

Real-time event-based particle image velocimetry for active flow control

Christian Willert¹, Luca Franceschelli², Enrico Amico³, Marco Raiola²,
Gioacchino Cafiero³ and Stefano Discetti²

¹DLR Institute of Propulsion Technology, German Aerospace Center, Köln, Germany

²Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés,
Madrid, Spain

³Mechanical and Aerospace Engineering Dept., Politecnico di Torino, Torino, Italy

E-mail: chris.willert@dlr.de

Abstract. This work investigates event-based vision (EBV) as a tool for real-time flow diagnostics in configurations analogous to two-dimensional, two-component particle image velocimetry. Owing to its reduced data stream compared to conventional frame-based imaging, EBV enables kilohertz-rate pseudo-framing and efficient processing on standard computing hardware. A pseudo-frame-based implementation called real-time event-based imaging velocimetry is presented, capable of delivering velocity fields at several hundred hertz with $\mathcal{O}(10^6)$ vectors per second.

The concept is experimentally demonstrated on a small-scale jet in water, where event rates above 100×10^6 events/s and online processing at 250 – 700 Hz are achieved depending on seeding and interrogation settings. Beyond these validation cases, two active flow control applications on a jet in air are illustrated: open-loop optimization of jet mixing using Bayesian optimization, and closed-loop control of a water jet using reinforcement learning. These results highlight EBV as a cost-effective and scalable sensing technology with strong potential for real-time feedback in flow control.

1 Introduction

The active manipulation of flow fields to achieve specific performance objectives is of great technological importance [1]. A fundamental requirement for the success of such strategies is the availability of measurement techniques that can provide meaningful flow information both quickly and reliably. For optimization tasks, the measurement system must be fast and data-efficient to allow rapid exploration of actuation strategies. In contrast, closed-loop control requires strict real-time capability, since actuation must continuously adapt to the instantaneous flow state.

Spatially resolved diagnostics such as particle image velocimetry (PIV) are attractive in this context because they deliver velocity fields rather than single-point signals, greatly increasing the observability of the flow. This broader view provides direct access to flow structures and their evolution, offering more informative feedback for both optimization and closed-loop strategies [2]. However, the application of PIV to real-time flow control is hindered by the large amount of data, the complexity of the required algorithms, and the limitations of framing cameras [3]. Reported real-time particle image velocimetry (RT-PIV) implementations based on conventional image cross-correlation typically achieve field rates on the order of 5–25 Hz [4–8]. Higher rates have been demonstrated using optical flow approaches on GPUs [9, 10], with Gautier and Aider [9] reporting up to 220 Hz. More recently, Bollt et al. [11] achieved kilohertz-rate velocity field processing by exploiting NVIDIA’s built-in GPU optical flow engine, refining



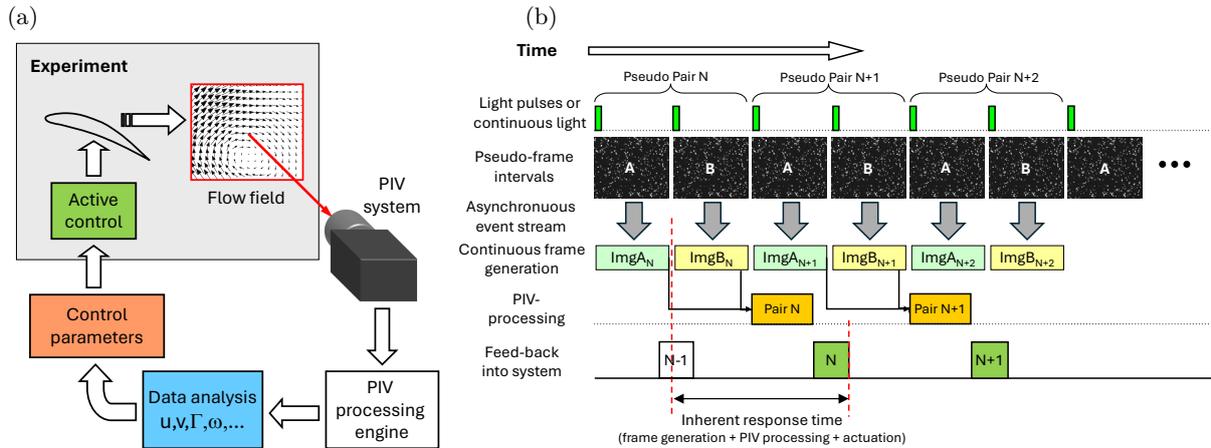


Figure 1: (a): possible implementation of active flow-control based on real-time flow field data (from [5]); (b): data stream for pseudo-frame-based, real-time event-based imaging velocimetry. Contrary to framing cameras the pseudo-frames are available immediately after the end of framing period, thereby reducing system latency.

the estimates through an additional cross-correlation step. When the flow configuration is well defined, processing can be accelerated by limiting the analysis to a reduced region of interest, by extracting data along a single line, or by using coarse-grained PIV estimates [12, 13]. Following this principle, Nonomura et al. [14] achieved closed-loop control at frame rates up to 2000 Hz by calculating only a few velocity vectors per frame and reconstructing the remaining flow field from them. However, such approaches are inherently application-specific and may overlook flow features outside the predefined region, which limits their general applicability in complex or turbulent flows. Nevertheless, the need to acquire and read out full image frames imposes an unavoidable latency in all framing-camera-based approaches.

Event-based vision (EBV) offers a promising alternative. Unlike conventional cameras, EBV sensors asynchronously report only changes in pixel intensity, generating a continuous stream of events with reduced data rate, high dynamic range, and microsecond-scale temporal resolution [15]. These properties make EBV well suited for real-time flow diagnostics. Yet, most EBV-based velocimetry methods reported to date operate in an off-line fashion, where computational cost is of lesser importance. A notable exception is the event-based three-dimensional Lagrangian particle tracking (3D-LPT) implementations of Rusch and Rösger [16], which reconstructs time-averaged flow fields in real time. However, they can only track a limited number of particles simultaneously ($\mathcal{O}(50-100)$), which is insufficient for dense velocity fields in flow control scenarios.

Two general strategies can be followed when processing EBV data (Fig. 1). The first is to group incoming events into isolated clusters (or blobs) representing individual particles and to track these clusters in space and time [17], resulting in Lagrangian trajectories [18, 19]. While appealing, this approach scales poorly with the number of particles: even with optimized implementations (e.g. OpenCV or `scipy.ndimage.label`), cluster detection for $\mathcal{O}(1000)$ particles requires several milliseconds and increases with particle density. The subsequent tracking step and the required interpolation to Eulerian fields for control applications further add to the computational effort.

The second strategy is to partition the incoming event stream into short, fixed-duration time slices, thereby forming so-called pseudo-frames [20–22]. All events within a slice are assigned the same timestamp, typically synchronized to the pulsed laser illumination, analogous to an image produced by a framing camera. Although this approach does not exploit the full asynchronous nature of the EBV sensor, the generation of pseudo-frames is considerably more efficient than continuously extracting and tracking clusters over time. Pseudo-frame generation scales with event rate, but much less strongly than cluster detection schemes, and the resulting frames can be processed using conventional PIV algorithms that are readily parallelized. This results in a high frame throughput that depends only mildly on the event data rate, that is, on the particle image density. Overall, the pseudo-frame strategy offers several advantages: correlation processing is inherently parallelizable, performance is less sensitive to seeding density, and latency is reduced. In particular, the second pseudo-frame of a pair is immediately available at the end of the slicing interval, whereas a conventional camera still requires an additional read-out step (see Fig. 1).

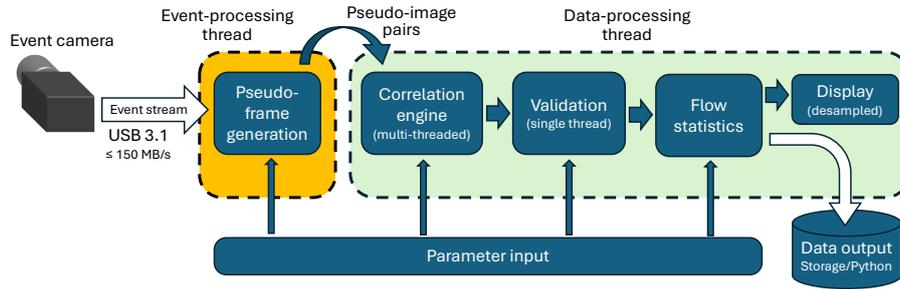


Figure 2: Processing pipeline for RT-EBIV.

