

Event-Based Noise Filtration with Point-of-Interest Detection and Tracking for Space Situational Awareness

Nikolaus Salvatore, Alan D. George
NSF SHREC Center
ECE Dept., University of Pittsburgh
4420 Bayard St., Suite 560,
Pittsburgh, PA 15213
{nds40, alan.george}@pitt.edu

ABSTRACT

This paper explores an asynchronous noise-suppression technique to be used in conjunction with asynchronous Gaussian blob tracking on dynamic vision sensor (DVS) data. The high dynamic range and fast acquisition time of DVS recordings enables the imaging of high-velocity targets despite ordinarily problematic lighting conditions. The technique presented here relies on treating each pixel of the sensor as a spiking cell that keeps track of its own activity over time, which in turn can be filtered out of the resulting sensor event stream by user-configurable threshold values. In addition, asynchronous blob tracking is supplemented with double-exponential smoothing prediction and Bezier curve-fitting in order to smooth tracker movement and interpolate target trajectory respectively. This overall scheme is intended to achieve asynchronous point-source tracking using a DVS for space applications. In the space environment, radiation effects are expected to introduce transient, and possibly persistent, noise into the asynchronous event-stream of the DVS. Given the large distances between objects in space, targets of interest may be no larger than a single pixel and can therefore appear similar to noise-induced events. In this paper, the asynchronous approach is experimentally compared to a more traditional approach applied to reconstructed frame data for both performance and accuracy metrics. The results of this research show that the asynchronous approach can produce comparable or even better tracking accuracy while also drastically reducing the execution time of the process by seven times on average.

Keywords

Dynamic vision sensor; event-based image processing; space situational awareness; biomimetic vision

1. INTRODUCTION

Recent years have seen the development of a new class of imaging sensors capable of replicating basic properties of biological vision, namely its focus on detecting changes within scenes. The first of these new neuromorphic vision systems was the DVS proposed in [1], which details the sensor's architecture and relation to biological analogs. The DVS functions by detecting logarithmic intensity of luminance changes through a series of photoreceptors and integrating and comparative circuits associated with each individual pixel. Luminance changes are detected via conventional

photodetectors such as those often found in active-pixel sensors. Once the luminance change at a given pixel induces a voltage beyond a certain predefined threshold, the cell will generate an event that encodes the (x,y) coordinates, polarity, and timestamp. Each of the pixel cells of the sensor monitors both positive and negative changes in luminance intensity, which is in turn reported in the positive or negative value of the polarity associated with each pixel event. It is also important to note that the events generated are not synchronized with the internal clock, creating the need for the timestamp recorded with each event. These events are then streamed to the onboard processor through a multiplexing technique referred to as address-event representation (AER) that serves to maintain the absolute order in which events occurred. This system and addressing scheme have been shown to register events on the microsecond scale, making the DVS ideal for applications requiring extremely fast response times. The overall benefit of this asynchronous, change-based architecture is superior power efficiency, temporal resolution, and dynamic range as well as drastically reduced data rate as compared to conventional cameras [1,2]. After several iterations on similar biomimetic vision systems, the asynchronous time-based image sensor (ATIS) has emerged as one of the more well-developed variations on the base DVS design. While it retains the strengths of the DVS architecture, the ATIS also encodes the relative intensity of luminance in the timing of events, enabling it to reconstruct full gray-scale images in addition to the binary events generated by the DVS. The ability to reconstruct variable-intensity images allows for more traditional image-processing and computer-vision techniques, while still leveraging the high-speed data acquisition of the asynchronous sensor [3].

The primary focus of this research is to exploit the capabilities of DVSs for object tracking within the context of space-based observation, particularly in a low-Earth orbit (LEO) environment. The space environment imposes a unique set of challenges on computer vision that impact both the software and underlying hardware involved. First, and perhaps most fundamental, of these challenges is the size, weight, power, and cost (SWaP-C) constraints placed on hardware. DVSs excel in this regard as they are both lightweight and power-efficient, while also requiring fewer computing resources to perform event-based image processing. These qualities make them ideal for deployment on space platforms that make use of embedded architectures with strict power constraints. Another challenge faced by object tracking in space is the relative visibility of certain objects of interest. A significant percentage of space debris is made up of objects mere centimeters in size, and which have exceptionally low reflectance while travelling at extremely high speeds. Despite their small size, the high speed of debris poses a serious danger to space platforms,

which stand to benefit greatly from autonomous methods of avoidance [4,5]. Although some large-scale, ground-based optical solutions have been able to track exceptionally small space debris, the high dynamic range (HDR) and temporal resolution of DVSs may prove useful for autonomous collision avoidance onboard space platforms. Furthermore, the change-based detection properties of DVSs, coupled with their HDR and relatively low data rate, could be ideal for detecting aerial objects on Earth. Recent work using DVSs in conjunction with ground-based telescopes has demonstrated the potential of event-based approaches to tracking celestial objects. This research compared several bio-inspired vision sensors, including a DVS and ATIS, for the purposes of tracking objects in both low-Earth orbit and geosynchronous orbit (GEO) during daytime lighting conditions, displaying the efficacy of the HDR of bio-inspired sensors [6].

Another challenge of particular interest in this research is the effect of radiation on spaceborne hardware. Radiation effects are typically classified as either transient or cumulative, from which correct operation may or may not be recoverable. Both types of radiation effects can pose significant issues for space missions, though transient single-event effects (SEEs) can often be managed by simple power-cycling of hardware or through standard hardware and software dependability techniques [7]. For imaging equipment, SEEs can manifest as salt-and-pepper noise that can be mitigated via common noise-filtering techniques when processing captured images. However, not all SEEs are transient, and some can cause permanent damage to sensing equipment that results in affected pixels being latched in an excited or unexcited state. Cumulative total ionizing dose (TID) effects are of relatively greater concern due to the long-term deleterious effects on hardware [8]. TID studies with CMOS cameras have shown that proton and heavy-ion radiation common in space can significantly impact photodiode responsivity and contribute to long-term degradation [9]. Other studies conducted on bipolar transistors, such as those that compose operational amplifiers, have shown significant impacts on voltage response with increasing TID [10]. Given the reliance of DVS upon CMOS components, they will likely experience many of the same types of degradation as other types of conventional CMOS cameras [11].

2. RELATED WORK

The asynchronous nature of DVS data introduces new avenues for image processing given that the pixel data does not exist in conventional image frames. It is possible to use conventional image-processing and computer-vision techniques by integrating the events over a predefined period of time and then reconstructing conventional frames from these events. Recent work has also shown that a hybrid approach can be taken where conventional feature detection algorithms, such as the Harris corner detector, are used to locate features which are then tracked in the asynchronous event stream [12]. However, processing the events entirely in an asynchronous manner can significantly reduce the execution time of various operations as well as decrease the overall computational load. To date, several common image processing algorithms have been modified to work solely within an asynchronous context. Event-based cluster trackers have proven to be quite effective at tracking large objects such as vehicles, which have clear positive and negative polarity edges, at extremely high framerate and with low computational resource usage [13]. Algorithms for asynchronous optical-flow calculation have exhibited superior speed and processing usage, while also retaining comparable performance to frame-based approaches [14]. An event-based Hough circle transform demonstrated the ability to perform high-

speed, multi-object tracking within the context of microparticle tracking [15]. Additionally, spiking neural networks (SNNs) have been shown to naturally complement the spiking nature of DVS pixels. In recent work, SNNs have been used to asynchronously detect and track lines via clusters of spiking neurons representing line parameters in the Hough space. However, this approach scales poorly with increasing DVS pixel array size due to the large increase in the number of neurons required [16].

