

Real-time Particle Image Velocimetry using Event-based Imaging

C. Willert¹, L. Franceschelli²

¹ DLR Institute of Propulsion Technology, Deutsches Zentrum für Luft- und Raumfahrt (German Aerospace Center), Köln, Germany

² Department of Aerospace Engineering, Universidad Carlos III de Madrid, Leganés, Madrid, Spain *
chris.willert@dlr.de

Abstract

The work addresses efforts in harnessing the potential of event-based vision (EBV) for the rapid retrieval of flow field data in configurations akin to two-dimensional, two-component particle image velocimetry (2d-2c PIV). The reduced data stream of EBV compared to conventional frame-based imaging allows for equivalent framing rates up into the kHz range and thereby can capture transient phenomena. The discussed implementation utilizes lower power cameras that directly stream to conventional off-the-shelf computing hardware and can achieve velocity field rates in the kHz range at a throughput of $O(10^6)$ vectors per second. The concept is demonstrated on a small jet in a water basin by sampling the flow at pseudo-frame rates of 2 kHz and seeding densities up to 0.04 particles per pixel (ppp). By varying laser pulsing rate, seeding density and laser energy a maximum achievable event data rate exceeding 100×10^6 events/s could be demonstrated at a pseudo-frame rate of 1 kHz. Online measurements continuously provided more than 250 velocity maps per second, each with 880 nodes, increasing to more than 700 Hz using pixel-binning and smaller sampling windows.

1 Introduction

With its reduced data rate and potentially rapid data processing capability, EBV is a strong candidate for use in active flow control scenarios, such as depicted in Fig. 1 for particle image velocimetry (PIV). For the most part, current implementation of EBV based flow diagnostics operate in an off-line mode, that is, event data is first collected from the experiment and then processed using algorithms in which computational cost is of lesser importance. The event-based three-dimensional Lagrangian particle tracking (3D-LPT) implementations by (Borer et al., 2017) and (Rusch and Rösger, 2023) can be considered an exception, since they reconstruct the (time-averaged) 3D flow field in real-time. However, only a very limited number of particles can be tracked simultaneously. On the background of leveraging real-time flow field data from event-camera data streams, the number density is on the low side and necessitates different approaches.

Within the field of PIV, the literature reports several 2D-2C real-time particle image velocimetry (RT-PIV) implementations based on conventional cross-correlation of image pairs (Siegel et al., 2003; Willert et al., 2010; Varon et al., 2019; Ouyang et al., 2022; McCormick et al., 2024) or by optical flow methods (Gautier and Aider, 2015; Gautier et al., 2015). Processing rates are coupled to the camera framing rate with image pairs being processed with a delay of at least one framing interval after the exposure of the second frame due to the subsequent read-out (see Fig. 2). The reported PIV field rates typically restrict themselves to the order of 5 Hz to 25 Hz. Gautier and Aider (2015) report PIV field rates up to 220 Hz/m. Very recently, kHz-rate velocity field processing was demonstrated by Bollt et al. (2025) using a GPU-based optical flow implementation. The present study investigates the potential of extending these real-time processing approaches replacing frame-based imaging with event-based sensing at the input end.

Real-time velocimetry based on EBV suggests two approaches, either particle tracking approaches (PTV) or ensemble processing methods such akin to cross-correlation based PIV. For either, the incoming event stream can be handled in two ways: The first approach is to group the incoming events into isolated clusters (blobs), representing individual particles. These clusters are then continuously tracked in time and space, resulting in Lagrangian particle tracks (Borer et al., 2017; Rusch and Rösger, 2023). The other approach sections the incoming data into time-slices of pre-defined duration thereby creating so called pseudo-images. All events in these frames are assigned with the same time-stamp, much

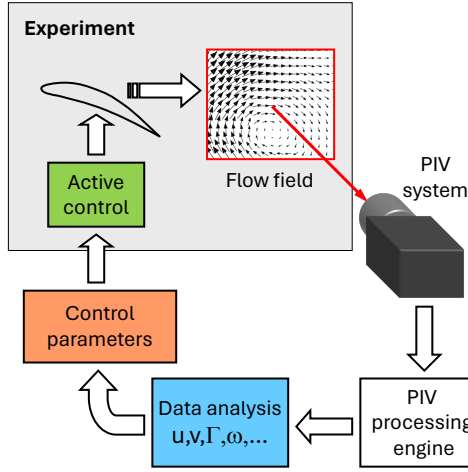


Figure 1: Possible implementation of active flow-control based on real-time flow field data (from Willert et al., 2010).

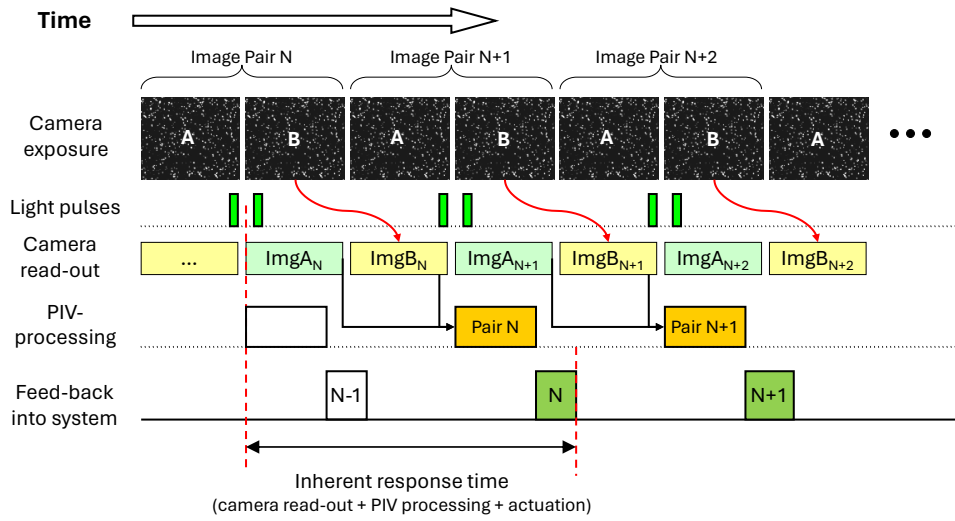


Figure 2: Data stream for correlation-based real-time PIV (from Willert et al., 2010)