The present study investigates the potential of this second approach by introducing a pseudo-frame-based real-time event-based imaging velocimetry (RT-EBIV) system. The implementation accumulates pseudo-image pairs at a defined slicing interval and processes them using a single-pass correlation algorithm with subsequent validation. Bench-top experiments on a rotating disk and a water jet are used to assess performance, demonstrating velocity field rates of several hundred Hertz and highlighting RT-EBIV as a viable tool for active flow control.

In addition to these validation cases, two control-oriented applications are briefly illustrated. The first concerns the optimization of an air jet flow, where event-based imaging velocimetry (EBIV) serves as a fast and data-efficient cost-function estimator for exploring actuation strategies. The second involves a closed-loop experiment on a submerged water jet, in which real-time velocity fields from RT-EBIV are used directly for feedback to adapt the actuation. These examples emphasize the dual potential of RT-EBIV: enabling rapid flow optimization on the one hand, and supporting closed-loop control on the other.

2 Implementation

A simplified outline of the processing steps for RT-EBIV is provided in Fig. 2. For optimal performance the software is realized in C/C++ making use of modern multi-processor architecture and advanced CPU commands (e.g. SIMD instructions). Incoming event-data is accumulated for a defined time interval to generate pseudo-frames, ideally synchronized to the pulses driving the light source (e.g. laser). Each completed image-pair is then handed over to the cross-correlation engine, that operates independently in a separate thread to perform a single-pass interrogation with a fixed window size and sample spacing, typically 32×32 pixel. As each correlation plane is independent of the others, this step can be performed in parallel using OpenMP instructions [23]. Here, optimum processing speed is achieved when the number of nodes is a multiple of the utilized cores; otherwise, a few unfinished threads block downstream processing until done.

The validation on the recovered displacement field is based on 3×3 neighborhood operators such as displacement magnitude, vector difference magnitude or normalized median filtering [24]. Online processed flow field data along with statistical information on the captured events is continuously stored in a circular buffer that can be saved for further analysis.

For the feasibility measurements presented herein, event cameras based on the 4th generation EBV sensor IMX636 are used, which is a co-development by Sony and Prophesee [25]. Initial measurements were performed with a Metavision EVK4 by Prophesee or SilkyEvCam-HD by Century-Arks, which are equivalent in packaging and size. Compared to these models, the recently introduced uEye-EVS by IDS Imaging features reduced latency, but it is otherwise identical in terms of its specifications. All cameras use an USB3.1 interface for data transmission and have a power consumption of the order of 1 W to 3 W. The IMX636 sensor features pixels with a size of $4.86 \mu\text{m}/\text{pixel}$ on a grid of 1280×720 pixels. Although the sensor is specified for a maximum throughput exceeding 1×10^9 events per second (1 GEV/s), numerous experiments by the authors have led to the conclusion that the devices can provide a sustained data rate of up to 50-60 MEV/s without losing events [20]; under certain conditions even 100 MEV/s can be reached. In the context of EBIV this is an important figure because it defines the maximum number of features (i.e. particle images) that can be reliably captured as described in the following section.

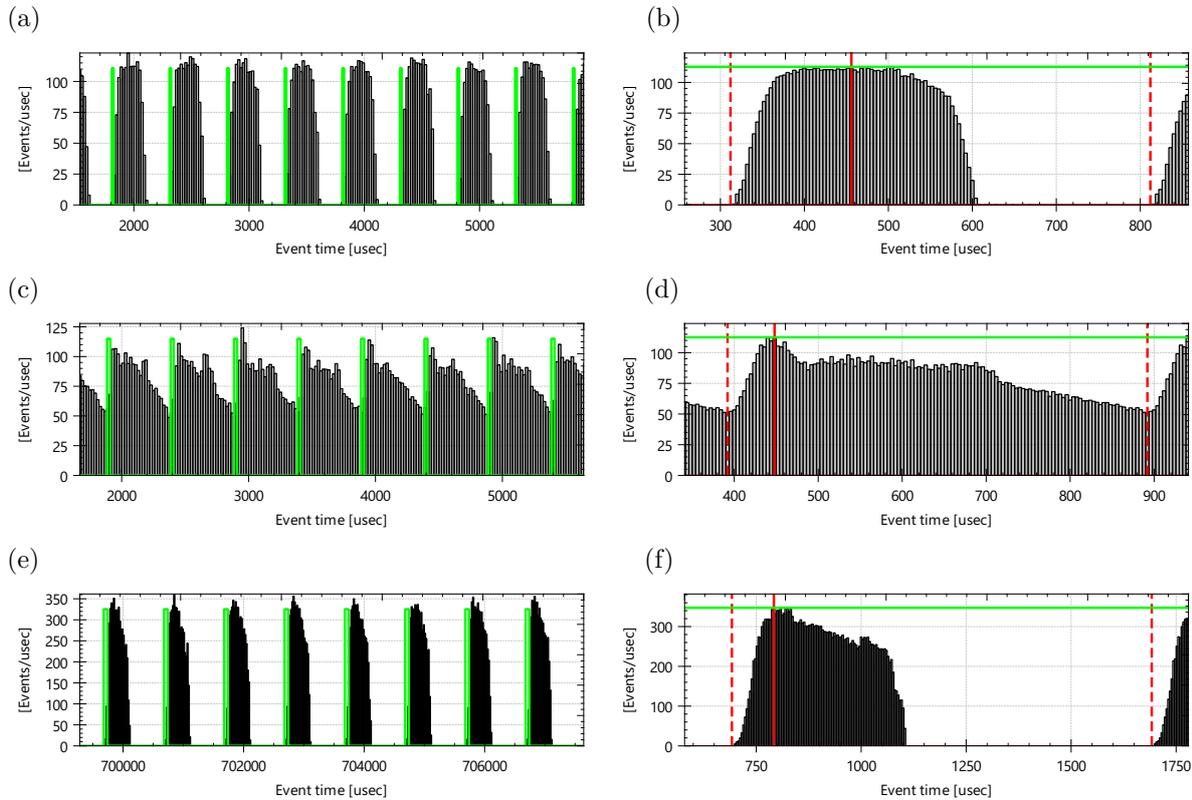


Figure 3: Histogram of event-data acquired from rotating disk setup with a light pulsing rate of 2 kHz at 54×10^6 Ev/s (a,b) and 81×10^6 Ev/s (c,d) and for 1 kHz at 105×10^6 Ev/s (e,f). Right column represents mean distribution over $\mathcal{O}(100)$ pulse periods, indicated by red dashed lines (20 μ s bin width).

3 Validation experiments

3.1 Rotating disk

The functionality of the RT-EBIV software was tested and debugged using a rotating disk setup that provides a simple and repeatable “flow” surrogate. The disk is coated with a scattering surface of small scatterers in the form of retro-reflecting glass beads of about 50 μ m embedded in clear acrylic. Maximum retro-reflective activity of the micro-beads is achieved by a LED ring light mounted on the camera objective, providing pulsed light illumination of desired frequency and intensity. Driven by a DC gear motor, the rotating disk provides stable long-term conditions. The setup also yields valuable insights into the maximum possible data rates that can be achieved with EBIV. For this, the setup was operated at different LED pulsing rates while observing the following constraints:

- no cross-talk (spillage) of events in the following frame (avoidance of situations as in Fig. 3c)
- minimal loss of triggered events due to readout limitations. These are noticeable as strip-like features in the pseudo-frames,
- absence of negative events in the data, which is achieved by de-tuning the EBV camera’s biases from their default values (here: High-pass filter=70, Diff-OFF=120).

As negative contrast change events are neither required nor suitable for EBIV, their avoidance provides more bandwidth for the capture and transmission of positive events. Event recordings of several seconds were acquired and post-processed to determine the maximum number of events that can be reliably captured at a given pulsing frequency.