One technique of particular relevance used in this research is the event-based Gaussian blob tracker introduced in [17]. This approach allows for both object detection and tracking by instantiating a series of trackers whose position and shape in the visual field are defined by the parameters of their bivariate Gaussian distribution. Every new event generated by the sensor is evaluated for its probability of belonging to each of the existing trackers, spawning a new tracker if no tracker has a score beyond a predefined threshold. An exponentially decaying activity score is also associated with each tracker to differentiate between active and inactive trackers.

Several means of noise filtration have also been adapted for use with asynchronous event streams. One such method exploits the temporal aspect of the event stream by requiring each incoming event to be supported by neighboring events within a certain predefined time threshold. A two-dimensional array of the camera's visual field records the timestamp of the most recently generated event within a window. When a new event is generated, the event is only passed on if the timestamp at the corresponding pixel coordinate has a timestamp more recent than the chosen support time [13]. An issue with using this method for space applications is that objects of low reflectance may not always induce enough luminance change to generate an event. As a result, objects appearing as a single pixel may not necessarily excite every pixel in its path of motion and thus true events could be eliminated by the filter. Event-based optical flow has also been adapted for noise filtration by approximating the "lifetime" of events defined as the time required for adjacent pixels to be excited. Events with lifetimes close to zero are considered noise and omitted from the resulting event stream. While this method is effective, it requires events to be stored within a temporal window and can be computationally intensive with increasing window sizes [18].

Another technique relies on the supposition that actions and objects of interest in the foreground will generate relatively more events than the background. The visual field of the sensor is divided into an arbitrary number of cells that then maintain an exponentially decaying record of the activity occurring within them. Only cells with activity scores higher than the average activity of all cells will pass their corresponding events through the filter [19]. A modified form of this noise suppression is employed in this research. An alternative approach has been developed that generates a sparse representation within the cells using the K-SVD algorithm, which has been used to denoise both individual frames and videos composed of reconstructed frames. Although this approach thoroughly denoises tested frames, it is computationally time-consuming and has not yet been shown to operate in real-time [20][21].

Even with the inclusion of noise filtration, event-based blob tracking solutions rely on position updates that pull the tracker in opposing directions even when accurately tracking objects as a whole. This non-uniform movement occurs because the events associated with an object do not necessarily have timestamps temporally ordered in the direction of the object's physical motion. These constant changes of direction necessitate some form of

predictive filter in order to smoothly track object trajectories. Kalman and extended Kalman filters have been established as extremely effective for object tracking, but the overhead imposed can severely impact execution time [22]. Given the high-speed data acquisition of the DVS, event-based algorithms must be able to cope with potentially high-activity scenes. One alternative approach to Kalman filtering makes use of double exponential smoothing prediction (DESP), a common data-forecasting method, to track head and hand movements. The experimental results showed comparable accuracy to the Kalman filtering approach, but with 135 times faster performance [23]. Assuming reduced or eliminated noise events and the generally consistent motion expected in space, this approach is well-suited for smoothing event-based tracker trajectories.

In addition to smoothing tracker trajectories, the error of the tracker position is assessed by several means adapted from previous works. One metric used to predict tracking failure is the forward-backward error typically employed with median flow tracking. This technique involves comparing the future and past trajectories of a given tracker and assigning an error value based on a chosen measure of distance. In an event-based context, this error can be used to suppress trackers activated by areas of constantly changing motion, which most likely do not correspond to objects of interest [24]. Lastly, a common metric for measuring tracker precision is the intersection over union (IoU) calculated between the given tracker and a ground-truth representation of the object of interest. This metric has been used extensively to compare the performance of different tracking algorithms and is used in this work to compare asynchronous and frame-based approaches with noise suppression [25].

3. APPROACH

The following section details the algorithm used for object tracking and noise suppression intended for space-based applications. The blob-tracking aspect of this research relies on the event-based Gaussian tracker technique established in [17] with specific emphasis on tracking single-point sources. This emphasis is necessitated by the large distances at which space objects are to be tracked in relation to the comparatively low resolution of the DVS (640×480). The event-based Gaussian trackers allow for both efficient object detection and tracking but would be susceptible to noise and attraction to non-target objects in space environments. The asynchronous approach attempts to mitigate this behavior by introducing tracker suppression in addition to noise suppression. Both asynchronous and conventional frame-based approaches are presented and then compared across a variety of metrics. The frame-based approach differs in that it accumulates events over a predefined interval of time and then constructs an image frame from the resulting array. The reconstruction of frames allows more conventional image-processing techniques to be applied, but it also increases the execution time and computational resources required.

3.1 Tracking Algorithm

As stated previously, each event generated by the DVS is represented by a vector of the form $e_i = [x_i, y_i, t_i, p_i]$, where the x and y values correspond to the pixel location, t to the timestamp of the event and p to the polarity of the event. The polarity may only take on a value of $p \in \{1, -1\}$, indicating an increase or decrease in luminance respectively. The motion of objects can then be visualized as a point-cloud of events modeled by a bivariate Gaussian distribution, $N(\mu, \Sigma)$. The parameters of the Gaussian distribution used to model the position and shape of a tracked object are defined as

$$\mathbf{u} = [x, y]^T \quad (1)$$

$$\Sigma = \begin{bmatrix} \sigma_x^2 & \sigma_{xy} \\ \sigma_{xy} & \sigma_y^2 \end{bmatrix}. \quad (2)$$

As each new event is generated, the probability of the event being associated with an active or inactive tracker is calculated as

$$p_i(\mathbf{u}) = \frac{1}{2\pi} |\Sigma_i|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{u}-\mathbf{u}_i)^T \Sigma_i^{-1}(\mathbf{u}-\mathbf{u}_i)} \quad (3)$$

where u_i denotes the corresponding locations of active and inactive trackers. The event is then associated with the tracker with the highest calculated p score and the parameters of the tracker are updated according to the weighted update calculation

$$\mathbf{u}_t = \alpha_1 \mathbf{u}_{t-1} + (1 - \alpha_1) \mathbf{u} \quad (4)$$

$$\Sigma_t = \alpha_2 \Sigma_{t-1} + (1 - \alpha_2) \Delta \Sigma \quad (5)$$

where

$$\Delta \Sigma = \begin{bmatrix} (x - u_{tx})^2 & (x - u_{tx})(y - u_{ty}) \\ (x - u_{tx})(y - u_{ty}) & (y - u_{ty})^2 \end{bmatrix}. \quad (6)$$

If no tracker has a calculated p score above a certain predefined threshold δp , a new tracker is instantiated with mean centered on the new event's location and covariance matrix initialized to predefined values. In this research, a δp of 0.001 was used while the starting values of the covariance matrix were varied experimentally. In addition to the parameters of the bivariate Gaussian, each of the trackers has an activity score that is updated according to the equation

$$A_i(t) = \begin{cases} A_i(t - \Delta t) e^{-\Delta t / \tau_1} + 1, & \text{if } p_i(\mathbf{u}) > \delta p \\ A_i(t - \Delta t) e^{-\Delta t / \tau_1}, & \text{otherwise} \end{cases} \quad (7)$$

where Δt is the time elapsed since the last tracker update and τ_1 is a constant chosen to tune the rate of tracker deactivation. If the activity A_i falls below a certain predefined threshold, the tracker will become inactive. In the asynchronous tracking approach, every event is processed as it is received. However, while the frame-based approach retains the timestamps associated with each event, it evaluates all events in order at discrete time steps after noise suppression has taken place. In [17], two activity thresholds were used to differentiate between deactivated and destroyed trackers. However, given that trackers in this research are intended for single-point sources that may have activities as low as one between updates, the activity threshold is chosen such that any activity is sufficient for the tracker to be considered active and no trackers are permanently destroyed. This can lead to large numbers of inactive trackers, but also ensures that small objects of interest will be tracked effectively.

