EventShiftFlow

Towards Hardware-Efficient FPGA-based Flow Estimation

Arianna Alonso Bizzi · Fernando Cladera · C. J. Taylor
University of Pennsylvania, Philadelphia, PA · alari@seas.upenn.edu

Penn Engineering
GRASP LABORATORY
General Robotics, Automation, Sensing & Perception Lab

ICRA
IEEE INTERNATIONAL CONFERENCE ON ROBOTICS AND AUTOMATION

Key Takeaway

EventShiftFlow recovers useful flow direction and magnitude using no dividers, no DSP blocks, and no floating-point arithmetic, using **under 2 kB of on-chip storage**. This represents **orders of magnitude lower resource usage** than EDFLOW (855 kB, 669 DSPs), enabling **real-time flow estimation on the smallest neuromorphic robotic platforms.**

Motivation

Event cameras detect per-pixel brightness changes asynchronously with **microsecond temporal resolution, high dynamic range, and low power consumption**. This makes them attractive for low-latency motion estimation on **size-, weight-, and power-constrained (SWaP) platforms** such as micro aerial vehicles.

Existing methods

Plane fitting, contrast maximization, and deep learning require **floating-point arithmetic, iterative optimization, or large memory buffers…**

…Infeasible on small chips or FPGAs!

Research question

How can we estimate motion using the smallest compute footprint possible, suitable for SWaP-constrained uses?

Algorithm: Bitvector-based motion estimation

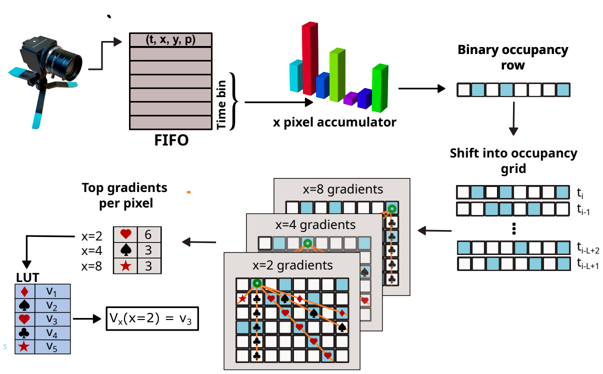

Events → fixed-duration time bins → 1-bit occupancy grid (thresholded event counts/pixel/axis) → velocity hypotheses scored by counting coincidences along diagonal traces through a shift-register grid - pure integer arithmetic. Uses only shift registers, counters, and comparators.

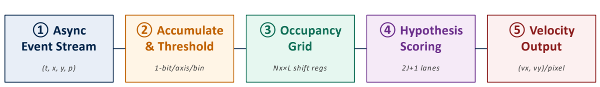

Streaming pipeline. X and Y axes processed identically in parallel.

The case for adaptive time bins

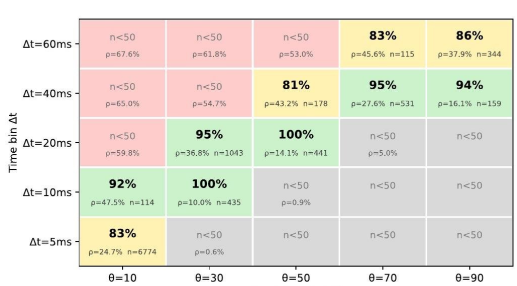

Results report directional accuracy on the Mueggler et al. Rotating Shapes Dataset, where ρ is occupancy density and n is the number of detections. Best performance at occupancy density $\rho \in [10\%, 40\%]$. Adaptive Δt via density feedback eliminates manual tuning.

