A Monocular Event-Camera Motion Capture System

Leonard Bauersfeld and Davide Scaramuzza

Robotics and Perception Group, University of Zurich, Switzerland
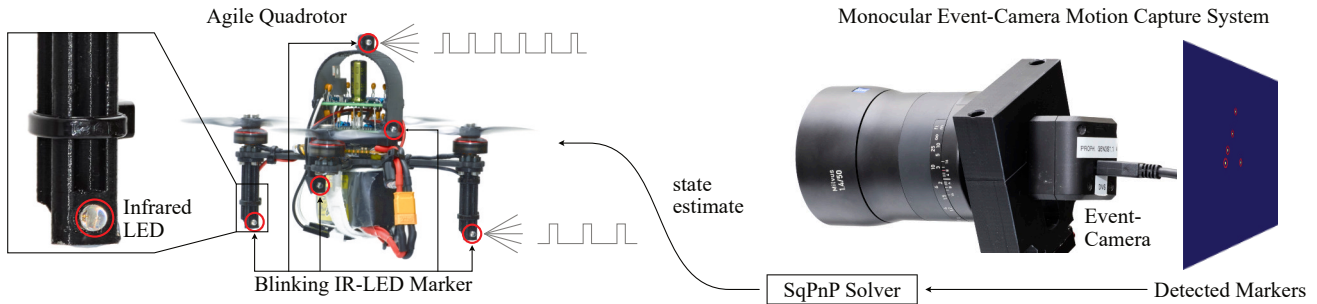
Fig. 1. Overview of the monocular event-camera motion capture system: the object to track (e.g. a small quadrotor) is equipped with $N >= 4$ infrared LED markers that blink at different frequencies. A single, calibrated event-camera is used to detect all markers and estimate the pose of the object by solving the perspective-n-points problem (PnP) with SqPnP [1]. The state estimate from this motion-capture system can then be used for closed-loop control.

Abstract— **Motion capture systems are a widespread tool in research to record ground-truth poses of objects. Commercial systems use reflective markers attached to the object and then triangulate pose of the object from multiple camera views. Consequently, the object must be visible to multiple cameras which makes such multi-view motion capture systems unsuited for deployments in narrow, confined spaces (e.g. ballast tanks of ships). In this technical report we describe a monocular event-camera motion capture system which overcomes this limitation and is ideally suited for narrow spaces. Instead of passive markers it relies on active, blinking LED markers such that each marker can be uniquely identified from the blinking frequency. The markers are placed at known locations on the tracking object. We then solve the PnP (perspective-n-points) problem to obtain the position and orientation of the object. The developed system has millimeter accuracy, millisecond latency and we demonstrate that its state estimate can be used to fly a small, agile quadrotor.**

I. INTRODUCTION

Motion capture systems are external camera systems that provide high-accuracy, low-latency pose estimates of an object by tracking markers attached to the object. The 3D position and orientation of the tracked object is then computed through multi-view triangulation. Such motion capture systems are ubiquitous in mobile robotics and they are used for platforms ranging from autonomous drones [2], [3], remotely controlled (RC) cars [4], and swimming vehicles [5] to manipulation [6]. However, such motion capture systems require the object to be in view of multiple cameras, making deployment in narrow, confined spaces extremely difficult.

Furthermore, the portability of the systems is negatively impacted by the long set-up time of the multi-camera system.

To overcome these drawbacks we developed the monocular, event-camera motion capture system shown in Fig. 1. In contrast to multi-camera motion-capture systems that rely on triangulation, a monocular system must solve the PnP (perspective-n-point problem) problem [7] to obtain the pose of the object in the camera frame. For a unique solution, at least four 3D $\leftrightarrow$ 2D correspondances must be known [8]. Put differently, it is not enough to detect the markers, but the markers must also be uniquely identified. Note, in the PnP-setting it is always assumed that the locations of the markers on the tracking object are known, for example from a CAD model.

Event-cameras are an excellent tool to overcome the above limitations and build a monocular motion capture system. By using blinking LEDs as active markers the markers can be easily detected by the event-camera. To uniquely identify each marker, the LEDs blink at different, known frequencies [9]. The blinking frequency for each detected marker is then measured from the event stream to associate the detection in the image with an LED marker. The high temporal resolution of event cameras makes it possible to use very fast blinking frequencies and obtain a low-latency system, which is ideally suited for challenging real-world robotics tasks.