In order to monitor the trajectory of the Gaussian blob tracker, its position is smoothed using DESP at regular intervals of time τ , which are chosen at runtime. With this method, the x and y positions of the blob tracker are treated as a time series of points modeled with a linear regression equation whose y-intercept and slope vary over time. At each multiple of time τ , two smoothing statistics are calculated as

$$\overrightarrow{\mathbf{S}\mathbf{u}}_\tau = \alpha_3 \overrightarrow{\mathbf{u}}_\tau + (1 - \alpha_3) \overrightarrow{\mathbf{S}\mathbf{u}}_{\tau-1} \quad (8)$$

$$\overrightarrow{\mathbf{S}\mathbf{u}}_\tau^{[2]} = \alpha_3 \overrightarrow{\mathbf{u}}_\tau + (1 - \alpha_3) \overrightarrow{\mathbf{S}\mathbf{u}}_{\tau-1}^{[2]} \quad (9)$$

where α_3 is an update parameter chosen to determine the degree of exponential decay and $\overrightarrow{\mathbf{u}}_\tau$ is the vector representing the blob

tracker's x and y location at time τ . The first smoothing statistic, \overline{Su}_τ , represents the smoothed average value of event positions associated with the tracker, while the second, $\overline{Su}_\tau^{[2]}$, captures the smoothed trend in event positions, i.e. the tracker's motion. With these smoothing statistics, the tracker position at time $\tau + 1$ is then forecasted according to the equation

$$\overline{u}_{\tau+1} = \overline{b}_0(\tau) + \overline{b}_1(\tau + 1) \quad (10)$$

where

$$\overline{b}_1(\tau) = \frac{\alpha}{(1-\alpha)} \left(\overline{Su}_\tau - \overline{Su}_\tau^{[2]} \right) \quad (11)$$

$$\overline{b}_0(\tau) = 2\overline{Su}_\tau - \overline{Su}_\tau^{[2]} - \tau \overline{b}_1(\tau) \quad (12)$$

The values calculated here, $\overline{b}_1(\tau)$ and $\overline{b}_0(\tau)$, represent the estimated slope and y-intercept respectively of the linear regression fitting the given tracker's position over time. Since these values vary over time, the α_3 parameter controls the extent to which new points affect the linear regression fit to the tracker's position. As a result, very small values on the order of 1×10^{-4} are chosen due to the large number of events involved in tracker updates. Finally, the smoothed trajectory of the tracked object is then interpolated between time steps with a cubic Bezier curve fit with the equation

$$\mathbf{B}(t) = (1-t)^3 \mathbf{U}_0 + 3(1-t)^2 t \mathbf{U}_1 + 3(1-t)t^2 \mathbf{U}_2 + t^3 \mathbf{U}_3 \quad (13)$$

where U_i indicates the $[x, y]^T$ position of points obtained from the smoothed tracker trajectory. In this equation, U_3 refers to the most recent tracker position in time obtained via smoothing, while the remaining points U_i refer to previous points in its trajectory.

Since new trackers are constantly spawned when an event has no tracker with p score above the given threshold, it is possible, and in practice quite likely, that multiple trackers will begin to overlap on the same object being tracked. To remedy this occurrence, the distances between and activities of active trackers are compared at each discrete timestep before curve-fitting takes place. Trackers that overlap with adjacent trackers of relatively higher activity are deactivated and omitted from curve-fitting.

3.2 Noise and Tracker Suppression

The noise-suppression techniques employed differ between the asynchronous and frame-based approaches. For the frame-based approach, a conventional image frame is built by tallying the presence of events at each pixel location regardless of polarity or frequency. Once an event is received with timestamp greater than or equal to a multiple of the integration time chosen, a simple summation kernel is convolved with the integrated image according to

$$g(x, y) = \sum_{u=-a}^a \sum_{v=-b}^b K(u, v) I(x - u, y - v) \quad (14)$$

where $I(x, y)$ is the integrated event frame and K is the kernel

$$K = \frac{1}{3} \begin{bmatrix} \mathbf{1} & \mathbf{1} & \mathbf{1} \\ \mathbf{1} & \mathbf{0} & \mathbf{1} \\ \mathbf{1} & \mathbf{1} & \mathbf{1} \end{bmatrix}. \quad (15)$$

Given that the outputs of this operation are fixed to integer values, $g(x, y)$ will only have values greater than zero when an event has at least three neighboring events within the integrated frame. This function therefore serves to mask any pixel location with fewer than three neighboring events occurring within the integration time τ .

The asynchronous approach uses a version of the event-based dynamic background suppression introduced in [19], except that every pixel is associated with its own cell rather than a large

number of pixels falling into the same cell. Reducing the number of events grouped for noise suppression is necessary for tracking targets that may be exceptionally small in size and induce small amounts of activity. In addition, cells are suppressed using a two-sided activity threshold rather than the average of all current cell activity. This change to the suppression technique serves to filter events occurring at pixels of both low and high activity. As an event is received, the corresponding cell is updated in a similar fashion as the activity associated with trackers as

$$A_u(t) = A_u(t - \Delta t) e^{-\Delta t/\tau_2} + 1 \quad (16)$$

where Δt is the time elapsed since the last pixel update, A_u is the activity associated with the pixel at $u = [x, y]^T$ and τ_2 is a constant chosen to tune pixel activity decay. Thresholds δ_1 and δ_2 are chosen such that only events occurring at pixels with activity $\delta_1 < A_u(t) < \delta_2$ will be processed by the Gaussian blob tracking algorithm. Negative polarity events are ignored in this filtering since motion of exceptionally small objects is expected to generate positive polarity events immediately followed by negative polarity events on their trailing edges. This programmatic filtering of noise events also serves to filter large regions of change in the sensor's view such as sections of the Earth that might lie in the sensor's view. Since the sensitivity of the hardware itself would need to be maximized in order to track distant, dim objects, portions of the Earth would generate a significant number of spatially and temporally close events. These events will be filtered by the asynchronous approach and therefore significantly reduce computational load and potential false positives in tracking.