like the image produced by a framing camera. The stream of pseudo-images can then either be subject to Lagrangian particle tracking (LPT) algorithms or processed using correlation-based algorithms (e.g. PIV mode). Output data for the latter are gridded Eulerian flow fields.

Initial efforts by the authors implementing the first approach (LPT), based on algorithms developed for 3D-LPT (Willert and Klinner, 2025, led to the observation that the processing speed scales inversely with the particle number within the camera's field of view (FOV). Even with fast, optimized segmentation algorithms (e.g. from OpenCV4 (Bradski, 2000)), the cluster detection step (`scipy.ndimage.label()`) barely achieved an update rate below 10 ms for $O(1000)$ particle images (= clusters). The subsequent tracking step also scales with the number of particles within FOV. Finally the Lagrangian particle tracks have to be re-interpolated onto some regular grid to allow flow field based analysis such as for reduced-order modeling in the context of flow control (Brunton and Noack, 2015).

While not really making direct use of the asynchronous nature of the data provided by the EBV camera, the generation of pseudo-frames from the event stream is considerably more efficient than trying to continuously extract (and track) clusters over time. This pseudo-frame generation does scale with event-rate but to a much lower degree than cluster detection schemes. Using approach 1 from above, the pseudo-frames can then be processed using conventional PIV algorithms, that can be easily parallelized in comparison to the previously mentioned LPT schemes. This results in a high frame throughput that only mildly depends on the event data rate, that is, particle image density.

Compared to framing camera based RT-PIV, the latency is reduced by at least one framing interval: the second pseudo-frame is immediately available at the end of the time-slicing period (see Fig. 3).

In the context of active flow control the continuously increasing speed of both imaging and computing hardware are enabling real-time PIV processing scenarios such as outlined in Fig. ??a (Willert et al., 2010; McCormick et al., 2024). The basic idea is to extract flow relevant information from real-time processed PIV data to provide flow-controlling parameters with minimal inference time (latency). The

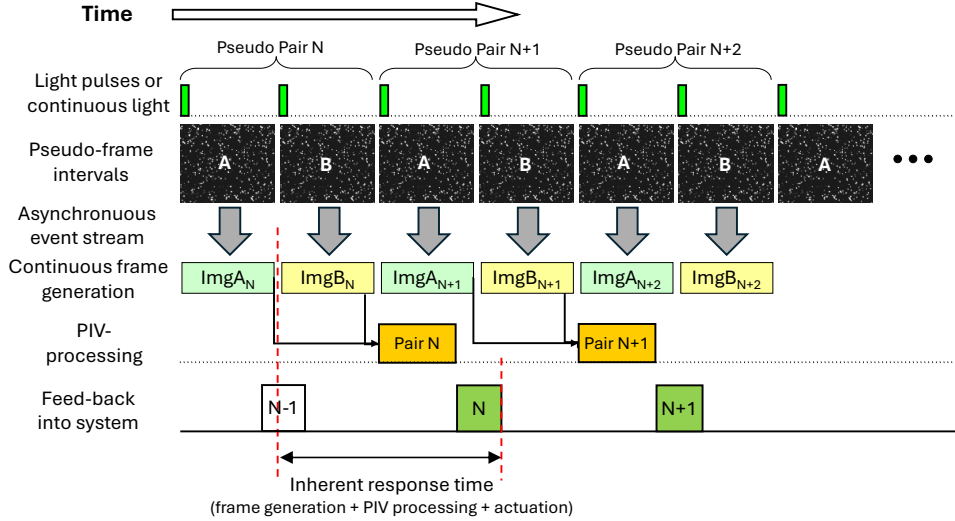


Figure 3: Data stream for pseudo-frame-based, real-time event-based imaging velocimetry (EBIV)

proposed contribution describes the implementation of a PIV system based EBV camera technology, capable of processing incoming imagery at high frame rates. Unlike conventional cameras, EBV cameras only detect and report intensity changes in the observed scene, providing an asynchronous data stream of contrast change events. With its reduced data rate, high dynamic range and potentially rapid data processing capability, EBV is a strong candidate for this purpose. To date, the real-time capability of EBV has not been fully exploited with the exception of the 3D-LPT implementations by (Rusch and Rösger, 2023) that reconstructs the (time-averaged) 3D flow field in real-time. However, their setup can only simultaneously track a small number of particles $O(50-100)$. The present work aims at acquiring and processing imagery containing $O(1000)$ particles necessitating approaches differing from conventional particle tracking.

While not really making direct use of the asynchronous nature of the data provided by the EBV camera, the generation of pseudo-frames from the event stream was found to be considerably more efficient than trying to continuously extract (and track) clusters over time. The pseudo-frames can then be processed using conventional PIV algorithms, that can be easily parallelized in comparison to LPT schemes.

Compared to RT-PIV implementations using framing cameras, the latency of real-time EBIV (RT-EBIV) is reduced by at least one framing interval: the second pseudo-frame is immediately available at the end of the time-slicing period (see Fig. 3) whereas the frame of conventional camera still needs to be read out. This is possible as the asynchronous events are accumulated on the pseudo-image array as they arrive from the sensor.