Figure 3 provides histograms of the acquired event data at a LED pulsing rate of 2 kHz for LED pulse durations of 15 μ s and 25 μ s (corresponding respectively to 54 Ev/ μ s and 81 Ev/ μ s on average) and at a pulsing rate of 1 kHz for LED pulse durations of 60 μ s (corresponding to 105 Ev/ μ s). For the shorter pulse duration the pulses are separated cleanly from one pulse period to the next (see Fig. 3a,b). The

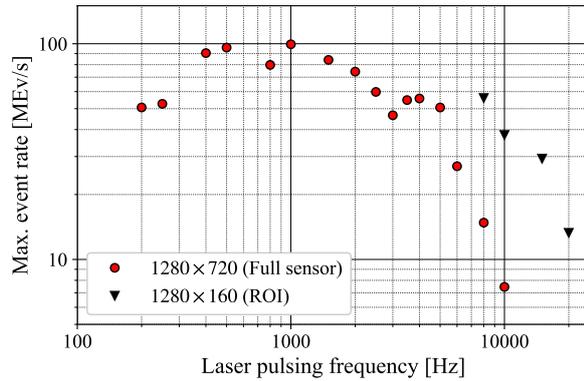


Figure 4: Maximum achievable event rate for different laser pulsing frequencies. Event rates for a reduced sensor ROI are indicated by the ∇ symbol.

longer pulse saturates the readout channel of the camera and results in event overflow from a given pulse period into the following (see Fig. 3c,d), at which point the performance of EBIV deteriorates considerably. The cleanly separated pulses at 54×10^6 Ev/s suggest that the event rate could be nearly doubled, however, at 81×10^6 Ev/s this is not the case (Fig. 3c,d). Nonetheless, pulsing the light at 1 kHz and 60 μ s duration results in a maximum event rate in excess of 100×10^6 Ev/s. Noteworthy is that the peak event rate exceeds 300×10^6 Ev/s at 1 kHz and decreases to about 120×10^6 Ev/s at 2 kHz. The mean data rate streamed from the camera is about 145 MB/s at 1 kHz and increases to 175 MB/s at 2 kHz which seems to indicate a sustained streaming limitation of about 150 MB/s. The discrepancy in the maximum achievable event rate at different pulsing frequencies is likely to be related to the event processing architecture on the camera and requires further investigation.

An overview of the maximum achieved event rate for different pulsing frequencies is given in Fig. 4 and indicates an optimum around 1 kHz reaching 100×10^6 Ev/s. At higher frequencies the finite latency begins to dominate, which leads to a rapid decrease in the number of particle images that can be captured reliably. It is presumed that at lower frequencies not all triggered pixels can be processed and are partially lost. The reduction of the sensor's active region of interest (ROI) in the vertical direction allows the maximum event rate to be increased, which is most likely due to a reduced number of rows being accessed during sensor readout. It should be noted that a ROI reduction in the horizontal direction does not lead to an increased event throughput.

In the context of particle imaging, the source density, generally expressed as particles per pixel (ppp), is an important performance metric for assessing algorithmic performance and achievable spatial resolution. Based on the maximum achievable event rates, the particle image count within the pseudo-frame can be predicted for different particle image diameters as plotted in Fig. 5(a). Normalization by the sensor array size yields the particle density in ppp (see Fig. 5b). Not surprisingly, the highest ppp is achieved at lower light pulsing frequencies: at 200 Hz a ppp approaching 0.3 is feasible assuming particle images of only 1 pixel in size. In practice, a particle image diameter of 1.5 to 2 pixel is preferred to reduce pixel locking artifacts, which reduces the maximum particle image density to the range of 0.1 ppp at 200 Hz. The number of particles captured per unit of time peaks at a pulsing frequency of 1 kHz and amounts to roughly 60 000 particle images of 1.5 px diameter per pseudo-frame which corresponds to about 0.06 ppp.

3.2 Validation on a free jet in water

More fluid-mechanically relevant conditions are produced with a small impeller-driven pump placed in a small water basin, as shown in Fig. 6 (a). The tank has a dimension of 300 mm length, 100 mm width and 200 mm height with the pump placed centrally at one end. The unconditioned flow exits from the pump nozzle (5.5 mm diameter) at about 1.5 mm, introducing large-scale convecting turbulent structures that can be tracked by the RT-EBIV system.

The water is seeded with spheroidal polyamide powder of 20 μ m median grain size (Arkema, Orgasol 202 D NAT 1). The domain of interest is illuminated by a light sheet of about 0.5 mm thickness generated with a low-cost pulse-width modulated diode laser (NEJE, E30130), which has a rated (CW) power of 5.5 W at a wavelength of 445 nm. For the measurements, a pulse width of 10 μ s to 20 μ s was sufficient to achieve optimal event rates from a light sheet of 40 mm height. The equivalent pulse energy is in the range of 100 μ J and reflects the high sensitivity of the EBIV sensor. The particle field is imaged at a magnification of 25.5 pixel/mm using a Nikon Micro Nikkor 55 mm macro lens with its aperture set to f/2.8, providing a field of view of 50.2 mm \times 28.2 mm. A representative snapshot of PIV processed jet

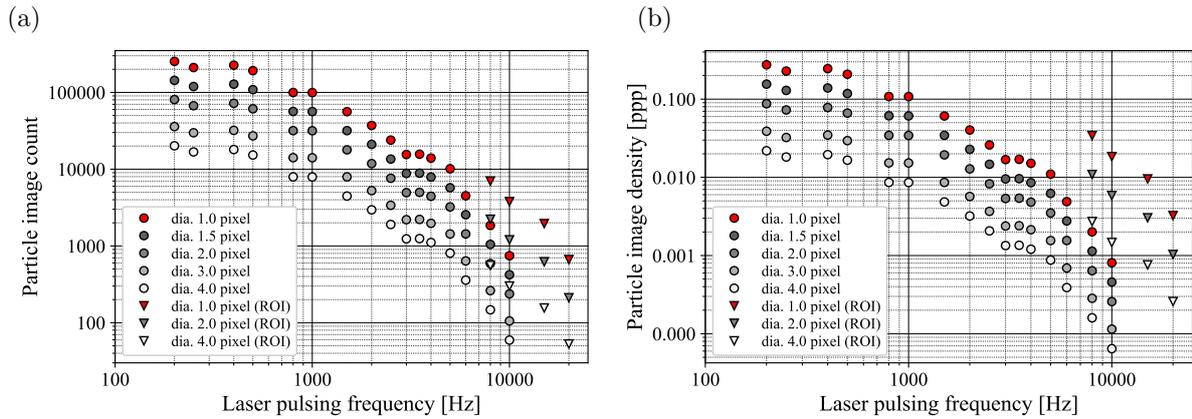


Figure 5: (a) maximum number of imaged particles at different laser pulsing frequencies and different particle image densities, (b) corresponding effective particle image density. ROI (∇) indicates an area of 1280×160 pixel.

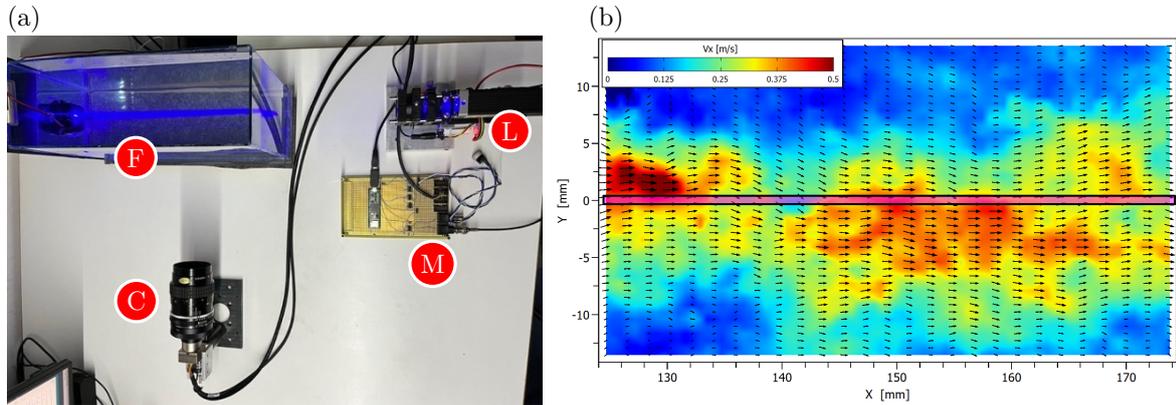


Figure 6: (a) Benchtop test setup of the implemented RT-EBIV configuration on a submerged jet water flow (F) using a Prophesee EVK4 camera (C) and pulse-modulated blue laser with light sheet optics (L). The camera and laser are synchronized with a micro-controller (M). (b) Snapshot of PIV-processed data captured by the setup. The red bar along the jet centerline indicates sampling ROI for space-time plots in Fig. 9.

flow data is provided in Fig. 6(b), whereas Fig. 7 shows typical pseudo-frames reconstructed from the event data stream at different event rates. The illumination is pulsed at 2 kHz, and each pseudo-frame captures five pulses for a total exposure time of 2.5 ms. Two operating conditions are shown: (a) event rate 25×10^6 Ev/s and (b) 75×10^6 Ev/s obtained by variation of the laser pulse energy.