The idea to use blinking LEDs in combination with an event camera for localization has been explored for over a decade [9]–[12], but most prior works focus on using the event camera on a mobile robot to localize w.r.t fixed markers. In this setting, the compute is constrained but markers can be large and consist of multiple LEDs. The availability of IMU measurements from the robot additionally simplifies the task when the event-camera is mounted the robot itself.

Our system is different as it uses a static event camera and active markers on the robot, and as such presents a much improved version of [9] featuring true 6-DOF tracking, a 50-fold improvement in accuracy, and a 4-fold improvement in update rate. To the best of the authors knowledge, it is the first true monocular motion-capture system providing millimeter-accuracy 6-DOF pose estimates at update rates as high as 1 kHz. This advance is enabled by modern sensors, an efficient, cache-friendly event-processing pipeline, utilizing a novel *signed delta-time volume* event representation in combination with the very accurate and robust SqPnP [1] algorithm to estimate the pose of the tracked object.

II. RELATED WORK

The field of monocular pose estimation is mostly re-searched in the context of robot localization and approaches can be categorized into three families:

1) standard cameras with fiducial markers (e.g. Aruco)
2) standard cameras with active point markers (e.g. LEDs)
3) event-cameras with active markers (e.g. blinking LEDs)

A. Pose Estimation with Standard Cameras

Most works on monocular pose estimation for mobile robots use standard cameras because those are readily avail-able on most robots. With a standard camera, the common approach [13]–[15] is to use fiducial markers (e.g. April-Tag [16], Aruco [17]) as such markers can robustly identified in RGB and grayscale images.

Fiducial markers have a minimum size because the camera must still be able to detect the structure of the marker. Therefore, using infrared LEDs has been proposed as an alternative solution [18] to further miniaturize the system. The LEDs can be detected by thresholding an IR-filtered image, but in this setup the points can't be uniquely identified as all LEDs look the same to the camera. The pose estimation algorithm thus has to exhaustively search through all possible combinations.

Independent of the marker type, standard cameras are not well-suited for a low-latency monocular motion capture system as the latency of the overall system is strongly limited by the framerate of the camera.

B. Pose Estimation with Event Cameras

In contrast to a frame-based camera an event-camera is inherently low-latency as events are output with microsecond latency. A major challenge is to design algorithms in a way that they are able to make use of this high update rate.

The first work [9] using event-cameras for localization—in a sense a direct predecessor of this work—uses blinking LEDs as active markers that can be identified based on their frequency. However, the pipeline is not able to estimate the location along the optical axis of the event-camera. Furthermore, it has noise levels on the order of 10 cm, and thus is unsuitable for closed-loop control.

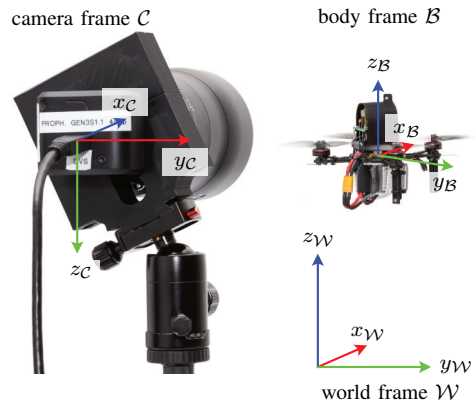More recent monocular event-camera pose-estimation pipelines [10], [12] rely on very large active markers, up to

Fig. 2. Definition of the camera frame ($z_\mathcal{C}$ is aligned with the optical axis), the body frame $\mathcal{B}$ and the world frame $\mathcal{W}$. The position of the active markers is defined in body-frame and must be known. The transform $T_{\mathcal{C}\mathcal{B}}$ is estimated through PnP and the pose of the camera in the world frame is assumed to be known.

$60 \times 60\,\text{cm}$ to achieve accurate localization. The markers are also containing multiple LEDs such that their identification is not just via frequency-detection, but similar to fiducial markers.

III. SYSTEM DESIGN

This section gives a brief overview of the developed monocular event-camera motion-capture system. Details re-garding the implementation of the blinking LED detection are given in Sec. V. The blinking LED circuit itself is discussed in Sec. VI.