The following describes an initial implementation of RT-EBIV that operates by accumulating pseudo-image pairs at a defined framing interval followed by a single-pass PIV algorithm with subsequent validation. Bench-top validations on a rotating disc and a turbulent water flow in a small tank provided an initial assessment of the achieved processing speed of several hundred Hz.

2 Implementation

A simplified outline of the processing steps for RT-EBIV is provided in Fig. 4. 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 (Chandra et al., 2001). Here optimum processing speed is achieved when number of nodes is a multiple of utilized cores; otherwise a few unfinished threads block downstream processing until done.

Validation on the recovered displacement field is based on 3×3 neighborhood operators such as displacement magnitude, vector difference magnitude or normalized median filtering (Westerweel and Scarano, 2005).

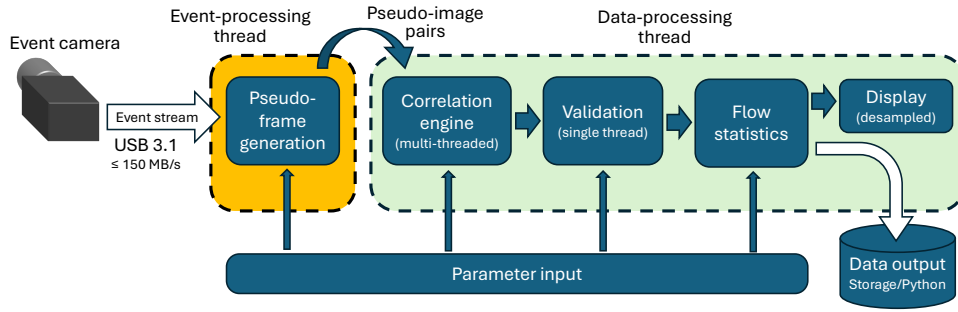


Figure 4: Processing pipeline for RT-EBIV

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 cameras based on the 4th generation EBV sensor IMX636 are used which is a co-development by Sony and Prophesee (Finatou et al., 2020). Initial measurements were performed with a Metavision EVK4 by Prophesee (F) or SilkyEvCam-HD by Century-Arks (JP) which are equivalent in packaging and size. Compared to these models, the recently introduced uEye-EVS by IDS Imaging (D) features a 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 on 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 MEV/s to 60 MEV/s without losing events; 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 (= particle images) that can be reliably captured. More comprehensive measurements were performed as part of the present investigation as described in Sec. 3.1.

2.1 Validation on a rotating disc

A simple and repeatable “flow” can be mimicked by a rotating disc covered with a scattering surface of small dots, sandpaper or retro-reflecting beads. Here, micro-beads of about $50 \mu\text{m}$ embedded in clear acrylic paint are sprayed onto a CD-ROM disk that was first coated with flat black paint. The disc is illuminated by the pulsed light of desired frequency and intensity from a LED ring light mounted on the camera objective for maximum retro-reflective activity of the micro-beads. Driven by a DC gear motor, the rotating disk provides stable long-term conditions.

2.2 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. 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 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 $20 \mu\text{m}$ 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 (CN), E30130) with a rated (CW) power of 5.5 W at a wavelength of 445 nm. For the measurements a pulse width of $10 \mu\text{s}$ to $20 \mu\text{s}$ was sufficient to achieve optimal event rates from a light sheet of 40 mm height. The equivalent pulse energy is in the range of $100 \mu\text{J}$ and reflects the high sensitivity of the EBV detector. 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 $50.2 \text{ mm} \times 28.2 \text{ mm}$.




Figure 5: Imaging setup with rotating disk to obtain maximum event-rate information using an IDS uEye EVS event camera (C). The rotating disk (D) is illuminated with LED light pulses from a ring light (R) using current pulse driver (P).

3 Results

3.1 Rotating disk

The functionality of the RT-EBIV software was tested and debugged using the rotating disk setup (see Sec. 2.1). It also provided valuable insight 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 into the following frame (avoidance of situations as in Fig. 6c)
- 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 EBIV 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 6 provides histograms of the acquired event data at a LED pulsing rate of 2 kHz for two different LED pulse durations (20 μ s and 25 μ s). For the shorter pulse duration the pulses are separated cleanly whereas the longer pulse results in event overflow from a given pulse period into the following (see Fig. 6c,d), at which point the performance of EBIV deteriorates considerably. The cleanly separated pulses at 54×10^6 Ev/s suggest that events can be captured at even higher event rates, however at 81×10^6 Ev/s this is not the case. However at a reduced pulsing frequency of 1 kHz a maximum event rate in excess of 100×10^6 Ev/s is achieved with the camera streaming at 146 MB/s. At 2 kHz and 81×10^6 Ev/s the effective data rate streamed by the camera increases to 175 MB/s which seems to indicate a streaming limitation around 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. 8 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 a reduced number of rows being accessed