Off-line processing Benchmarks of the water jet flow were performed using both live event streaming as well as off-line processing to assess the throughput in terms of velocity maps per second. Off-line processing was used to quantify the maximum possible throughput in the absence of the fixed pseudo-framing interval. The assessment is based on 2s recordings acquired at 1 kHz laser pulsing with a pulse width of 20 μ s resulting in a positive event data rate of 31×10^6 Ev/s. The corresponding 2000 pseudo-frames of the 2s sequences are continuously streamed by the RT-EBIV program on two different host platforms based on either X64-processor (Intel Xeon, 18 cores) or ARM64-processor (Apple M4 Pro, 10 cores). Since the graphical rendering of the data consumes considerable computational resources, display updates are reduced to $\mathcal{O}(10\text{Hz})$. Benchmarks in terms of velocity fields per second (fps) are obtained using multi-core processing with and without a validation step enabled. The assessment is repeated on 2×2 binned images, in which each 2×2 block of adjacent pixels is merged into one effective pixel. This spatial binning improves the signal-to-noise ratio and reduces computational load at the expense of

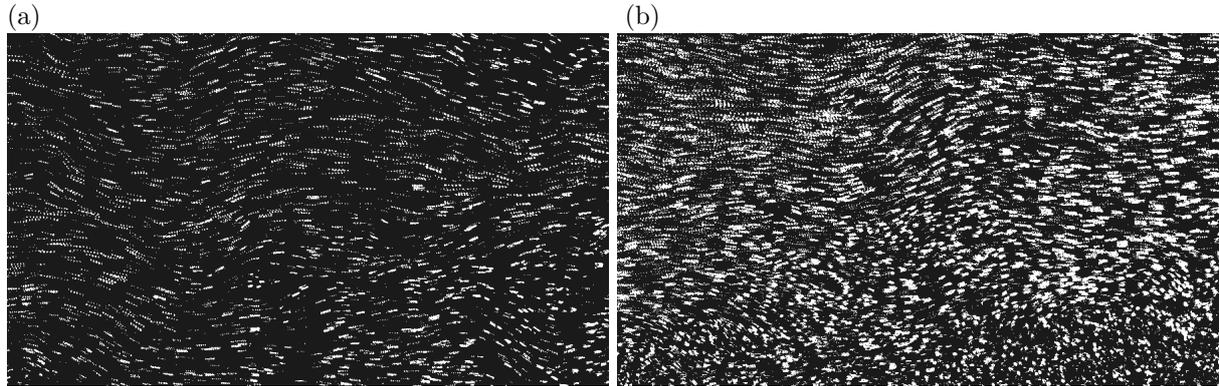


Figure 7: pseudo-frames obtained of the water jet flow at a laser pulsing frequency of 2 kHz capturing 5 pulses (2.5 ms) at 25×10^6 Ev/s (a) and 75×10^6 Ev/s (b).

Table 1: Performance of RT-EBIV in “free running mode” for different processing configurations using pre-recorded data using either an X64-processor (Intel Xeon W-2295 CPU, 18(36) cores, 3.00 GHz) or ARM64-processor (Apple M4 Pro, 10(14) cores). Frame pair rate (fp) is half the pseudo-framing rate; 2×2 binning is applied to both image and sampling window. Test case: turbulence in water, FOV 1280×720 pixel, $f_{\text{Laser}} = 1000$ Hz, 31×10^6 Ev/s. Values in bold indicate processing at or above pseudo-framing rate.

sample	nodes	valid.	throughput (x64)			throughput (arm64)		
			[fp/s]	[vec/s]	speed	[fp/s]	[vec/s]	speed
multi-threaded performance								
32×32	880	on	328	288 640	0.16×	816	718 080	1.63 ×
32×32	880	off	361	317 680	0.72×	907	798 160	1.81 ×
32×32 (binned)	880	on	524	461 120	1.05 ×	1425	1 254 000	2.85 ×
32×32 (binned)	880	off	610	536 800	1.22 ×	1750	1 540 000	3.50 ×
single threaded performance								
32×32	880	on	86	75 680	0.66×	182	160 160	0.36×
32×32	880	off	87	76 560	0.17×	187	164 560	0.37×
32×32 (binned)	880	on	295	259 600	0.59×	623	548 240	1.25 ×
32×32 (binned)	880	off	311	273 680	0.62×	673	592 240	1.35 ×

spatial resolution.

Table 1 provides an overview of the velocity field throughput that can be achieved using pre-recorded event data. For the 1 kHz framing rate used here, true real-time processing can be achieved when the *speed* quantity in Table 1 exceeds unity, which is the case for the ARM64 processor whereas the X64 processor can only achieve this for 2×2 binned pseudo-images. The single-threaded validation step consumes roughly 10% to 20% of the processing resources. Image binning increases processing speed by about 80% which is mainly due to the $4 \times$ smaller PIV samples. The estimates reported in Table 1 have an uncertainty of about 10% and are strongly influenced by background activities of the multi-process host operating system (MS Window 11 Pro).

On-line processing Actual RT-EBIV measurements are performed on the same water jet setup using a pseudo-framing rate of 2 kHz and a fixed event rate of 33×10^6 Ev/s. For the assessment, 5000 velocity maps are sampled with different processing configurations. Because frame generation occupies dedicated CPU threads, the current frame pair cannot be transferred to the PIV processing thread immediately, resulting in a short hand-off delay. This delay varies over time depending on how well the independent threads align with each other and can be retrieved from the actual frame counts. Figure 8 shows histograms of the frame-pair interval, defined as the number of pseudo-frames that elapse between two

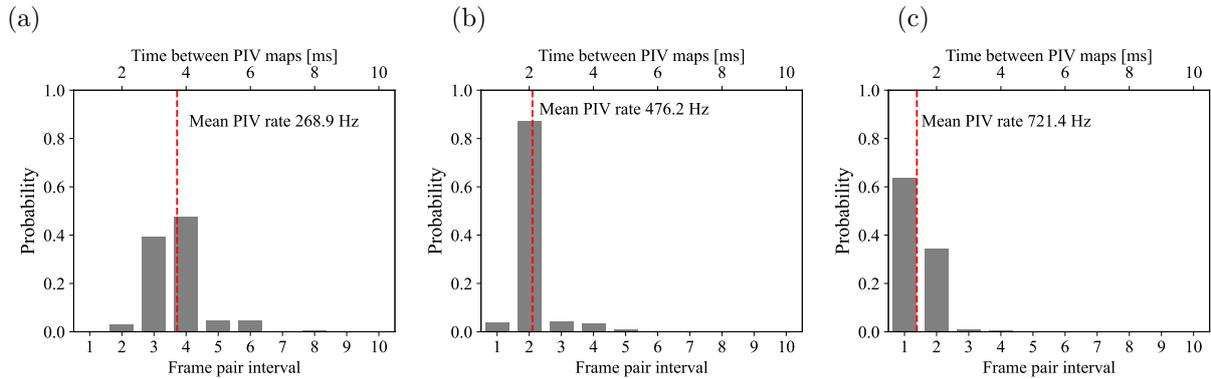


Figure 8: Histograms of the frames pair intervals while online processing live event data stream at 33×10^6 Ev/s at a pseudo-framing rate of 2000 Hz; (a): full field processing (880 nodes), (b): ROI (395 nodes), (c): binned pseudo-images processed with 24×24 sampling (390 nodes). Mean PIV-field rate is indicated by dashed red line.

consecutive PIV snapshots. Because snapshots are computed while new pseudo-frames arrive at a fixed rate, multiple pseudo-frames often elapse before the next snapshot is available. The histogram quantifies the effective processing cadence and its variability.