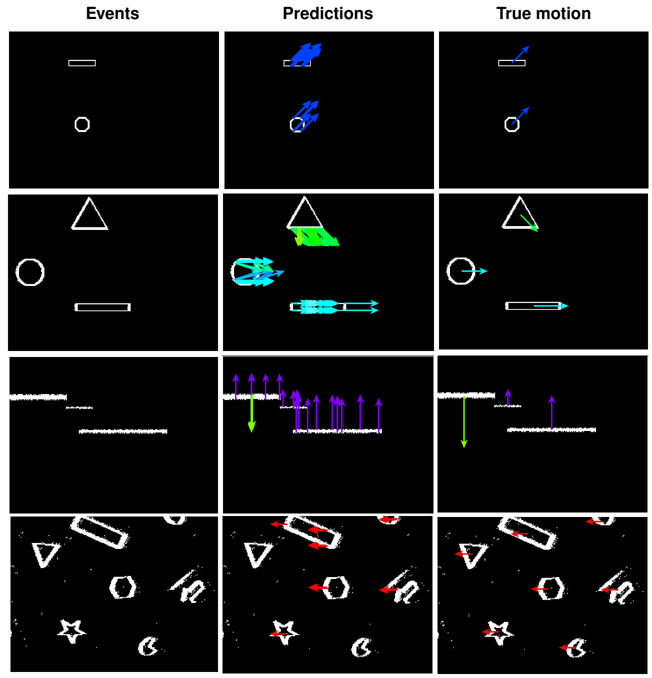

Events, predicted flow, and ground-truth for 6 test scenarios. Rows 1–3 synthetic, Row 4 real data (Mueggler et al.)

Key: *Length of arrow ∝ speed*

Algorithm Parameters

Symbol	Parameter	Value
Δt	Time bin duration	5–50 ms
θ_e	Event threshold	30–100
L	Temporal depth	16 bins
θ_s	Score threshold	0.3 L
J	Hypotheses	±15 px/bin
β	Min in-bounds steps	4

Resource Comparison

Method	Block RAM	DSPs	Storage	FP/Div
EDFLOW [Liu'22]	390 (855 kB)	669	855 kB	Yes
Aung et al. [2018]	138	16	~270 kB	Yes
EventShiftFlow (ours)	**0**	**0**	**< 2 kB**	**No**

99.5% directional accuracy (real data)

< 2 kB on-chip storage (both axes)

210 ns per-pixel latency @ 100 MHz

0 DSP blocks & dividers

Artix-7 FPGA Prototype

FPGA Device	Artix-7 xc7a100tftg256-2
Tool	Vivado 2025.2
Clock	100 MHz (+0.030 ns slack)
Scoring Latency	2400 cycles = 24 µs @ 100 MHz
Total on-chip Power	0.142 W

Slice LUTs	13,326 (~21% of device)
Slice Registers	5,517 (~4%)
Block RAM	0 tiles
DSP Blocks	0
Axis	x-axis (prototype)

Prototype includes full UART-to-motion path (event framing, AXI FIFO, LED debug). Core motion-scoring datapath uses only a fraction of the reported resources.

Limitations

Velocity quantization

Velocity estimates are limited to the tested delay hypotheses. Increasing resolution requires more hypotheses and hardware resources.

Sensitivity to scene complexity

Robustness in cluttered and dynamic environments remains to be validated.

Sensitivity to event density

Performance depends on occupancy density within each time bin. Requires adaptive time-binning.

Future Work

2-Axis FPGA prototype

Extend x-axis design to full 2D; validate with a live event camera stream.

Temporal regularisation

Majority voting, exponential smoothing, or score bonus for the previous winner. Add "momentum".

Polarity-aware scoring

Exploit sign consistency of moving edges to improve discrimination at no additional memory cost.

References: Liu & Delbruck, EDFLOW, IEEE TCSVT 2022 · Zhu et al., EV-FlowNet, RSS 2018 · Gehrig et al., E-RAFT, 3DV 2021 · Mueggler et al., RPG Dataset, IJRR 2017 · Rebecq et al., E2VID, CVPR 2019

Project: alonsobizzi.github.io/eventshiftflow-site