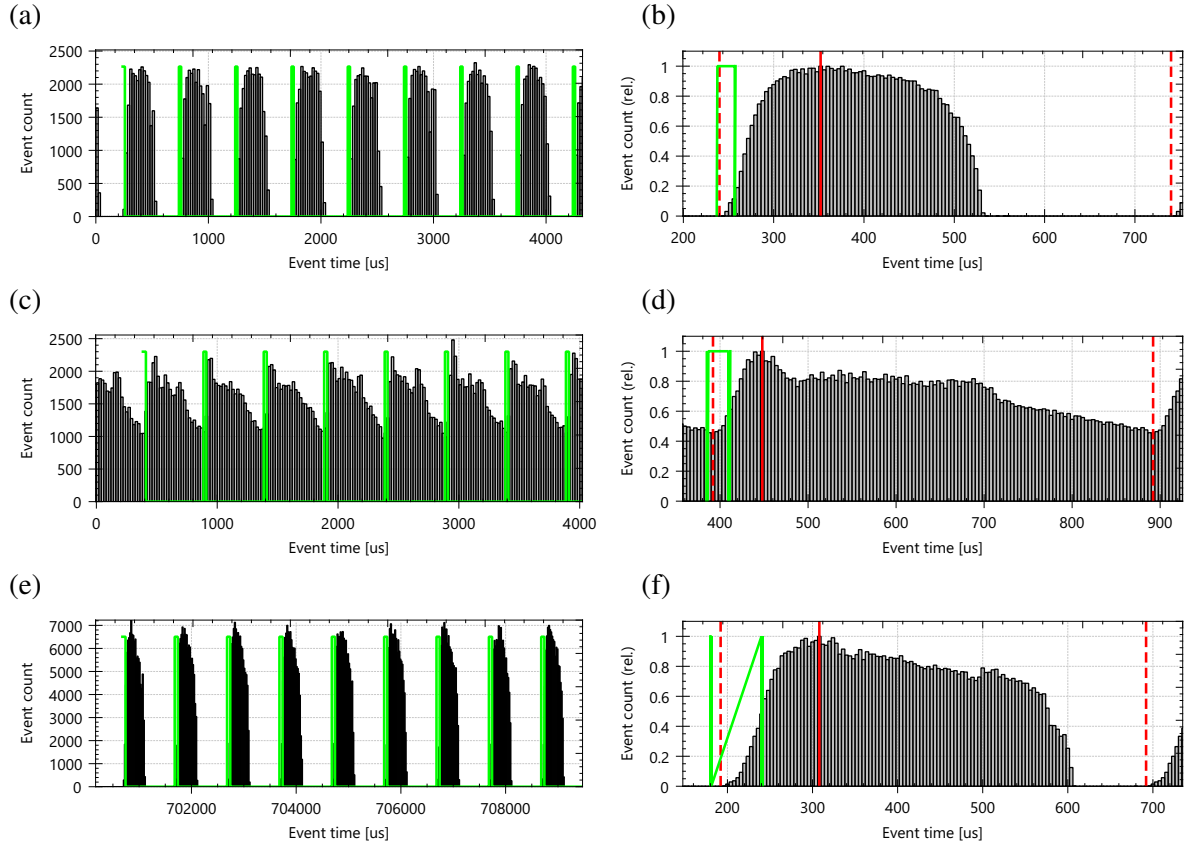


Figure 6: Histogram of event-data acquired from rotating disk setup with a light pulsing rate of 2 kHz at 54×10^6 Ev/s (a,b) and 81×10^6 Ev/s (c,d) and for 1 kHz at 105×10^6 Ev/s (e,f). Right column represents mean distribution over $O(100)$ pulse periods, indicated by red dashed lines ($20 \mu\text{s}$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

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. 8(a). Normalization by the sensor array size yields the particle density in ppp (see Fig. 8b). Not surprising, the highest ppp is achieved at lower light pulsing frequencies: at 200 Hz a ppp approaching 0.3 is feasible if the particle images are 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. 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 Water jet experiment

Off-line processing

Benchmarks of the water jet flow were performed using both live event streaming as well as off-line processing in terms of velocity maps per second. Off-line processing was used to quantify the maximum possible throughput in absence of the imposed pseudo-framing interval. The assessment is based on a 2 s recording acquired at a laser pulsing at 1 kHz and pulse width of $20 \mu\text{s}$ and positive event data rate of 31×10^6 Ev/s. The corresponding 2000 pseudo-frames 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 $O(10 \text{ Hz})$. Benchmarks in terms of velocity fields per second (fps) are obtained using multi-core processing with and without validation step enabled. The assessment is repeated on 2×2 binned images.




Figure 7: Maximum achievable event rate for different laser pulsing frequencies.

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 one 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 activity of the multi-process host operating system (MS Window 11 Pro).

Table 1: Performance of RT-EBIV in “free running mode” for different processing configurations using pre-recorded data using either an X64-processor (Intel(R) Xeon(R) W-2295 CPU, 18(36) cores, 3.00GHz) 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, $f_{\text{Laser}} = 1000 \text{ Hz}$, $31 \times 10^6 \text{ Ev/s}$. 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	$0.16\times$	816	718 080	$1.63\times$
32×32	880	off	361	317 680	$0.72\times$	907	798 160	$1.81\times$
32×32 (binned)	880	on	524	461 120	$1.05\times$	1425	1 254 000	$2.85\times$
32×32 (binned)	880	off	610	536 800	$1.22\times$	1750	1 540 000	$3.50\times$
single threaded performance								
32×32	880	on	86	75 680	$0.66\times$	182	160 160	$0.36\times$
32×32	880	off	87	76 560	$0.17\times$	187	164 560	$0.37\times$
32×32 (binned)	880	on	295	259 600	$0.59\times$	623	548 240	$1.25\times$
32×32 (binned)	880	off	311	273 680	$0.62\times$	673	592 240	$1.35\times$

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 $33 \times 10^6 \text{ Ev/s}$. For the assessment 5000 velocity maps are sampled with different processing configurations. The independently operating frame generation and processing threads along with the fixed framing rate imposes a delay in the hand-off of the current frame-pair to the PIV processing thread. 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. Fig. 12 provides histograms of the actual number of frame-pairs (e.g. velocity maps) sampled from the continuous image stream.

