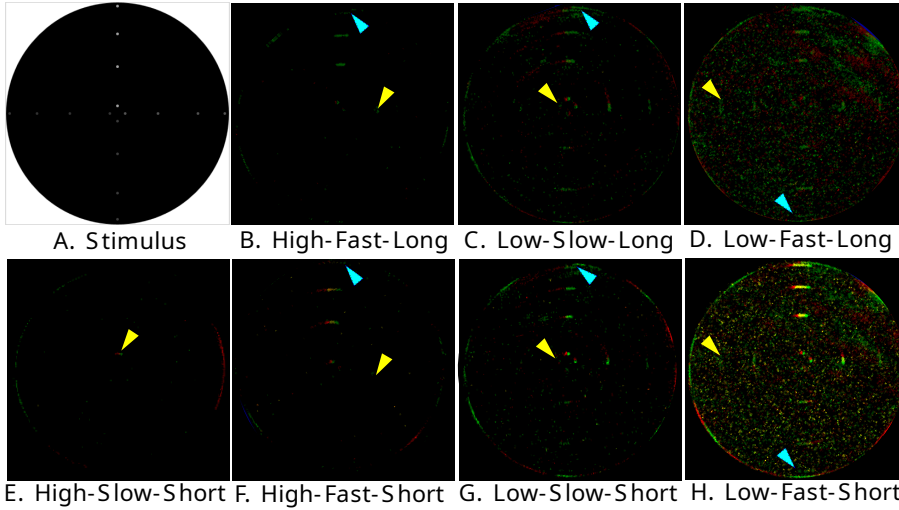
A. Stimulus   B. High-Fast-Long   C. Low-Slow-Long   D. Low-Fast-Long

E. High-Slow-Short   F. High-Fast-Short   G. Low-Slow-Short   H. Low-Fast-Short

Figure 10. Multivariate biasing effects. $10\,\mathrm{ms}$ accumulations of DVS events are shown while viewing a $125\,\mathrm{RPM}$ motorized disk (**A**) with gray dots of varied contrast and speed (radial distance) under $15\,\mathrm{lx}$ illumination with a f/1.8 lens. Events are green (ON) and red (OFF), and full scale rendering is 3 events. Each image (**B-H**) represents a combination of threshold (High/Low), bandwidth (Fast/Slow), and refractory period (Long/Short) biases. The leading edge of the dimmest (yellow) and fastest (blue) detectable dot is identified by an arrow. If multiple dots of the same speed are detected, only the dimmest is indicated – if multiple of the same brightness, only the fastest.

consecutive events at an edge indicating a brighter object).

5. Multivariate Bias Dependency

As DVS are applied to increasingly diverse applications, it is important to consider both the low-level behaviors controlled by sensor biases and the multivariate effects of their combination. Fig. 10 provides a visualization of how varied bias combinations can have a drastic influence on sensor response to the same dynamic scene. The scene consists of a motorized disk with small white dots of different contrast and radial distance, rotating at $125\,\mathrm{RPM}$ under dim indoor lighting conditions, using a scene illuminance of $\approx 15\,\mathrm{lx}$.

From this small sample set, its clear that varied bias combinations result in drastically different performance. From left to right, the sensor becomes gradually more sensitive to faster and dimmer spots, but the noise also increases. In this scenario, the combination of bandwidth and threshold determines the fastest and dimmest detectable objects in a scene. The refractory period limits how much information is collected from each dot. For example, in Fig. 10**D**, the long refractory period makes it more difficult to differentiate between objects of varied brightness compared to **H**, where the brightest dots clearly produce more events. All three parameters influence noise, and nuanced interactions can have a significant influence on overall noise rates, as demonstrated in [23].

Although these visual examples are instructive, in a real application, there are more than just two sensing criteria to consider. To facilitate selecting an optimal bias configuration for widely varied scenarios, we created an online spreadsheet[1] that recommends a combination of event camera biases based on six criteria that describe task requirements and scene limitations. The selected criteria are *data priority* (i.e., whether higher event counts that include high-fidelity brightness information, or a data sparse representation of the visual field with lower event counts is preferred), *sensor motion*, *background illumination*, *object size*, *object contrast*, and *object speed*. Each criteria is split into two categories, resulting in 64 categorical combinations, and for each we propose a set of bias "tweaks" to use as a starting point for optimizing event camera performance. Because different event camera models and software interfaces have slightly different bias formats, we describe bias tweaks in relative terms (i.e. fast-mid-slow bandwidth, low-mid-high sensitivity, long-mid-short refractory period).

6. Conclusion

The results and discussion presented in this paper bring light to the internal operation of the DVS pixel, by showing how low-level effects such as noise and leakage affect signal, and by exploring how biasing can be optimized when these effects are considered. Moreover, we propose a novel way to interpret events as a combination of signal, noise, and leakage, and an interpenetration of leakage as a variation of TC threshold in time. These approaches can lead to more optimal modeling and utilization of the DVS. The proposed Bias Optimization Tool summarizes information described in the paper, and provides the DVS user with an easy way to optimize biasing based on the task and scene, considering interaction of the main variables playing a role in the process.

---

[1]CVPR 2023 - DVS Bias Optimization Tool: https://docs.google.com/spreadsheets/d/1XaS3hkcj1bSG5gaMnlAy89rbsomILDgu/edit#gid=1310047800

References

[1] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128×128 120 dB 15 μs latency asynchronous temporal contrast vision sensor," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 2, pp. 566–576, 2008.