A. Prerequisites

In order to obtain accurate 3D pose estimation of the tracked objects, the event-camera must be calibrated and the location of the blinking LEDs must be known. Furthermore, the transform $T_{\mathcal{C}\mathcal{W}}$ between the camera frame $\mathcal{C}$ and the world frame $\mathcal{W}$ must be known. The coordinate systems are illustrated in Fig. 2.

To calibrate the camera, we follow the approach from [19] where the calibration is performed by first converting the event stream into event frames and then calibrating these using Kalibr [20]. The calibration also estimates the lens distortion and in this work we rely on a double-sphere distortion model because of its accuracy and computation efficiency [21]. The event-camera is mounted horizontally on a stable tripod such that the transform from camera-frame $\mathcal{C}$ to world-frame $\mathcal{W}$ is fixed and known.

To ensure that the locations of the blinking LEDs are known and fixed, a 3D-printed LED holder is used. The locations of the LEDs are then directly known from the CAD model. The blinking frequency of each LED is also known, see Sec. VI for details on the circuit design.

B. LED detection

Robustly detecting blinking LEDs in an event-stream in real-time is the critical component of the motion-capture sys-tem. Events are processed in batches between 1 ms (1 kHz)

and 2.5 ms (400 Hz), depending on the desired pose update rate. In a first step, all pixels whose event rate is below a threshold are discarded. The threshold is calculated based on the assumption that each LED period at least triggers two events and that a transition is detected with a given probability (e.g. 80 %).

For all pixels with a sufficient event rate, the average period and standard deviation is calculated. For details on this process, see Sec. V. After identifying the average period, neighboring pixels with similar periods are clustered together. If the average period of the cluster is closer than $25\,\mu s$ to one of the expected blinking frequencies from the LED, it is matched to that LED. The centroid of each LED is tracked with a constant-velocity particle filter.

C. Pose Estimation

To estimate the pose from the detected LED centroids, the centroids are undistorted first. Then, the PnP-problem is solved with SQPnP [1]. We chose this algorithm over other well-known algorithms such as EPnP [22] because SQPnP is fast enough and globally optimal [1].

The pose-estimate in the camera-frame is finally transformed into the world-frame and published via ROS.

IV. EXPERIMENTAL RESULTS

Often a motion capture systems are chosen because of their low noise levels and low latency which is ideally suited for robot control. Therefore, we demonstrate the performance of the developed system, by flying the small drone depicted in Fig. 1 in closed-loop with the event-camera motion capture being the only source of state estimation. We use a Prophesee Gen 3.1 event-camera (640×480 px, 3/4 inch sensor) with either a 25 mm or a 50 mm lens resulting in a horizontal FoV (field of view) of 22 deg and 11 deg, respectively.

A. Pose Estimation Noise

In a first experiment the drone is rigidly placed at various distances in front of the camera. Then, 10 seconds of data are recorded with the event-camera and we calculate the standard deviation of the pose estimate. Since the drone is static, all deviations from the mean are only due to noise. The results of this analysis summarized in Fig. 3. A few interesting observations can be made which are subsequently discussed in detail:

1) The noise levels along the $z_\mathcal{C}$ axis are much larger compared to the $x_\mathcal{C}$ and $y_\mathcal{C}$ axis.
2) The SqPnP [1] algorithm achieves a much better performance than EPnP [22].
3) For SqPnP the noise levels in position scale quadratically with the distance from the camera, while the orientation noise scales linearly.