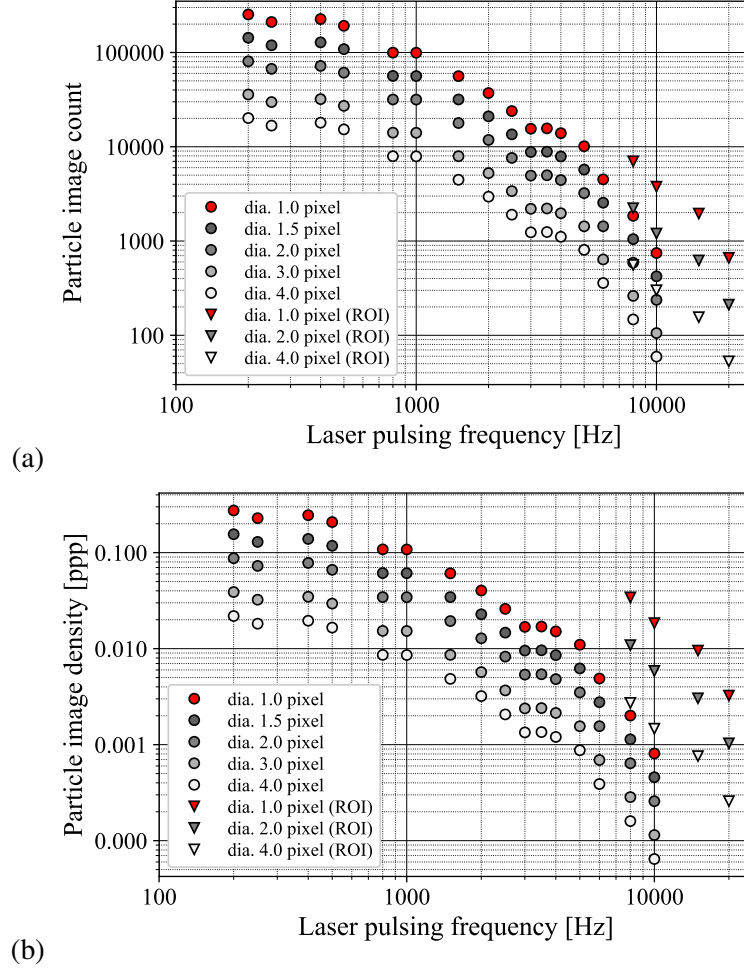


Figure 8: (a) maximum number of imaged particles at different laser pulsing frequencies and different particle image densities, (b) corresponding effective particle image density. ROI (∇) indicates an area of 1280×160 pixel.

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. 12a) 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. 12b). Image-binning further increases the processing speed: fields with $26 \times 15 = 390$ nodes are processed at a continuous rate of 720 Hz (Fig. 12c). In this case nearly two thirds of the generated frame pairs are processed without skipping.

Figure 13 visualize the acquired velocity field data in the form of space-time diagrams, so called kymograms. Here velocity data is sampled along the jet axis (x) for each available time-step and assembled into the diagrams shown in Fig. 13. Vertical stripes in the kymograms indicate a sporadic loss of signal particularly visible in the validation subplot of Fig. 13(b). The inclined features in the diagrams not only reveal the spatio-temporal coherence of the acquired data but show the convective nature of flow structures passing through the field of view. The change of slope with distance (x) from the nozzle indicate the deceleration of the flow. Although recorded at the same operating conditions (2000 Hz pseudo-framing), the slope of the features differs between Fig. 13(a) and Fig. 13(b). This is due to the increase of sampling rate from 270 Hz to 480 Hz (It should be noted that the horizontal axis of the kymograms represents the field index and not time).

Further information on achieved data quality is obtained by sampling the velocity field at a fixed position in space. Figure 14 shows the instantaneous velocity estimate (round markers), the mean velocity in the 3×3 node neighborhood (blue line) and the mean along the horizontal axis. Velocity spikes in particular in Fig. 14(b) indicate signal drop-out not detected by the validation step. Based on previous experiments (e.g. Willert, 2023; Franceschelli et al., 2025) the displacement estimate uncertainty is on




Figure 9: benchtop testing of the implemented RT-EBIV configuration on a turbulent water flow (F) using a Prophesee EVK2 camera (C) and pulse-modulated blue laser. Computer screen displays raw event imagery (E) and continuously updated flow statistics (S).




Figure 10: 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).

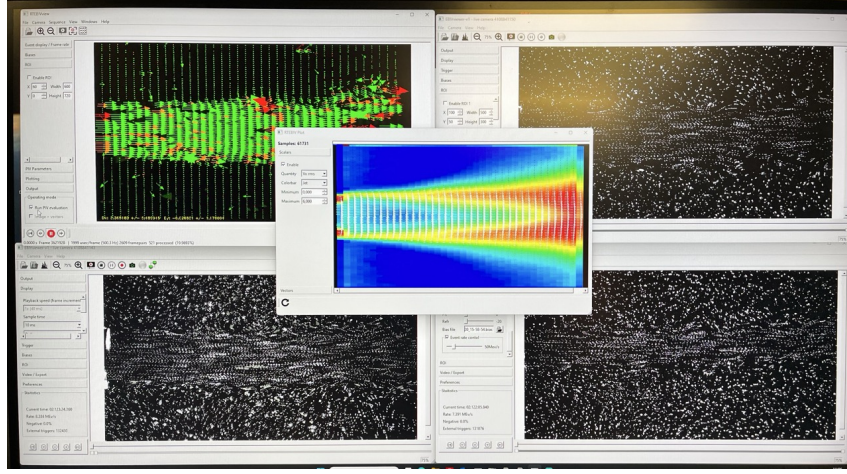


Figure 11: Image of computer screen showing continuous streaming of 4 synchronized event cameras in a 3d-PTV configuration at a laser pulsing frequency of 500 Hz. RT-EBIV processing is performed at top-left with the center window displaying the accumulated streamwise fluctuation of the jet after $N \approx 50000$ samples.