At full image resolution (1280×720 pixel), velocity fields containing $40 \times 21 = 880$ nodes and 32×32 pixel sampling can be provided at nearly 270 Hz using 8 cores of an Intel i7 processor. The histogram (Fig. 8a) reveals that most fields are obtained at frame-pair intervals of 3 or 4, while a few can be sampled at shorter intervals. The longer intervals are caused by occasional unsteadiness in the event stream leading to drop-out. By reducing the field of view, the number of nodes is reduced, which is reflected in an increased velocity field rate of nearly 480 Hz of $79 \times 5 = 395$ nodes (Fig. 8b). Image-binning further increases the processing speed: fields with $26 \times 15 = 390$ nodes are processed at a continuous rate of 720 Hz (Fig. 8c). In this case, nearly two-thirds of the generated frame pairs are processed without skipping.

Figure 9 shows the spatio-temporal evolution of the velocity components along the centerline of the jet (x -axis), also known as kymogram. They are obtained by stacking the velocity profile data from each processed frame-pair, indicated by the horizontal strip in Fig. 6(b), and provide an intuitive visualization of spatio-temporal flow dynamics and data continuity. As shown previously in Fig. 8(c), the fraction of dropped frame pairs is reduced to about 30%, resulting in an almost continuous time series with a mean sampling rate of 720 Hz. This improvement is achieved by pixel-binning the data and by reducing the number of nodes to 26×15 . In binned mode, the sampling window of 24×24 pixels corresponds to 48×48 pixel sample in non-binned space, but is processed nearly twice as fast. The larger sample size enhances the robustness of the cross-correlation processing and increases the data yield, as reflected in a high validation rate in Fig. 9(c).

4 EBIV in active control configurations

This section presents two examples of active flow control applications of EBIV, highlighting its use in both open-loop and closed-loop configurations. In the first case, EBIV provides a fast and data-efficient cost-function estimator for the optimization of mixing in an air jet, where actuation parameters are tuned through Bayesian optimization. In the second case, EBIV supplies real-time velocity fields for feedback control of a jet flow, where actuation is adapted through a Reinforcement Learning (RL) algorithm. Both studies are ongoing, and only a brief description of the experimental setups together with preliminary results is provided. A schematic overview of the two configurations is given in Fig. 10.

4.1 Air-jet open-loop optimization

The first application explores the use of EBIV within an open-loop optimization framework aimed at maximizing the turbulent kinetic energy (TKE) of a jet flow, which serves as a quantitative measure of its mixing efficiency. A detailed description of this study is provided in Franceschelli et al. [26], while only a concise overview is given here. The overall concept is illustrated in Fig. 10(a). In this configuration, EBIV provides spatially resolved velocity fields from which we compute a reward proportional to the

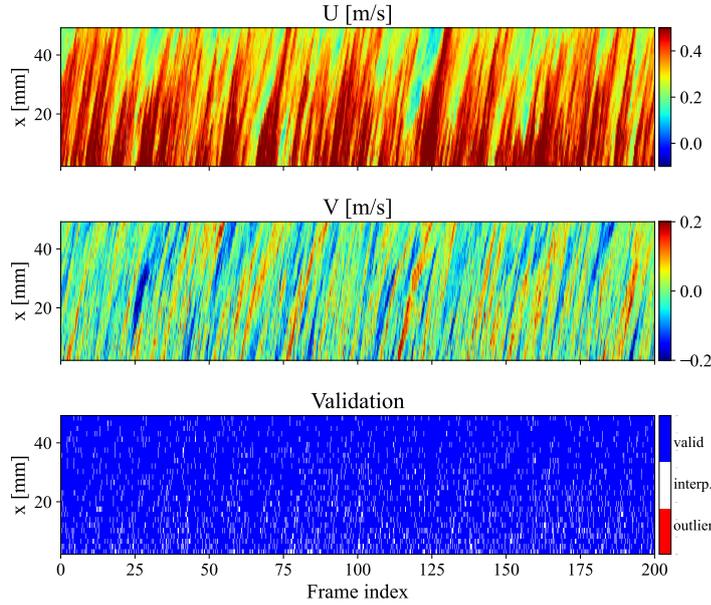


Figure 9: Kymograms of both velocity components obtained from 2×2 binned pseudo-images at a pseudo-framing rate of 2000 Hz. Data is sampled along the jet axis using 24×24 pixel sampling (26 nodes). Mean flow field sample rate is 721 Hz.

spatially integrated TKE over the observed region, used as a proxy for mixing. A Bayesian optimization algorithm then iteratively adjusts the actuation parameters to maximize this reward.

Unlike previous studies that relied on point-wise measurements to define the reward [27, 28], the use of EBIV allows the reward to be formulated from spatially distributed flow information, providing a more representative measure of the jet dynamics. The reduced data stream and high sensitivity of EBIV further make it possible to explore complex actuation strategies rapidly and with relatively low-cost instrumentation (illumination is provided by a 20 W diode laser originally designed for laser engraving applications).

The experiment is conducted on a round air jet at $Re_D \approx 2500$, with the laser sheet placed in the jet's longitudinal mid-plane and EBIV operated at 1.5 kHz. Actuation is provided by six zero-net-mass-flow (ZNMF) synthetic jet (SJ) arranged around the nozzle exit and by two acoustic drivers in the stagnation chamber. The SJ were independently characterized, and their driving amplitude was adjusted so that the peak velocity of the pulsating jet at one diameter downstream of each exit orifice was 1.5 m/s. The optimization variables are the actuation frequencies f_i and their relative phase delays ϕ_i , expressed in non-dimensional form as Strouhal numbers $St_i = f_i D/U$. Bayesian optimization is used to maximize a TKE-based reward, $y(\mathbf{x})$, where $\mathbf{x} = (f_i, \phi_i)$ collects the actuation frequencies and phase delays. The reward is the spatial average of the ensemble-averaged TKE over a masked subset Ω of the EBIV-resolved field, excluding the low-velocity regions:

$$y(\mathbf{x}) = \frac{1}{|\Omega|} \int_{\Omega} \overline{\text{TKE}}(x, y) \, d\Omega, \quad \Omega = \{(x, y) \mid \bar{u}(x, y) \geq 0.2 \bar{u}_{\max}\}. \quad (1)$$

Here, $\overline{\text{TKE}}(x, y) = \frac{1}{2}(\overline{u'^2}(x, y) + \overline{v'^2}(x, y))$, and the overbar denotes ensemble averaging of EBIV snapshots. The reference velocity \bar{u}_{\max} is the maximum mean streamwise velocity within the EBIV-resolved field, i.e., $\bar{u}_{\max} = \max(\bar{u}(x, y))$. The $0.2 \bar{u}_{\max}$ threshold filters out quiescent regions where TKE may be artificially elevated.

Each evaluation is based on an average of 500 randomly sampled EBIV pseudo-frames over a five-second acquisition. To improve robustness, every actuation setting is repeated twice, and a third measurement is performed if the results fall outside a $\pm 3.5\%$ confidence interval. The evaluation of a single actuation setting requires approximately 1.5–2 minutes, including acquisition and processing.

Although the setup consists of six actuators, the results shown here focus on a simplified case that serves as a preliminary step before extending the optimization to the full system. Specifically, only the two actuators located in the EBIV measurement plane (actuators 3 and 6 in Fig. 10(a)) are considered. The optimization is performed with respect to their actuation frequencies, f_3 and f_6 , and relative phase delay, $\Delta\phi$. Since the measurements are centered on the jet axis, the actuation space is symmetric:

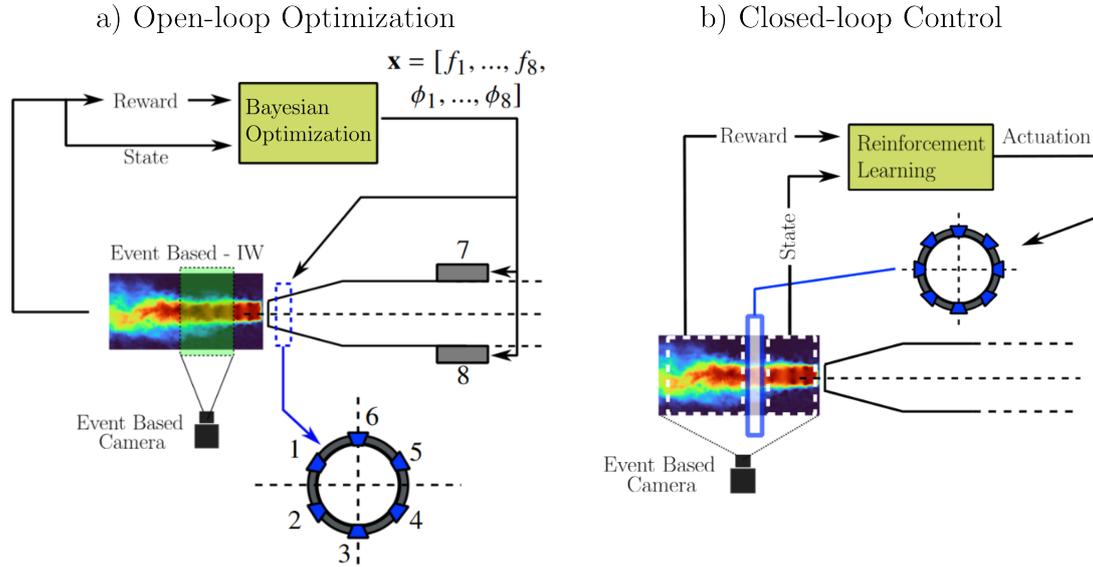


Figure 10: Schematic representation of two RT-EBIV-based active flow control applications: (a) open-loop optimization of an air jet and (b) closed-loop control of a jet flow.