Since the intent of this research is to track objects that may be as small as a single pixel in the DVS's view, parameters are chosen for the Gaussian blob tracker algorithm such that trackers will even be spawned for regions of low activity. Setting the activation threshold so low can in turn lead to some noise events behaving similarly to objects of interest and therefore escaping noise suppression. As a result, the trackers themselves are also filtered by two separate metrics. First, the Gaussian tracker algorithm is modified such that negative polarity events are accumulated separately from negative polarity events as

$$N_i(t) = \begin{cases} N_i(t - \Delta t) e^{-\Delta t/\tau_1} + 1, & \text{if } p_i(u) > \delta p \\ N_i(t - \Delta t) e^{-\Delta t/\tau_1}, & \text{otherwise} \end{cases} \quad (17)$$

where $N_i(t)$ is the negative polarity activity of tracker i at time t . A threshold is chosen at runtime such that trackers with negative polarity activity below this threshold will not be subject to DESP or trajectory curve fitting. These trackers are essentially omitted as not corresponding to an object of interest. This suppression serves to omit both trackers that may be associated with noise events as well as trackers that may be attracted by larger objects, such as sections of the Earth, which are not the intended target of this research. Second, a forward-backward error statistic is calculated at each multiple of the integration time τ when performing DESP. The Euclidean distance between the next predicted location $u_{\tau+1}$ of the smoothed trajectory and the last update location of the tracker is calculated and added to a running average of the tracker's forward-backward error as shown below.

$$E_t = E_{t-1} + \sqrt{(x_{\tau+1} - x_t)^2 + (y_{\tau+1} - y_t)^2} \quad (18)$$

$$\overline{E}_t = \frac{E_t}{n_{events}} \quad (19)$$

As with the negative polarity activity, an error threshold is chosen such that, if a tracker's forward-backward error exceeds the threshold, it will be deactivated and hidden from trajectory fitting.

It should be noted that the frame-based approach does not employ tracker suppression due to spatially solitary events always being filtered. As a result, the frame-based method’s ability to track point-source objects is limited, but the number of trackers instantiated is greatly reduced in comparison to the asynchronous method. The entirety of both algorithms is described with the pseudocode detailed in Algorithm 1 and 2.

Algorithm 1: Frame-based Tracking Algorithm

```

Initialize lastFrame and 2D image array;
for each incoming event,  $e(x,y,t,p)$  do
  Store  $e(x,y,t,p)$  at  $[x,y]$  position of 2D image array;
  if  $(t - \text{lastFrame}) \geq \text{integrationTime}$  then
    Perform noise suppression with Eq. 14 and 15 on image array;
    for each event in image array do
      Determine closest tracker according to Eq. 3;
      Update all tracker activities according to Eq. 7;
      Promote(or demote) trackers to active layer (to hidden layer);
      Update position/shape of tracker with highest  $p$  according to Eq. 4 and 5;
      Apply double-exponential smoothing to updated tracker position (Eq. 8-12);
    end
    Fit Bezier curve to smoothed active tracker positions (Eq. 13);
    Draw trackers, trajectories and events on output image frame;
  end
end

```

Algorithm 2: Asynchronous Tracking Algorithm

```

Initialize lastFrame and 2D array of pixel activity cells;
for each incoming event,  $e(x,y,t,p)$  do
  Update pixel cell at  $(x,y)$  according to Eq. 16;
  if  $\delta_1 \leq A_n(t) \leq \delta_2$  or event has negative polarity then
    Determine closest tracker according to Eq. 3;
    Update all tracker activities according to Eq. 7 and 17;
    Promote(or demote) trackers to active layer (to hidden layer);
    Update position/shape of tracker with highest  $p$  according to Eq. 4 and 5;
    Apply double-exponential smoothing to updated tracker position (Eq. 8-12);
    Calculate tracker forward-backward error and add to average according to Eq. 18 and 19;
  end
  if  $(t - \text{lastFrame}) \geq \text{integrationTime}$  then
    Draw events on output image frame;
    for each active tracker do
      if tracker  $N_i(t) \geq \text{neg. polarity threshold}$  and forward-backward error average  $\leq \text{error threshold}$  then
        Fit Bezier curve to smoothed active tracker positions (Eq. 13);
        Draw tracker and trajectory and events on output image frame;
      end
    end
  end
end
Flush image frame and set lastFrame to last event's  $t$ ;
end

```

3.3 Experimental Setup

In order to compare the asynchronous and frame-based approaches and assess their ability to track point source objects, experimental testing was conducted with varied velocity, luminance contrast, and distance to a point-source of interest. In this experiment, the point-source of interest consisted of a laser pointer projected onto a background illuminated by two direct current (DC), variable-intensity LED lamps. The lighting intensity of the lamps was varied at several discrete levels in order to assess tracking with different contrast levels. DC lighting was necessary, in addition to the room being darkened, due to the DVS’s tendency to detect flickering from alternating current (AC) lighting. The laser pointer was mounted onto a stepper motor controlled via microcontroller board in order to control the angular velocity of the target at discrete microstepping levels.

Tables 1 and 2 list the contrast ratios and speeds measured for each of the discrete levels used in testing. Tests were conducted using each combination of contrast ratio and speed mode for a total of 30 trials under base conditions. To study the effect of having a smaller target area, the distance from the camera to the target was then increased and additional tests were conducted with all contrast ratios and the Full, Quarter, and Slowest speed modes to cover the breath of speeds chosen. These 15 additional trials are denoted among the results with the corresponding contrast level and “Incr. Distance” in order to distinguish them. Tests were conducted at 2.5 meters and 5.0 meters to target respectively due to the constraints of the testing area. Since the targets in the 45 preceding trials follow a simple, horizontal and linear path, an additional 11 tests were conducted that included sudden and/or persistent changes in the direction of motion as well as multiple targets of interest. These additional tests are described in Table 3, several of which were repeated multiple times with different patterns of movement.

Table 1: Contrast Conditions

Contrast #	Contrast Ratio
1	∞
2	25.87
3	6.23
4	3.01
5	1.09

Table 2: Speed Conditions

Speed Mode	Angular Velocity (Deg/s)
“Full”	455.7
“Half”	232.26
“Quarter”	136.88
“Eighth”	63.16
“Sixteenth”	25.68
“Slowest”	7.4

Table 3: Additional Tests

Experiment (# of Trials)	Description
Multiple Targets (3)	Several targets with sudden movements and collisions between them
High Contrast Non-Linear (3)	Non-linear, varying-speed targets at the highest contrast, i.e. no backlighting
Low Contrast Non-Linear (3)	Non-linear, varying-speed targets at contrast level 4 (chosen due to accuracy loss at the lowest contrast level)
Changing Direction (1)	Target with persistent changes in motion, a spiraling motion, at highest contrast level
Increased Distance (1)	Non-linear, varying-speed target with 5m distance to target and highest contrast level

The ground-truth for each test was generated using built-in OpenCV functions to locate contours and calculate minimum enclosing circles on noise-suppressed event frames [26]. Intersection-over-union and active tracking time metrics were then calculated by comparing the intersection of the event-based Gaussian blob trackers with the set of minimum enclosed circles. It should be noted that although the Gaussian trackers are elliptical in shape, the IoU metric was calculated using a circle circumscribed with the tracker’s larger axis as its diameter. This approach results in a decrease in the IoU scores calculated but should not impact the active tracking time measured in each experiment since active tracking is determined by the tracker simply intersecting with the ground-truth circles across frames. Contrast levels were calculated as the ratio of laser pointer luminance to background lighting luminance as measured via a luxmeter.

4. RESULTS

The performance of the frame-based and asynchronous approaches was assessed with each combination of contrast, speed, and distance conditions previously presented as well as the additional 11 non-linear movement trials. Before evaluating accuracy across the experiments, threshold values used for noise-filtering in the asynchronous method were first tuned by inspection for each experiment. Filtering thresholds were chosen in order to maximize tracking accuracy in terms of percentage of actively tracked frames and with regard to the different speeds of the target. In general, slower targets necessitate a larger upper-bound threshold, δ_2 , since they will repeatedly trigger events in a small range of pixels. In the same respect, the lower-bound threshold, δ_1 , can be raised as well in order to filter additional non-target events without interfering with accurate target tracking. Due to the frame-based method filtering events according to the number of spatial and temporal neighbors, no tuning for the frame-based methods was required.