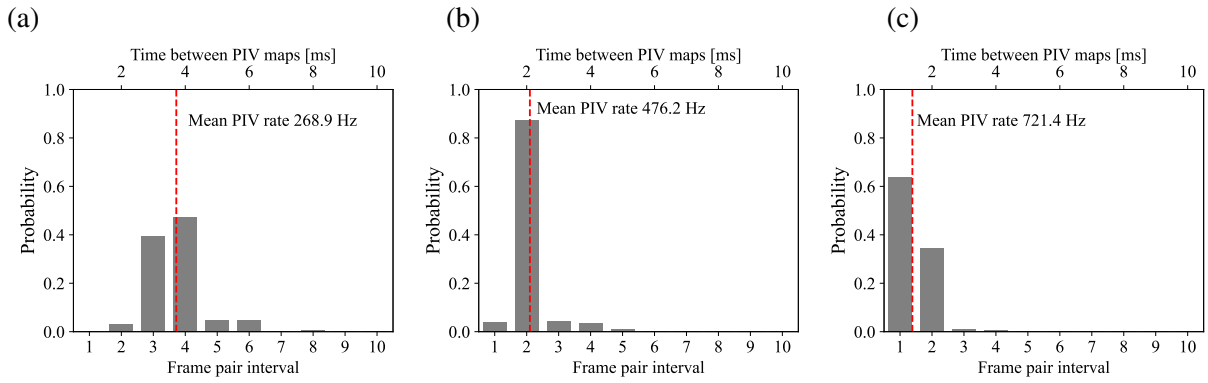
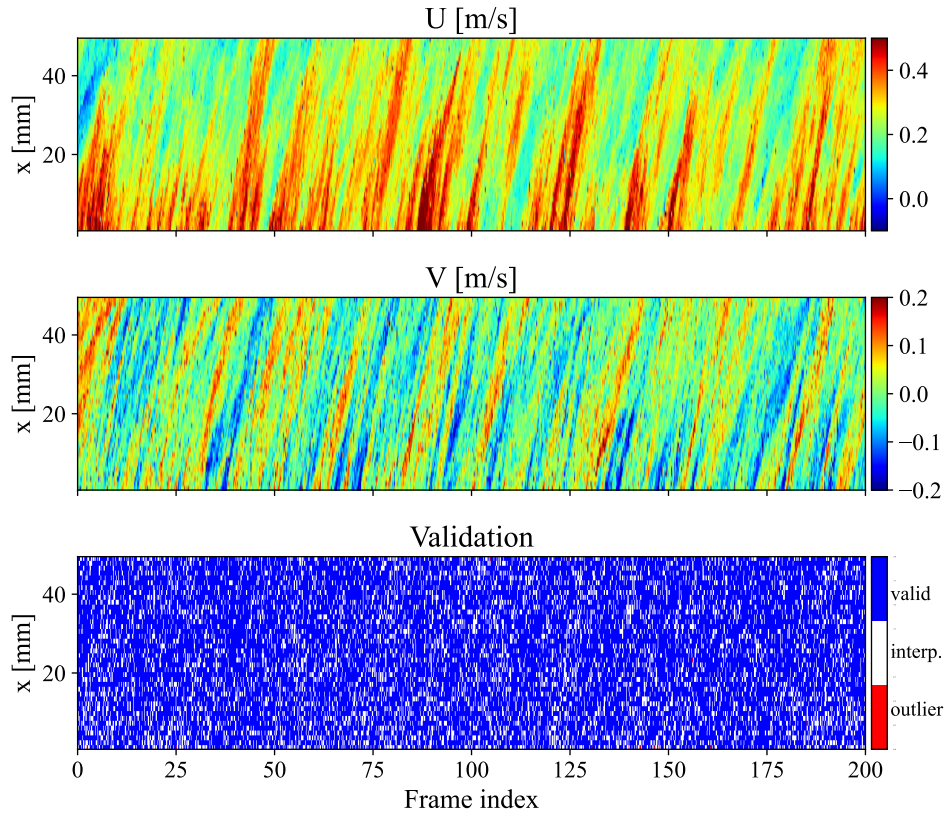


Figure 12: Histograms of the frames pair intervals while online processing live event data stream at 33×10^6 Ev/s at a pseudo-framing rate of 2000 Hz; (a): full field processing (880 nodes), (b): ROI (395 nodes), (c): binned pseudo-images processed with 24×24 sampling (390 nodes). Mean PIV-field rate is indicated by dashed red line.

(a)



(b)




Figure 13: Kymograms of both velocity components obtained by online processing of live event data stream at 40×10^6 Ev/s at a pseudo-framing rate of 2000 Hz. Data is sampled along the jet axis; (a): full field processing (40 nodes, 269 Hz), (b) ROI (79 nodes, 476 Hz)

the order of a 0.25 pixel or about 2 cm/s and is due to the binary nature of the imagery along with the rather rudimentary cross-correlation scheme.