$\mathbf{x} = (f_3, f_6, \Delta\phi)$ is equivalent to $(f_6, f_3, -\Delta\phi)$ and it yields the same y . The optimization map is therefore symmetric with respect to the bisector, and the search can be restricted to one half of the (f_3, f_6) plane.

A clear region of enhanced performance appears around $St \approx 0.3$, corresponding to the preferred axisymmetric mode of the jet and consistent with its natural dynamics [29, 30]. Interestingly, the optimal configuration is obtained when one SJ is driven at this frequency while the other operates at approximately twice the value, forming a subharmonic–harmonic pairing that maximizes the TKE and thus the mixing efficiency. The convergence history of the Bayesian optimization is shown in Fig. 11(b). The cost function $y(\mathbf{x})/U_\infty^2$ is plotted over 400 iterations i . The red line indicates the historical maximum, i.e., the best cost function value identified up to the i -th iteration, while the yellow shaded region corresponds to the initial 50 exploratory actuations used to initialize the optimization. After this exploration stage, a consistent increase in the cost function is observed, confirming the robustness of the optimization process. Under the optimal actuation strategy, the controlled jet exhibits a marked enhancement of TKE already in the near field ($x/D < 3$) and develops an asymmetric flow structure consistent with the asymmetric forcing. Importantly, this TKE increase does not arise from indiscriminate amplification of small-scale turbulence; rather, the actuation promotes large-scale, spatially coherent structures.

Actuation power and efficiency. The rms electrical power of each actuator is estimated as $P_{\text{rms}} = V_{\text{rms}}^2/R$, where V_{rms} is the rms drive voltage to the loudspeaker and $R = 3.6 \Omega$ is the circuit resistance. For the optimal setting, the two actuators operating at $St = (0.34, 0.70)$ consume $P_{\text{rms}} = (0.26, 0.78) \text{ W}$, corresponding to a total electrical power of 1.04 W.

The fluidic power supplied by each SJ to the flow can be estimated as

$$P = \frac{1}{2} \rho A u_u^3, \quad (2)$$

where ρ is the air density, A is the actuator exit area and u_u is a characteristic velocity [31]. There is no single, widely accepted definition of u_u in the literature. In this work, we set $u_u = U_0$, following the definition proposed by Greco et al. [32] but adapted to our measurement configuration. The reference velocity U_0 is obtained from the phase-averaged centreline velocity of the synthetic jet during the ejection half-cycle as

$$U_0 = \frac{1}{\tau} \int_0^{\tau/2} u_c(t) dt, \quad (3)$$

where $\tau = 1/f$ is the actuation period and $u_c(t)$ is the hot-wire velocity measured on the SJ centreline at one SJ diameter downstream of the exit. This location is assumed to provide a good approximation

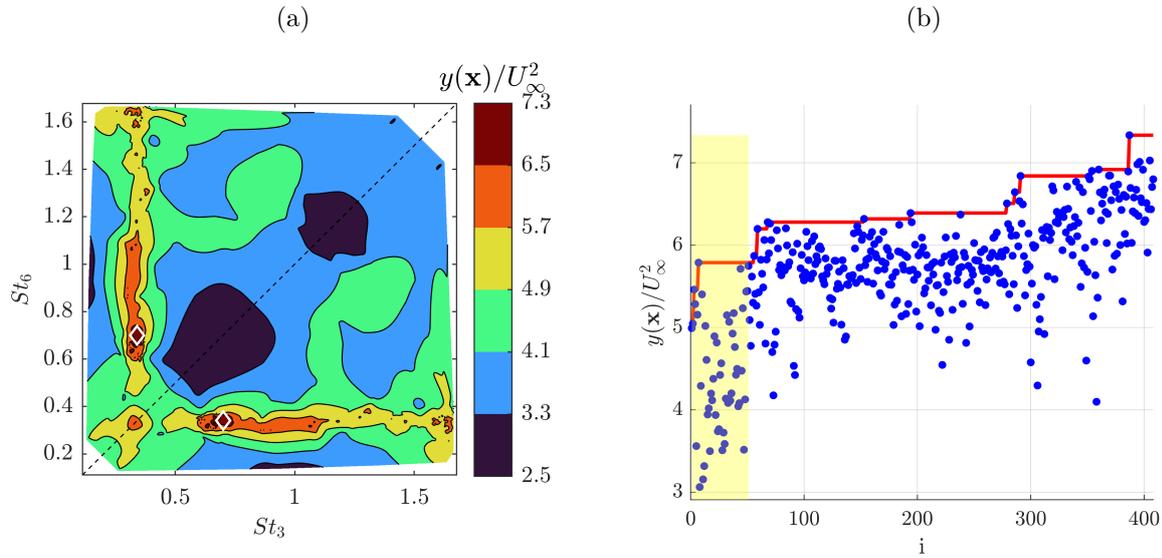


Figure 11: Open-loop optimization results in the air jet experiment: (a) frequency optimization map for the two SJ actuators, expressed in terms of Strouhal numbers. Because the two actuators are arranged symmetrically about the jet axis, the cost function is invariant under exchanging their labels and reversing the phase; the map is therefore symmetric by construction with respect to the diagonal $St_3 = St_6$. (b) Cost-function history over 400 iterations of the Bayesian optimization process, normalized by the squared value of the jet bulk velocity U_∞^2 . The red line represents the best cost-function value identified up to the i -th iteration. The yellow shaded region corresponds to the initial 50 exploratory actuations used to initialize the Bayesian optimization algorithm.

of the exit velocity, since the mass flow rate in the ejection phase is nearly conserved up to this position and entrainment effects remain small [32]. Such definition, which relies solely on the centreline value and implicitly assumes a uniform exit profile, tends to overestimate the actual kinetic energy flux incident on area A and, therefore, can be considered as a conservative (lower-bound) estimate of the control efficiency. In the present experiments, the peak velocity at the actuator exit is kept approximately constant for all actuation frequencies. Although this does not strictly enforce a constant $U_{0,c}$, if the temporal waveform of the jet does not change significantly with f , $U_{0,c}$ remains approximately constant and so does the estimated fluidic power input, $P \approx 1.2 \times 10^{-6}$ W, across all actuation frequencies. This actuation produces an average increase of the turbulent kinetic energy in the observed region by a factor of approximately 15.8. More specifically, the mean TKE in the observed region increases from about $38.7 \text{ m}^2/\text{s}^2$ in the uncontrolled case to approximately $610.0 \text{ m}^2/\text{s}^2$ under optimal actuation.

Although the electrical power input to the actuators can be measured accurately, and the fluidic power of each synthetic jet can be estimated from its characterization, it is still not straightforward to define a meaningful control efficiency: only a limited portion of the jet is observed, so the associated increase in flow energy (and power) cannot be reliably integrated over the full three-dimensional field. For this reason, we do not introduce a global efficiency metric and restrict the discussion to local changes in TKE and actuation power.

4.2 RT-EBIV in closed-loop control

The second application investigates the use of RT-EBIV for closed-loop flow control of a submerged water jet with a nozzle diameter of $D = 30$ mm operating at $Re_D \approx 9000$. The overall control concept is illustrated in Fig. 10(b). In this configuration, RT-EBIV serves two purposes. First, it provides real-time estimates of the flow state in the region immediately downstream of the nozzle, within the first two diameters. Second, it evaluates the reward further downstream, beyond the actuation region, thus enabling a direct link between the actuation and its impact on the jet development. Actuation is provided by eight radially oriented water valves positioned at approximately $2D$ from the nozzle exit.