4.1 Noise Filtration

Figure 1 depicts the percentage of events that were filtered as noise by the asynchronous and frame-based approaches in each of the experiments. As target speed does not affect the number of events registered overall, filtered event percentages are averaged across all speed modes in experiments with the same contrast ratio.

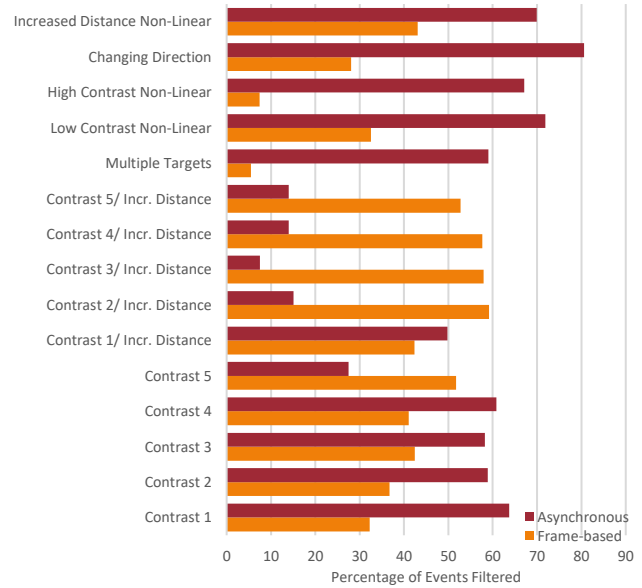


Figure 1: Noise Filtration by Experiment

Across most experiments, the asynchronous approach showed much larger percentages of filtered events than the frame-based approach. In some of the additional trials, the asynchronous approach reached up to approximately 80% average events filtered, such as in the case of the Changing Direction trial. Conversely, the frame-based approach had a maximum of only about 60% filtered events occurring in the trial with second contrast level and increased distance. The fact that the two approaches percentage of filtered events differed greatly between the different types of experiments also indicates that the two approaches differ in the context in which events are filtered. Since the frame-based approach filters events with few neighbors over time, this behavior suggests that the trials with increased distance exhibited far more solitary events as reflected in the much larger percentage of filtered events. By contrast, the multiple target trials had many repetitive events occurring spatially and temporally close, resulting in extremely low filtering rates for the frame-based method but much higher rates with the asynchronous method. While the total number of events observed in each experiment is largely determined by the length of the recording, variations in event numbers were observed with differing contrast levels as well. Although the 11 separate trials had varying recording lengths, the original contrast and velocity-controlled trials had roughly the same recording time across each. This fact indicates that lower contrast between the target area and background resulted in larger numbers of events that can be attributed to erroneous events occurring in the background. In ground applications, this phenomenon could be mitigated by significantly raising the sensitivity thresholds of the sensor itself, but this would not be possible on a space platform where maximum sensitivity is required to detect exceedingly dim and distant objects.

4.2 Performance

Figure 2 depicts the performance in average frames per second (FPS) for the two approaches as a function of the number of events generated. These performance numbers are drawn from each of the experiments previously detailed. It should be noted that the number of events shown is before noise filtration, which partially explains the difference in FPS for trials with approximately the same number

of events. All trials were run using a four-core Intel Core i5-8250U 1.6 GHz processor. The Bezier curve-fitting, frame drawing and frame-based noise-suppression portions of the algorithm were all parallelized using the C++ OpenMP API. Since the frame-based method involves applying a kernel convolution to the entirety of reconstructed frames, the same number of calculations must be made regardless of the number of events in the event stream. As a result, the performance would be expected to be relatively constant across all experiments. Although execution times do appear to be consistent in the frame-based approach, some trials that exhibited larger or smaller numbers of unfiltered events, as well more trackers being instantiated, resulted in small disparities in average execution time. Conversely, the asynchronous approach shows much greater variation in FPS versus the number of events, which is a result of the larger number of events filtered before tracking is performed. The number of inactive trackers instantiated also affects asynchronous performance, further adding to the FPS variation for experiments with approximately the same number of events.

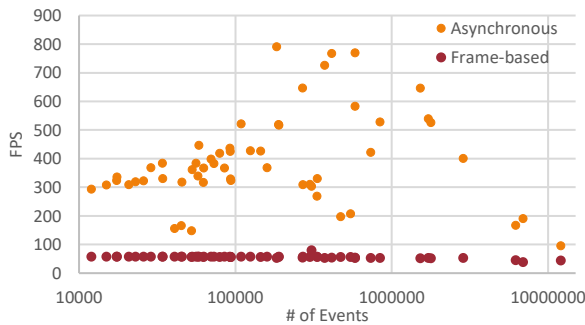


Figure 2: Performance in FPS vs. Number of Events

4.3 Tracker Instantiation and Suppression

Figures 3 and 4 display the number of trackers generated for each experiment and number of trackers suppressed in the asynchronous case respectively. In Figure 3, the number of trackers generated is reported as an average across all experiments with the same contrast level as the number of trackers instantiated is primarily dependent on the contrast. By contrast, Figure 4 shows the average number of trackers suppressed with experiments grouped by target speed since the asynchronous tracker suppression varies more with the speed of the target rather than brightness contrast. It should be noted that the vast majority of the trackers instantiated are inactive throughout most of the experiments and do not necessarily indicate a false positive in tracking.

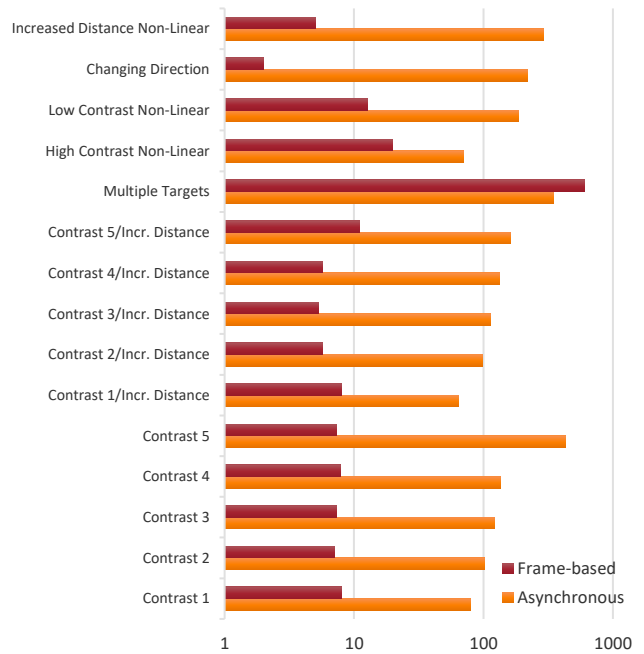


Figure 3: Average Number of Trackers Generated per Experiment Group

As expected, the asynchronous approach spawned considerably more trackers in each trial as a result of some noise events passing through filtration. However, there are a few notable exceptions, such as the multiple target trials, where the much higher percentage of events filtered by the asynchronous method resulted in comparatively fewer trackers being generated. Several trials at the slowest speed also exhibited no trackers instantiated by the frame-based method, which represented a complete failure to track the target. This failure was due primarily to the exceedingly small number of events generated by the target, which were in turn filtered by the frame-based noise suppression.