Figure 15 provides kymograms obtained using 2×2 pixel binning pseudo-images. As pointed out in Fig. 12(c) the data drop-out (skipped frame pairs) is reduced to about 30 % leading to nearly time-contiguous data with a mean sampling rate of 720 Hz. This is achieved by reducing the number of nodes to 26×15 by simultaneously increasing the sample size to 48×48 pixel (no overlap). In binned mode, the sample size reduces to 24×24 . The enlarged samples also improve data yield as reflected in the higher validation rate of Fig. 15(c) compared to the previous two operating conditions using 32×32 pixel sampling. Also, the previously observed sporadic loss of signal is not present.

4 Discussion and Concluding Remarks

Using both off-line data as well as actual online measurements of a simple jet flow indicate that RT-EBIV can provide velocity fields in excess of 250 Hz. The pseudo-framing rate, which for pulsed-EBIV translates to the pulsing frequency of the light source (e.g. laser/LED), along with the density of the sampled velocity map primarily determine the throughput rate. At pulsing frequencies up to 500 Hz image pairs can be processed without skipping frames. Measurements at higher flow velocities require a higher pulsing frequency which results in frame-pairs to be skipped. Even so, the latency - namely the time between pseudo-frame capture and the availability of the velocity map - is not impacted. Compared to framing-based cameras, the overall latency is reduced by the availability of the second (pseudo)frame immediately after the end of the time-slice while a framing camera still needs to be read out.

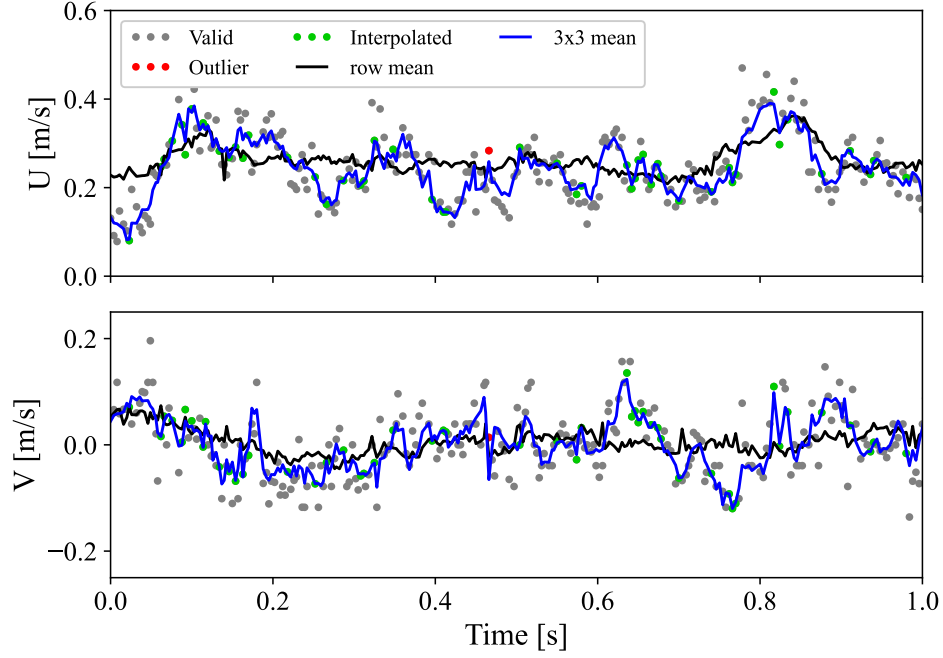
The previously reported values depend on the host processor architecture (e.g. type and number of utilized cores). The background activity of the host operating system was also found to impact the performance. Noteworthy is the significant gain in processing speed using an ARM64 processor for which camera drivers or not yet available.

A further pathway for increased processing speed could be the hardware-level generation of 1-bit pseudo-frame pairs using FPGA processing directly attached to the output of the EBV sensor (Blachut and Kryjak, 2023). Along this line Gao et al., 2024 reported throughput rates exceeding 1000 fps. However, care needs to be taken to avoid the additional frame transfer latency that is always present in conventional frame-based imaging. Evaluation kits for EBV application development on embedded systems including FPGA processing capabilities are commercially available (e.g. AMD/Xilinx Kria KV260 with Sony/Prophesee IMX636 EBV sensor). This could be coupled with GPU-based processing using implementations such as introduced by Bollt et al., 2025.

References

- Blachut K and Kryjak T (2023). High-definition event frame generation using SoC FPGA devices. In *2023 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA)*, p. 106–111. IEEE. doi:10.23919/spa59660.2023.10274447. URL <http://dx.doi.org/10.23919/SPA59660.2023.10274447>
- Bollt S A, Foxman S H, and Gharib M (2025). RapidPIV: Full flow-field khz piv for real-time display and control. doi:10.48550/arXiv.2504.17987. 2504.17987, URL <https://arxiv.org/abs/2504.17987>
- Borer D, Delbruck T, and Rösgen T (2017). Three-dimensional particle tracking velocimetry using dynamic vision sensors. *Experiments in Fluids* 58(165). doi:10.1007/s00348-017-2452-5
- Bradski G (2000). The OpenCV Library. *Dr. Dobb's Journal of Software Tools*
- Brunton S L and Noack B R (2015). Closed-loop turbulence control: Progress and challenges. *Applied Mechanics Reviews* 67(5):050801. ISSN 0003-6900. doi:10.1115/1.4031175. URL <https://doi.org/10.1115/1.4031175>
- Chandra R, Dagum L, Kohr D, Menon R, Maydan D, and McDonald J (2001). *Parallel programming in OpenMP*. Morgan kaufmann
- Finatou T, Niwa A, Matolin D, Tsuchimoto K, Mascheroni A, Reynaud E, Mostafalu P, Brady F, Chotard L, LeGoff F, Takahashi H, Wakabayashi H, Oike Y, and Posch C (2020). 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)*, pp. 112–114. doi:10.1109/ISSCC19947.2020.9063149