[2] T. Finateu, A. Niwa, D. Matolin, K. Tsuchimoto, A. Mascheroni, E. Reynaud, P. Mostafalu, F. Brady, L. Chotard, F. LeGoff, H. Takahashi, H. Wakabayashi, Y. Oike, and C. Posch, "5.10 a 1280×720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86μm pixels, 1.066GEPS readout, programmable event-rate controller and compressive data-formatting pipeline," in *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2020, pp. 112–114. DOI: 10 . 1109 / ISSCC19947.2020.9063149.

[3] C. Brandli, R. Berner, M. Yang, S. Liu, and T. Delbruck, "A 240 × 180 130 dB 3 μs latency global shutter spatiotemporal vision sensor," *IEEE Journal of Solid-State Circuits*, vol. 49, no. 10, pp. 2333–2341, 2014. DOI: 10.1109/JSSC.2014.2342715.

[4] Y. Suh, S. Choi, M. Ito, J. Kim, Y. Lee, J. Seo, H. Jung, D. H. Yeo, S. Namgung, J. Bong, S. Yoo, S. H. Shin, D. Kwon, P. Kang, S. Kim, H. Na, K. Hwang, C. Shin, J. S. Kim, P. K. J. Park, J. Kim, H. Ryu, and Y. Park, "A 1280×960 dynamic vision sensor with a 4.95-μm pixel pitch and motion artifact minimization," in *2020 IEEE International Symposium on Circuits and Systems (IS-CAS)*, 2020, pp. 1–5. DOI: 10.1109/ISCAS45731.2020. 9180436.

[5] G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza, "Event-based vision: A survey," en, *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 1, pp. 154–180, Jan. 2022.

[6] *User guide - biasing dynamic sensors*, en, https://gitlab. com / invation / invation - docs / blob / master / Advanced % 20configurations / User _ guide_ - _Biasing.md, Accessed: 2022-2-9.

[7] *Biases — metavision intelligence docs 3.1.2 documentation*, en, https : / / docs . prophesee . ai / stable / hw / manuals/biases.html, Accessed: 2023-3-17.

[8] P. Lichtsteiner, "An AER temporal contrast vision sensor," Ph.D. dissertation, 2006.

[9] D. P. Moeys, "Analog and digital implementations of retinal processing for robot navigation systems," Ph.D. dissertation, 2016.

[10] C. Li, "Two-stream vision sensors," Ph.D. dissertation, 2017.

[11] R. Graca and T. Delbruck, "Unraveling the paradox of intensity-dependent DVS pixel noise," in *2021 International Image Sensor Workshop (IISW)*, Sep. 2021. DOI: 10.48550/ARXIV.2109. 08640.

[12] R. Graca, B. McReynolds, and T. Delbruck, "Optimal biasing and physical limits of DVS event noise," in *2023 International Image Sensor Workshop (IISW)*, May 2023. DOI: 10.48550/arXiv. 2304.04019.

[13] T. Delbruck, R. Graca, and M. Paluch, "Feedback control of event cameras," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2021, pp. 1324–1332. DOI: 10.1109/CVPRW53098.2021.00146.

[14] M. S. Dilmaghani, W. Shariff, C. Ryan, J. Lemley, and P. Corcoran, "Control and evaluation of event cameras output sharpness via bias," Oct. 2022. arXiv: 2210.13929 [cs.CV].

[15] B. J. McReynolds, R. P. Graca, and T. Delbruck, "Experimental methods to predict dynamic vision sensor event camera performance," *Optical Engineering*, vol. 61, no. 7, p. 074103, 2022. DOI: 10.1117/1.OE.61.7.074103.

[16] H. Rebecq, D. Gehrig, and D. Scaramuzza, "ESIM: An open event camera simulator," *Conf. on Robotics Learning (CoRL)*, Oct. 2018.

[17] Y. Hu, S.-C. Liu, and T. Delbruck, "V2e: From video frames to realistic DVS events," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, Jun. 2021, pp. 1312–1321.

[18] D. Joubert, A. Marcireau, N. Ralph, A. Jolley, A. van Schaik, and G. Cohen, "Event camera simulator improvements via characterized parameters," *Frontiers in Neuroscience*, vol. 15, 2021, ISSN: 1662-453X. DOI: 10.3389/fnins.2021.702765.

[19] S. Guo and T. Delbruck, "Low cost and latency event camera background activity denoising," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 785–795, 2023. DOI: 10.1109/TPAMI.2022.3152999.

[20] Y. Nozaki and T. Delbruck, "Temperature and parasitic photocurrent effects in dynamic vision sensors," *IEEE Transactions on Electron Devices*, vol. 64, no. 8, pp. 3239–3245, Aug. 2017, ISSN: 1557-9646. DOI: 10.1109/TED.2017.2717848.

[21] G. Taverni, D. Paul Moeys, C. Li, C. Cavaco, V. Motsnyi, D. San Segundo Bello, and T. Delbruck, "Front and back illuminated dynamic and active pixel vision sensors comparison," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 65, no. 5, pp. 677–681, 2018.

[22] T Delbrück and C. Mead, "Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits," 1995.

[23] B. McReynolds, R. Graca, and T. Delbruck, "Exploiting alternating DVS shot noise event pair statistics to reduce background activity," in *2023 International Image Sensor Workshop (IISW)*, May 2023. DOI: 10.48550/arXiv.2304.03494.