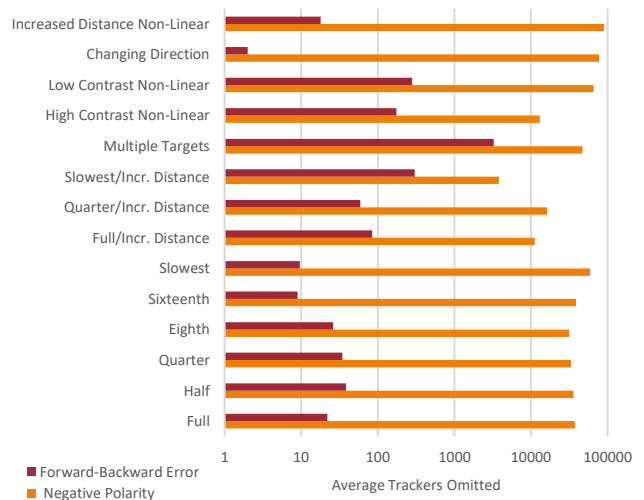


Figure 4: Number of Trackers Suppressed in Asynchronous Approach

It is important to note that the number of tracker omissions is cumulative across each trial, hence the number of omissions being greater than the number of trackers originally initialized. Furthermore, negative polarity thresholding is applied prior to forward-backward error thresholding, leading to the larger number

of omissions being attributed to negative polarity. In some trials, the negative polarity threshold successfully suppressed all erroneous trackers, leaving none to be suppressed as a result of forward-backward error. Additionally, there were several instances in the trials with the slowest target speed and increased distance where no trackers were suppressed by negative polarity thresholding. During these trials, the exceptionally slow targets exhibited few or no negative polarity events during movement, which necessitated reducing the negative polarity threshold to zero.

4.4 Experimental Tracking Accuracy

Figures 5 and 6 show the general trends in tracking accuracy with averaged values across experiments with the same target speed. Tracking accuracy was measured as the intersection over union with the ground-truth position of the target per frame, while active tracking time refers to the percentage of frames where an active tracker intersected with the ground-truth target position. Note that active tracking time specifically means that the tracker was active and not suppressed, so frames that were not actively tracked may have still contained an inactive or suppressed tracker on target. This impact of tracker suppression on active tracking time can be observed in the relatively lower active tracking time of the asynchronous method for lower velocity targets. Although the target tracker was suppressed in some frames, the asynchronous method maintained relatively the same or better IoU than the frame-based method with the same velocity targets.

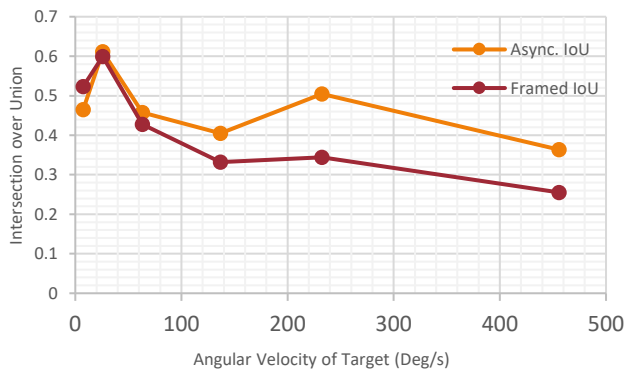


Figure 5: Intersection over Union vs. Target Speed

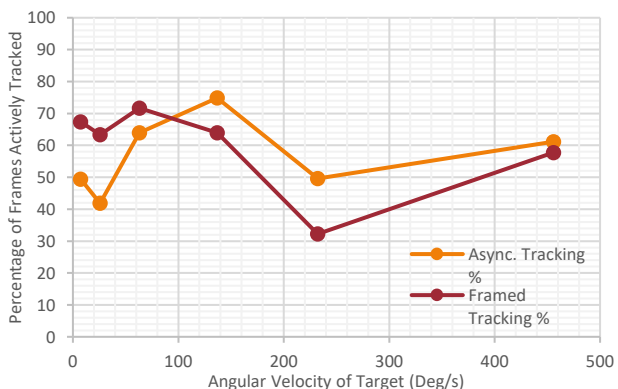


Figure 6: Percentage of Actively Tracked Frames vs. Target Speed

Figures 7 and 8 show the same accuracy metrics measured for the 11 additional trials with multiple targets and changes of target direction. Since these trials did not have precise measurements of target velocity, they are presented per experiment. The additional

trials showed similar trends with the asynchronous approach showing better IoU measures across almost all of the experiments, but with comparable or relatively smaller percentage of actively tracked frames. While the frame-based method exhibited some instances of slightly better IoU, the asynchronous approach showed greatly improved IoU for trials with multiple targets.

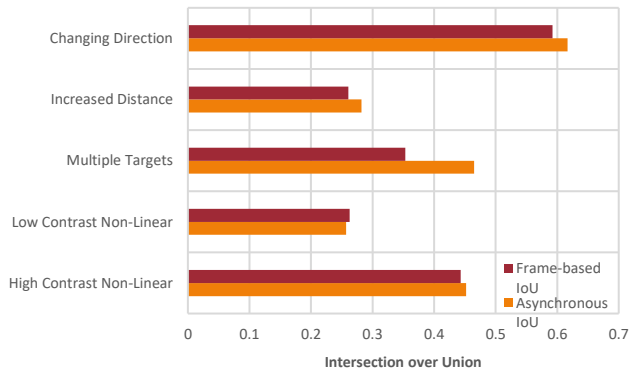


Figure 7: Intersection Over Union per Experiment

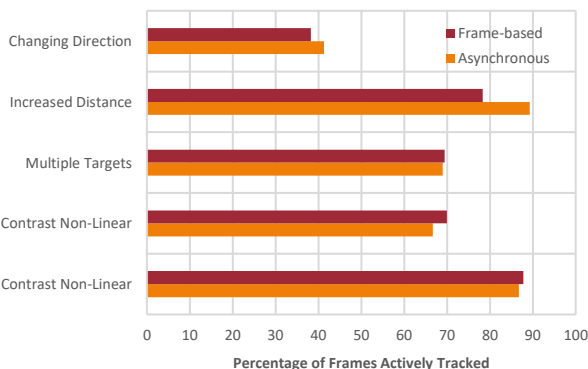


Figure 8: Percentage of Actively Tracked Frames per Experiment

Although the algorithm parameters were tuned to maximize tracking accuracy in the previous experiment analysis, additional metrics were recorded to assess tracking effectiveness within a range of tracking parameters. The graphs displayed in Figures 9 and 10 show averaged accuracy results with varied position update parameters for the trackers. The position update parameter refers to α_1 from Eqn. 4, which controls the magnitude of tracker displacement when a new event is associated with it. In general, larger α_1 values cause trackers to require greater number of events to result in significant movement of the tracker. This additional analysis was conducted with trial recordings at the highest level of contrast, i.e. no background lighting, but with varied speed and distance, in order to determine the effect of the position update factor on tracker accuracy. Disregarding the base differences between the accuracies of the two approaches, both methods showed the same general trends as a result of increasing the value of the position update parameter. In general, the range of 0.4 to 0.7 showed minor differences in both IoU and actively tracked frames, but both accuracy measures showed significant improvement beginning at a value of 0.8, with some exceptions. Interestingly, the increased distance trials showed comparable or even better accuracy with much smaller update values and diminishing accuracy with increased values. This trend can be explained by the much smaller number of events seen in these trials, and thus the

corresponding targets were composed of far fewer events as well. As a result, the position update parameter must be smaller such that fewer events are required for the trackers to accurately track the targets. Since these reported values are averaged across all trials, the average accuracy metrics are somewhat skewed due to the relatively low accuracies measured in the slowest target trials.