The velocity fields are processed online using the RT-EBIV software at an effective rate of about 100 Hz, which is sufficient for feedback-based experiments in this setup. Control is performed using a RL

approach [33], specifically the Soft Actor-Critic (SAC) algorithm, chosen for its robustness to measurement noise and off-policy data [34]. We deliberately selected a simple jet-deflection reward as a sanity-check objective to verify the end-to-end effectiveness of RT-EBIV-in-the-loop control before addressing more complex goals.

An illustrative case is presented in Fig. 12, where the control objective was to deflect the mean jet upward. The left frame shows a velocity field from the initial training episode, with no clear deflection; after approximately 110 episodes (right frame), the controller induces a consistent lateral displacement. The red markers indicate the actuator actions in the EBIV measurement plane. In parallel, Fig. 13 reports the episode reward, $r = \frac{1}{|\Omega_r|} \int_{\Omega_r} v(x, y) d\Omega_r$, defined as the spatial mean of the vertical velocity in a fixed window downstream of the actuation region, normalized by the maximum obtained reward r_{max} . The reward fluctuates around zero during early exploration and rises sharply after ~ 110 episodes, matching the onset of sustained deflection. These results confirm the feasibility of using EBIV-based real-time sensing in reinforcement-learning-driven closed-loop flow control.

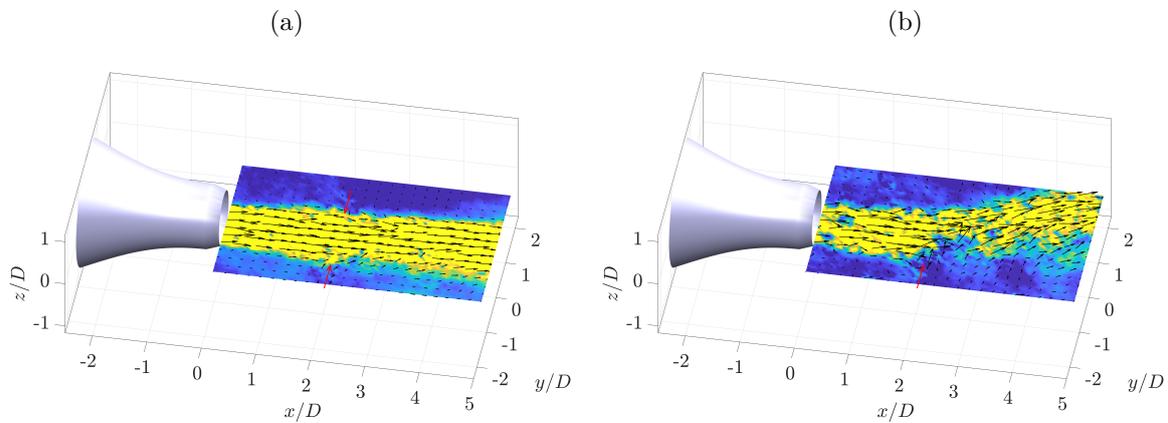


Figure 12: Closed-loop control of a submerged water jet ($D = 30$ mm, $Re_D \approx 9000$) using RT-EBIV. (a) Velocity field from the first training episode, showing no clear deflection. (b) Velocity field after approximately 110 training episodes, demonstrating successful deflection of the jet in the positive y -direction. Red markers indicate actuator actions in the EBIV measurement plane.

5 Discussion and Concluding Remarks

The present work has demonstrated that RT-EBIV can reliably deliver velocity fields at rates of several hundred Hertz, both in off-line and real-time operation. In particular, online measurements of a submerged jet flow confirmed sustained throughput rates in excess of 250 Hz, depending primarily on the pseudo-framing rate (set by the pulsing frequency of the light source) and on the density of the sampled velocity map. For pulsing frequencies up to 500 Hz, frame pairs were processed without skipping, while at higher frequencies some frames were dropped due to processor limitations. Importantly, the system

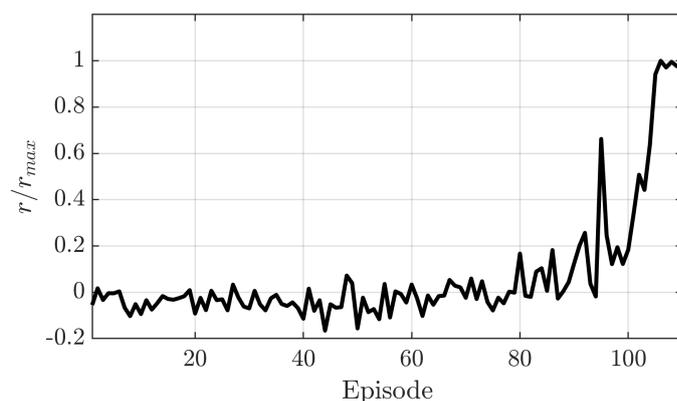


Figure 13: Training reward history in the closed-loop water jet experiment. Episode reward r , normalized by its maximum r_{max} , defined as the spatial average of $v(x, y)$ over a downstream window Ω_r .

latency—defined as the time between pseudo-frame capture and availability of the corresponding velocity field—remains essentially unaffected, since the second pseudo-frame is available immediately after the slicing interval. This constitutes a key advantage over framing-based cameras, which require an additional read-out step.

The achievable performance was found to depend strongly on the host processor architecture. ARM64 platforms showed a significant gain in processing speed compared to conventional X64 systems, although EBV camera drivers for such platforms are not yet available. Background operating-system activity was also observed to affect throughput, underlining the relevance of lightweight and dedicated processing environments.

Beyond these feasibility studies, this work has explored two first applications of EBIV in active flow control. In an open-loop configuration, EBIV served as a fast and data-efficient cost-function estimator, enabling Bayesian optimization of SJ actuation for enhanced mixing in an air jet. In a closed-loop configuration, RT-EBIV supplied real-time velocity fields at 100 Hz, providing both flow-state estimation and reward evaluation for a submerged water jet controlled with RL. Although preliminary, both examples highlight EBIV's dual potential: accelerating the exploration of actuation strategies and enabling real-time feedback control.

Looking forward, further improvements in performance are expected from dedicated hardware implementations. A promising pathway is the hardware-level generation of one-bit pseudo-frames directly on FPGA boards attached to the EBV sensor, which could eliminate transfer bottlenecks and further reduce latency [35]. Along this line Gao et al. [36] reported throughput rates exceeding 1000 fps. Evaluation kits coupling EBV sensors with FPGA (e.g. AMD/Xilinx Kria KV260 with Sony/Prophesee IMX636) are already commercially available and could provide a platform for compact, low-latency flow control systems. Combined with high-throughput GPU-based algorithms [11], such developments have the potential to extend RT-EBIV from laboratory demonstrations to robust implementation in practical active flow control applications.

Acknowledgments

SD, MR and LF acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 949085, NEXTFLOW ERC StG). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

- [1] Mohamad Gad-El-Hak. Flow control. *Applied Mechanics Reviews*, 42(10):261, 1989.
- [2] Steven L. Brunton and Bernd R. Noack. Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Reviews*, 67(5):050801, 08 2015. ISSN 0003-6900. doi: 10.1115/1.4031175.
- [3] M. Raffel, C.E. Willert, C.J. Kähler, F. Scarano, S.T. Wereley, and J. Kompenhans. *Particle Image Velocimetry: A Practical Guide (3rd Edition)*. Springer-Verlag, Berlin Heidelberg, 2018. ISBN 978-3-319-68852-7. doi: 10.1007/978-3-319-68852-7.
- [4] Stefan Siegel, Kelly Cohen, Tom McLaughlin, and James Myatt. Real-time particle image velocimetry for closed-loop flow control studies. In *41st Aerospace Sciences Meeting and Exhibit, Reno (NV)*, 2003. doi: 10.2514/6.2003-920. AIAA 2003-920.
- [5] C. Willert, M. Munson, and M. Gharib. Real-time particle image velocimetry for closed-loop flow control applications. In *15th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, 5-8 July 2010.*, 2010. URL <https://elib.dlr.de/64688/>.
- [6] Elliott Varon, Jean-Luc Aider, Yoann Eulalie, Stephie Edwige, and Philippe Gilotte. Adaptive control of the dynamics of a fully turbulent bimodal wake using real-time PIV. *Experiments in Fluids*, 60: 124, 2019. doi: 10.1007/s00348-019-2766-6.
- [7] Z. Ouyang, H. Yang, J. Lu, and ZhouPing Yin. Real-time adaptive particle image velocimetry for accurate unsteady flow field measurements. *Science China Technological Sciences*, 65:2143–2155, 2022. doi: 10.1007/s11431-022-2082-4.