(a)



(b)

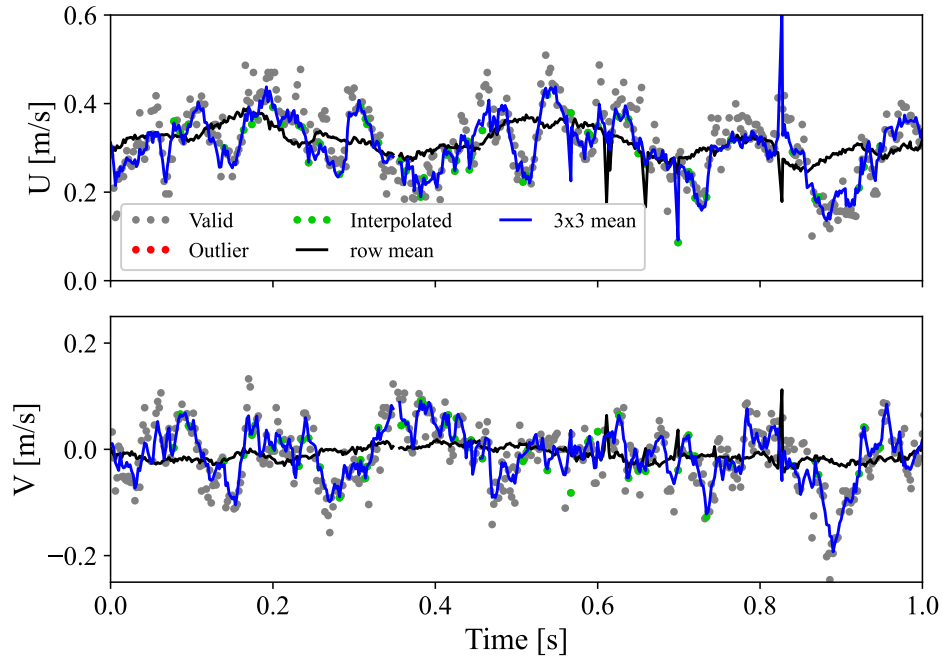


Figure 14: Velocity records of both velocity components over time obtained by online processing of live event data stream at 40×10^6 Ev/s at a pseudo-framing rate of 2000 Hz. Data is sampled at the center of field of view; (a): full field processing (880 nodes, 269 Hz), (b) ROI, 1280×240 (395 nodes 476 Hz).




Figure 15: 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.

- Franceschelli L, Willert C E, Raiola M, and Discetti S (2025). An assessment of event-based imaging velocimetry for efficient estimation of low-dimensional coordinates in turbulent flows. *Experimental Thermal and Fluid Science* 164:111425. ISSN 0894-1777. doi:10.1016/j.expthermflusci.2025.111425. URL <https://www.sciencedirect.com/science/article/pii/S0894177725000196>
- Gao Y, Zhang B, Ding Y, and So H K H (2024). 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*, p. 246–257. Association for Computing Machinery, New York, NY, USA. ISBN 9798400704185. doi:10.1145/3626202.3637558. URL <https://doi.org/10.1145/3626202.3637558>
- Gautier N and Aider J (2015). Real-time planar flow velocity measurements using an optical flow algorithm implemented on GPU. *Journal of Visualization* 18:277–286. doi:10.1007/s12650-014-0222-5
- Gautier N, Aider J L, Duriez T, Noack B R, Segond M, and Abel M (2015). Closed-loop separation control using machine learning. *Journal of Fluid Mechanics* 770:442–457. doi:10.1017/jfm.2015.95
- McCormick F, Gibeau B, and Ghaemi S (2024). Reactive control of velocity fluctuations using an active deformable surface and real-time PIV. *Journal of Fluid Mechanics* 985:A9. doi:10.1017/jfm.2024.292
- Ouyang Z, Yang H, Lu J, and Yin Z (2022). Real-time adaptive particle image velocimetry for accurate unsteady flow field measurements. *Science China Technological Sciences* 65:2143–2155. doi:10.1007/s11431-022-2082-4
- Rusch A and Rösgen T (2023). TrackAER: real-time event-based quantitative flow visualization. *Experiments in Fluids* 64:136. doi:10.1007/s00348-023-03673-0
- Siegel S, Cohen K, McLaughlin T, and Myatt J (2003). Real-time particle image velocimetry for closed-loop flow control studies. In *41st Aerospace Sciences Meeting and Exhibit, Reno (NV)*. doi:10.2514/6.2003-920. AIAA 2003-920
- Varon E, Aider J L, Eulalie Y, Edwige S, and Gilotte P (2019). Adaptive control of the dynamics of a fully turbulent bimodal wake using real-time PIV. *Experiments in Fluids* 60:124. doi:10.1007/s00348-019-2766-6
- Westerweel J and Scarano F (2005). Universal outlier detection for piv data. *Experiments in Fluids* 39(6):1096–1100. ISSN 0723-4864. doi:10.1007/s00348-005-0016-6. URL <http://dx.doi.org/10.1007/s00348-005-0016-6>
- Willert C (2023). Event-based imaging velocimetry using pulsed illumination. *Experiments in Fluids* 64:98. doi:10.1007/s00348-023-03641-8
- Willert C, Munson M, and Gharib M (2010). 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*. URL <https://elib.dlr.de/64688/>
- Willert C E and Klinner J (2025). Dynamic wall shear stress measurement using event-based 3d particle tracking. *Experiments in Fluids* 66:32. doi:10.1007/s00348-024-03946-2