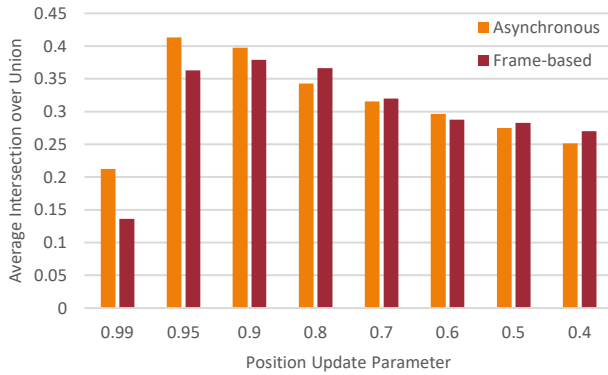


Figure 9: Intersection Over Union vs. Position Update Parameter




Figure 10: Percentage of Actively Tracked Frames vs. Position Update Parameter

4.5 Visualized Tracking Results

For demonstration purposes, Figure 11 (a-c) depicts the reconstructed frame output with ground-truth, frame-based, and asynchronous tracking results superimposed for several trials, while Figure 12 (a-c) shows the corresponding unfiltered frames. Positive- and negative- polarity events are indicated by black and grey pixels respectively. The ground-truth minimum enclosed circle is denoted by a red circle, while the frame-based and asynchronous trackers are overlaid with brown ellipses. Each of the trackers that is deemed to be tracking the target also displays a blue Bezier curve fit to its predicted trajectory obtained via DESP. As referenced earlier, both asynchronous and frame-based have many inactive trackers that are hidden in the reconstructed frames though only active trackers are displayed. The events displayed in the frames are the result of the asynchronous filtering technique and several erroneous active trackers can be seen in each. These erroneous trackers belong only to the asynchronous approach; however, these trackers have no estimated trajectory and promptly become inactive between frames. Trajectories are drawn for both the asynchronous and frame-based approaches, which can be seen to diverge most significantly in the trials containing multiple targets.

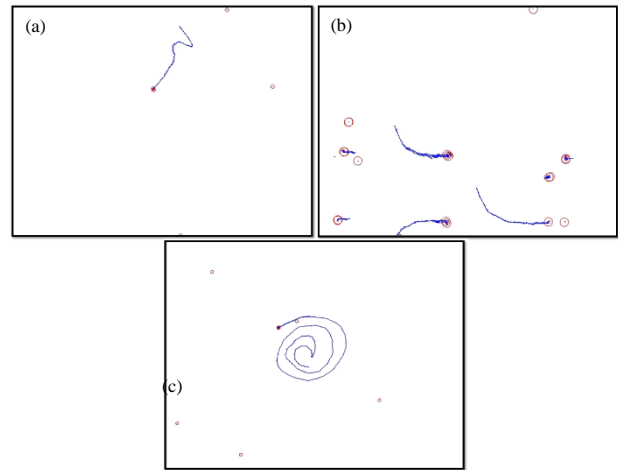


Figure 11: Filtered and Tracked Frames (a) High Contrast 3 Exp. (b) Multiple Targets 1 Exp. (c) Changing Direction Exp.

5. (a) DISCUSSION

(b)

Overall, the experimental results exhibited trends that were to be expected when comparing the asynchronous and frame-based approaches. In terms of average performance, the asynchronous approach had an average of $7\times$ better performance but had a much larger variation in relation to the total number of events. The reason for this disparity is primarily the differing forms of noise suppression used by both algorithms. The frame-based approach will always perform approximately the same number of comparisons whereas the asynchronous approach leverages the sparsity of the event stream to dramatically reduce computation time as evidenced by the differing trends in Figure 3. The two forms of noise suppression also resulted in very different ratios of unfiltered to filtered events, with the asynchronous approach filtering out many more events across all experiments. The multiple targets trials are an outstanding example of this with the asynchronous approach exhibiting an average of about 60% of filtered events compared to the frame-based approach's mere 6%. This large discrepancy is most likely due to the asynchronous approach also filtering events that may not necessarily be noise, but constituent events of potential targets. However, for the purposes of this research, filtering redundant events can reduce computation time without negatively impacting the tracking performance itself.

In addition to large differences in the number of filtered events, the filtering techniques also differed greatly in terms of the number of trackers instantiated. The asynchronous algorithm exhibited many more trackers being instantiated across all trials as compared to the frame-based approach. According to both algorithms, new trackers will be created if no pre-existing tracker has a p score above the predefined threshold. The asynchronous approach having more instantiated trackers would indicate more solitary events passing through the filter, though these trackers are largely hidden as they are correctly deemed to be associated with noise. In the asynchronous approach, the negative-polarity threshold served to suppress the majority of erroneous trackers. However, there were several trials in which the target had few or no negative-polarity events associated with it and which relied entirely on the forward-backward error threshold to filter noise-attracted trackers. These occurrences highlighted the need for both forms of tracker suppression as well as the necessity of tuning the suppression thresholds to the behavior of the intended target.

In regard to tracker accuracy, the asynchronous approach displayed either comparable or, in some cases, superior results in the IoU measurements. While the IoU scores for both approaches did fall below what might be considered an acceptable level of 50% in several trials, these results can be explained by a number of factors. As mentioned earlier, the IoU metric was calculated using the circle circumscribed around the larger axis of each elliptical Gaussian tracker. This approximation was made to ease the calculation involved in the metric as well as due to the assumption that the target is circular and would not cause significant change in tracker shape. However, this assumption was not always necessarily the case, especially with the additional non-linear and multiple target trials. As a result, the area of the tracker would be overestimated and cause the IoU metric to be much smaller than expected.

Another issue with the IoU metric is the use of reconstructed frames to form the ground-truth comparison. Since the ground-truth position of the target is evaluated per frame, events must be accumulated over the chosen integration time in order to

reconstruct a conventional image frame. However, sufficiently fast-




Figure 12: Original Unfiltered Frames (a) High Contrast 3 Exp. (b) Multiple Targets 1 Exp. (c) Changing Direction Exp

moving targets, such as those in the highest speed trials, would exhibit motion blur due to events from multiple positions of the target falling into the same reconstructed frame. While the event-based tracking would track these targets accurately due to the serial nature of the tracker update, the ground-truth comparison would fit a larger than necessary circle around the target and thus artificially decrease the IoU metric. Nonetheless, the IoU metric still shows that the asynchronous noise suppression did not result in a loss of tracker accuracy compared to the frame-based approach. Conversely, the percentage of actively tracked frames does seem to favor the frame-based approach across many of the experiments. This general trend is largely due to the tracker suppression employed by the asynchronous approach, where the target tracker may be made hidden if the negative-polarity activity falls below threshold or the forward-backward error rises above threshold. However, the tracker is not destroyed and positions that may not be actively tracked are in most cases correctly interpolated by the smoothing and curve-fitting operations. Even so, the asynchronous approach showed comparable active tracking time across most experiments.