- [8] Findlay McCormick, Bradley Gibeau, and Sina Ghaemi. Reactive control of velocity fluctuations using an active deformable surface and real-time PIV. *Journal of Fluid Mechanics*, 985:A9, 2024. doi: 10.1017/jfm.2024.292.
- [9] N. Gautier and J.L. Aider. Real-time planar flow velocity measurements using an optical flow algorithm implemented on GPU. *Journal of Visualization*, 18:277–286, 2015. doi: 10.1007/s12650-014-0222-5.
- [10] N. Gautier, J.-L. Aider, T. Duriez, B. R. Noack, M. Segond, and M. Abel. Closed-loop separation control using machine learning. *Journal of Fluid Mechanics*, 770:442–457, 2015. doi: 10.1017/jfm.2015.95.
- [11] Scott A. Bollt, Samuel H. Foxman, and Morteza Gharib. RapidPIV: Full flow-field khz piv for real-time display and control. *arXiv preprint arXiv:2504.17987*, 2025. doi: 10.48550/arXiv.2504.17987.
- [12] Caroline Braud and Alex Liberzon. Real-time processing methods to characterize streamwise vortices. *Journal of Wind Engineering and Industrial Aerodynamics*, 179:14–25, 2018. doi: 10.1016/j.jweia.2018.05.006.
- [13] Naoki Kanda, Chihaya Abe, Shintaro Goto, Keigo Yamada, Kumi Nakai, Yuji Saito, Keisuke Asai, and Taku Nonomura. Proof-of-concept study of sparse processing particle image velocimetry for real time flow observation. *Experiments in Fluids*, 63(9):143, 2022. doi: 10.1007/s00348-025-04039-4.
- [14] Taku Nonomura, Chihaya Abe, Ryo Naramura, and Yasuo Sasaki. Real-time feedback control of flow velocity field using sparse processing particle image velocimetry and plasma actuators. *Experiments in Fluids*, 66(7):136, 2025. doi: 10.1007/s00348-025-04039-4.
- [15] Guillermo Gallego, Tobi Delbrück, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J Davison, Jörg Conrath, Kostas Daniilidis, et al. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, 2020. doi: 10.1109/TPAMI.2020.3008413.
- [16] Alexander Rusch and Thomas Rösgen. TrackAER: real-time event-based quantitative flow visualization. *Experiments in Fluids*, 64:136, 2023. doi: 10.1007/s00348-023-03673-0.
- [17] Zhenjiang Ni, Cécile Pacoret, Ryad Benosman, Siohoi Ieng, and Stéphane Régnier. Asynchronous event-based high speed vision for microparticle tracking. *Journal of Microscopy*, 245(3):236–244, 2012. doi: 10.1111/j.1365-2818.2011.03565.x.
- [18] Yuanhao Wang, Ramzi Idoughi, and Wolfgang Heidrich. Stereo event-based particle tracking velocimetry for 3d fluid flow reconstruction. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020*, pages 36–53, Cham, 2020. Springer International Publishing. ISBN 978-3-030-58526-6. doi: 10.1007/978-3-030-58526-6_3.
- [19] Christian E Willert and Joachim Klinner. Dynamic wall shear stress measurement using event-based 3d particle tracking. *Experiments in Fluids*, 66:32, 2025. doi: 10.1007/s00348-024-03946-2.
- [20] Christian Willert. Event-based imaging velocimetry using pulsed illumination. *Experiments in Fluids*, 64:98, 2023. doi: 10.1007/s00348-023-03641-8.
- [21] Luca Franceschelli, Christian E. Willert, Marco Raiola, and Stefano Discetti. An assessment of event-based imaging velocimetry for efficient estimation of low-dimensional coordinates in turbulent flows. *Experimental Thermal and Fluid Science*, 164:111425, 2025. ISSN 0894-1777. doi: 10.1016/j.exptthermflusci.2025.111425.
- [22] Jiajun Cao, Xin Zeng, Sen Li, Chuangxin He, Xin Wen, and Yingzheng Liu. A novel event-based ensemble particle tracking velocimetry for single-pixel turbulence statistics. *Experimental Thermal and Fluid Science*, page 111554, 2025. doi: 10.1016/j.exptthermflusci.2025.111554.
- [23] Rohit Chandra, Leo Dagum, David Kohr, Ramesh Menon, Dror Maydan, and Jeff McDonald. *Parallel programming in OpenMP*. Morgan Kaufmann, 2001.
- [24] Jerry Westerweel and Fulvio Scarano. Universal outlier detection for piv data. *Experiments in Fluids*, 39(6):1096–1100, 2005. ISSN 0723-4864. doi: 10.1007/s00348-005-0016-6.

- [25] Thomas Finateu, Atsumi Niwa, Daniel Matolin, Koya Tsuchimoto, Andrea Mascheroni, Etienne Reynaud, Pooria Mostafalu, Frederick Brady, Ludovic Chotard, Florian LeGoff, Hirotsugu Takahashi, Hayato Wakabayashi, Yusuke Oike, and Christoph Posch. 5.10 - A 1280×720 back-illuminated stacked temporal contrast event-based vision sensor with 4.86 μm pixels, 1.066GEPS readout, programmable event-rate controller and compressive data-formatting pipeline. In *2020 IEEE International Solid-State Circuits Conference - (ISSCC)*, pages 112–114, 2020. doi: 10.1109/ISSCC19947.2020.9063149.
- [26] Luca Franceschelli, Enrico Amico, Marco Raiola, Christian Willert, Jacopo Serpieri, Gioacchino Cafiero, and Stefano Discetti. Event-based imaging velocimetry for jet flow control. In *16th International Symposium on Particle Image Velocimetry (ISPIV 2025)*, Tokyo, Japan, June 2025. URL <https://elib.dlr.de/215871/>.
- [27] Johann Moritz Reumschüssel, Yiqing Li, Philipp Maximilian zur Nedden, Tianyu Wang, Bernd R Noack, and Christian Oliver Paschereit. Experimental jet control with bayesian optimization and persistent data topology. *Physics of Fluids*, 36(9), 2024. doi: 10.1063/5.0217519.
- [28] Tamir Shaqarin, Zhutao Jiang, Tianyu Wang, Chang Hou, Guy Y Cornejo Maceda, Nan Deng, Nan Gao, and Bernd R Noack. Jet mixing optimization using a bio-inspired evolution of hardware and control. *Scientific Reports*, 14(1):25952, 2024. doi: 10.1038/s41598-024-75688-4.
- [29] A.K.M. Fazle Hussain and K.B.M.Q. Zaman. The ‘preferred mode’ of the axisymmetric jet. *Journal of Fluid Mechanics*, 110:39–71, 1981. doi: 10.1017/S0022112081000608.
- [30] Ephraim Gutmark and Chih-Ming Ho. Preferred modes and the spreading rates of jets. *Physics of Fluids*, 26(10):2932–2938, 1983. doi: 10.1063/1.864058.
- [31] Paweł Gil and Piotr Strzelczyk. Performance and efficiency of loudspeaker driven synthetic jet actuator. *Experimental Thermal and Fluid Science*, 76:163–174, 2016. doi: 10.1016/j.expthermflusci.2016.03.020.
- [32] Carlo Salvatore Greco, Andrea Ianiro, Tommaso Astarita, and Gennaro Cardone. On the near field of single and twin circular synthetic air jets. *International journal of heat and fluid flow*, 44:41–52, 2013.
- [33] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996. doi: 10.1613/jair.301.
- [34] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, pages 1861–1870. PMLR, 2018.
- [35] Krzysztof Blachut and Tomasz Kryjak. High-definition event frame generation using SoC FPGA devices. In *2023 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*, page 106–111. IEEE, September 2023. doi: 10.23919/spa59660.2023.10274447.
- [36] Yizhao Gao, Baoheng Zhang, Yuhao Ding, and Hayden Kwok-Hay So. A composable dynamic sparse dataflow architecture for efficient event-based vision processing on FPGA. In *Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays, FPGA ’24*, page 246–257, New York, NY, USA, 2024. Association for Computing Machinery. ISBN 9798400704185. doi: 10.1145/3626202.3637558.