With respect to varying the Gaussian tracker position update parameter, both approaches showed similar trends with respect to the resulting tracking accuracy. As previously mentioned, the position update weight tunes the number of events that are required to shift the center of the tracker. As a result, the noise suppression employed and size and speed of the target influence the optimal parameter value required to maximize tracking accuracy. During testing, the frame-based approach showed peak accuracy with position update parameters, α_1 , between 0.9 and 0.95, while the asynchronous approach exhibited better accuracy with value 0.8 for several trials. This disparity is most likely due to the asynchronous noise suppression, which as previously mentioned likely removed target events in addition to noise. As a result, the update factor would need to be lower in order to weight each event of the target greater when updating position and maintain accuracy. This behavior indicates that the asynchronous approach is more sensitive to changes in algorithm parameters and would likely need to be tuned according to the expected targets of interest.

6. CONCLUSIONS

This research has introduced an asynchronous noise and tracker suppression scheme for Gaussian blob point-source tracking in a space environment. Experimental results show comparable or superior tracking accuracy with the asynchronous approach, while also exhibiting dramatically better performance as compared to the frame-based approach. Furthermore, unlike the frame-based and other previous event-based noise-suppression techniques, the asynchronous approach has the ability to track objects that may only appear as a single pixel in the DVS's view. Given the asynchronous method's reliance on multiple parameters, it is necessary for the algorithm to be tuned for detecting and tracking specific types of objects in terms of speed, size and relative luminance. Additionally, future modifications to the asynchronous method could be made that suppress trackers based on velocity, number of events, tracking time, etc. in order to track specific types of objects in space or on Earth.

7. ACKNOWLEDGMENTS

This work was supported by the NSF SHREC Center members and by the IUCRC Program of the National Science Foundation under Grant No. CNS-1738783. The authors would like to thank Dr. Ryad Benosman and Prohsee for providing the DVS used in this work.

8. REFERENCES

- [1] M. Mahowald, "VLSI analogs of neuronal visual processing: a synthesis of form and function," 1992.
- [2] T. Delbruck, B. Linares-Barranco, E. Culurciello, and C. Posch, "Activity-driven, event-based vision sensors," *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, 2010.
- [3] C. Posch, D. Matolin, and R. Wohlgenannt, "An asynchronous time-based image sensor," *2008 IEEE International Symposium on Circuits and Systems*, 2008.
- [4] F. Bennet, C. D'Orgeville, Y. Gao, W. Gardhouse, N. Paulin, I. Price, F. Rigaut, I. T. Ritchie, C. H. Smith, K. Uhlendorf, and Y. Wang, "Adaptive optics for space debris tracking," *Adaptive Optics Systems IV*, 2014.
- [5] J. Sharma, "Space-Based Visible Space Surveillance Performance," *Journal of Guidance, Control, and Dynamics*, vol. 23, no. 1, pp. 153–158, 2000.
- [6] G. Cohen, S. Afshar, B. Morreale, T. Bessell, A. Wabnitz, M. Rutten, and A. V. Schaik, "Event-based Sensing for Space Situational Awareness," *The Journal of the Astronautical Sciences*, vol. 66, no. 2, pp. 125–141, Mar. 2019.
- [7] F. Wang and V. D. Agrawal, "Single Event Upset: An Embedded Tutorial," *21st International Conference on VLSI Design (VLSID 2008)*, 2008.
- [8] G. Hopkinson, "Radiation effects in a CMOS active pixel sensor," *IEEE Transactions on Nuclear Science*, vol. 47, no. 6, pp. 2480–2484, 2000.
- [9] V. Goiffon, M. Estribeau, O. Marcelot, P. Cervantes, P. Magnan, M. Gaillardin, C. Virmondois, P. Martin-Gonthier, R. Molina, F. Corbiere, S. Girard, P. Paillet, and C. Marcandella, "Radiation Effects in Pinned Photodiode CMOS Image Sensors: Pixel Performance Degradation Due to Total Ionizing Dose," *IEEE Transactions on Nuclear Science*, vol. 59, no. 6, pp. 2878–2887, 2012.
- [10] S. Buchner, D. McMorrow, N. Roche, L. Dusseau, and R. L. Pease, "The Effects of Low Dose-Rate Ionizing Radiation on the Shapes of Transients in the LM124 Operational Amplifier," *IEEE Transactions on Nuclear Science*, vol. 55, no. 6, pp. 3314–3320, 2008.
- [11] G. Reitz, R. Beaujean, E. Benton, S. Burmeister, T. Dachev, S. Deme, M. Luszik-Bhadra, and P. Olko, "Space radiation measurements on-board ISS—the DOSMAP experiment," *Radiation Protection Dosimetry*, vol. 116, no. 1–4, pp. 374–379, 2005.
- [12] D. Tedaldi, et. al. "Feature detection and tracking with the dynamic and active-pixel vision sensor (DAVIS)." *Second International Conference on Event-based Control, Communication, and Signal Processing (EBCSP)*. 2016.
- [13] T. Delbruck, "Frame-free dynamic digital vision," in *International Symposium on Secure-Life Electronics Advanced Electronics for Quality Life and Society*, 2008.
- [14] R. Benosman, S.-H. Ieng, C. Clercq, C. Bartolozzi, and M. Srinivasan, "Asynchronous frameless event-based optical flow," *Neural Networks*, vol. 27, pp. 32–37, 2012.
- [15] Z. Ni, C. Pacoret, R. Benosman, S. Ieng, and S. Régnier, "Asynchronous event-based high speed vision for microparticle tracking," *Journal of Microscopy*, vol. 245, no. 3, pp. 236–244, 2011.
- [16] S. Seifozzakerini, et. al. "Event-Based Hough Transform in a Spiking Neural Network for Multiple Line Detection and Tracking Using a Dynamic Vision Sensor." *BMVC*. 2016.
- [17] X. Lagorce, C. Meyer, S.-H. Ieng, D. Filliat, and R. Benosman, "Asynchronous Event-Based Multikernel Algorithm for High-Speed Visual Features Tracking," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 26, no. 8, pp. 1710–1720, 2015.
- [18] Mueggler, Elias, et al. "Lifetime estimation of events from dynamic vision sensors." *IEEE international conference on Robotics and Automation (ICRA)*. IEEE, 2015.
- [19] J.-M. Maro, G. Lenz, C. Reeves, and R. Benosman, "Event-based Visual Gesture Recognition with Background Suppression running on a smart-phone," *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, 2019.
- [20] X. Xie, et al. "DVS image noise removal using K-SVD method." *Ninth International Conference on Graphic and Image Processing (ICGIP 2017)*. Vol. 10615. International Society for Optics and Photonics, 2018.
- [21] X. Xie, J. Du, G. Shi, H. Hu, and W. Li. "An Improved Approach for Visualizing Dynamic Vision Sensor and its Video Denoising," *Proceedings of the International Conference on Video and Image Processing (ICVIP 2017)*. Association for Computing Machinery. pp. 176–180, 2017.
- [22] D. Comaniciu, V. Ramesh, and P. Meer, "Kernel-based object tracking," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 5, pp. 564–575, 2003.
- [23] J. J. Laviola, "Double exponential smoothing: an alternative to Kalman filter-based predictive tracking," *Proceedings of the workshop on Virtual environments 2003 - EGVE 03*, 2003.

- [24] Z. Kalal, K. Mikolajczyk, and J. Matas, "Forward-Backward Error: Automatic Detection of Tracking Failures," *2010 20th International Conference on Pattern Recognition*, 2010.
- [25] T. Böttger, P. Follmann, and M. Fauser, "Measuring the Accuracy of Object Detectors and Trackers," *Lecture Notes*

in Computer Science Pattern Recognition, pp. 415–426, 2017.

- [26] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.