The $z_\mathcal{C}$-axis is the optical axis of the camera and hence the $z_\mathcal{C}$ coordinate can only be inferred from the scale of the object. Assuming that elongation of the object along the optical axis is small compared to the distance to the camera (i.e. the object is nearly flat), a well-known result from stereo-vision applies [7]: for a given inter-marker distance
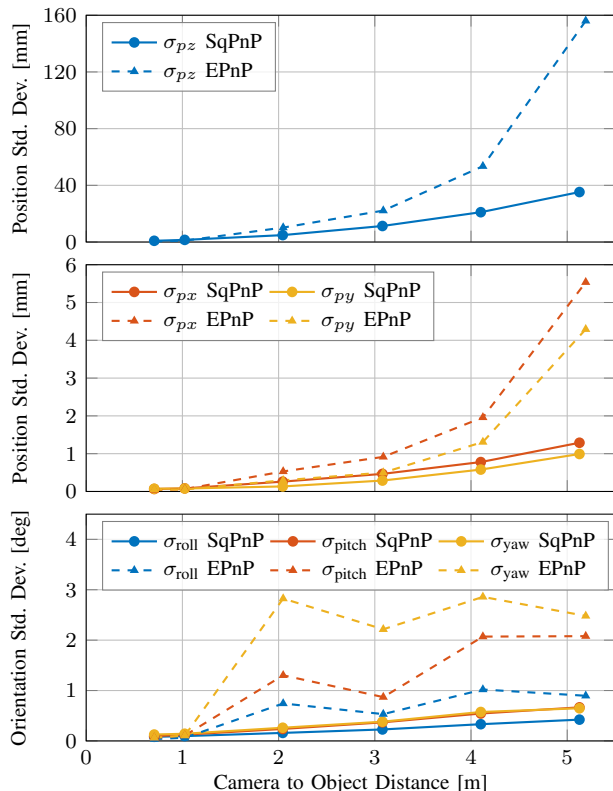

Fig. 3. The object is placed at distances between 70 cm and 5 m statically in front of the camera (with the 25 mm lens). The plots show the standard deviation in the position measurement ($z_\mathcal{C}$ and $x_\mathcal{C}$, $y_\mathcal{C}$) as well as the orientation measurements. We can clearly see that SqPnP [1] outperforms EPnP [22] by a large margin.

d, focal length f and a marker detection with uncertainty σ_u (in pixels) the depth uncertainty σ_{pz} scales with the square of the distance z as

$$\sigma_{pz} = \frac{\partial z_\mathcal{C}}{\partial u} \cdot \sigma_u = \frac{b \cdot f}{z_\mathcal{C}^2} \cdot \sigma_u \sim \frac{1}{z^2} \ . \tag{1}$$

For the positional errors in $x_\mathcal{C}$ and $y_\mathcal{C}$ we observe a similar quadratic dependency on the camera-object distance, however with much less noise. Intuitively, this makes sense as the translation along $x_\mathcal{C}$ and $y_\mathcal{C}$ are directly observable from each marker and thus the estimate is much more accurate.

The position noise plots also highlight the superior performance of SqPnP for this task: the optimization-based approach is able to estimate the position with much less variance given the same input data. This discrepancy becomes even larger when considering the orientation estimation shown at the bottom plot of Fig. 3. SqPnP dramatically outperforms EPnP which performs between two and four times worse. Interestingly, we observe that EPnP shows a large but nearly constant orientation uncertainty after 2 m, whereas SqPnP shows a linear increase in the noise standard deviation. Note that we do not compare against EPnP with nonlinear refinement as the OpenCV implementation requires at least six points for iterative refinement.
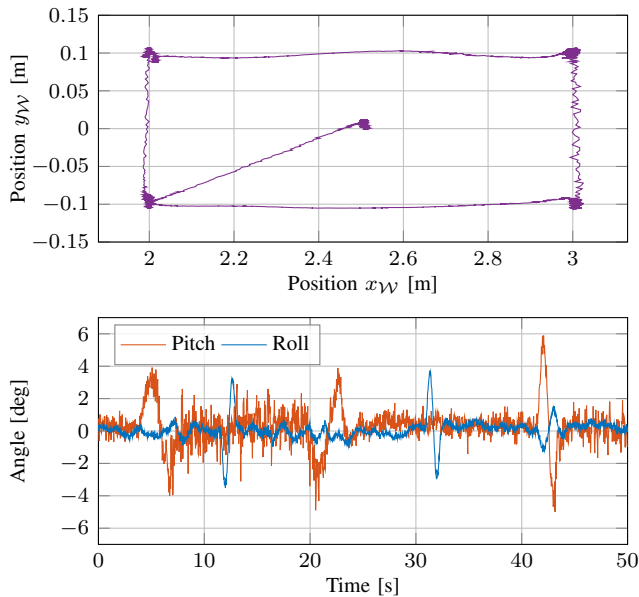
Fig. 4. Closed-loop experiments: the drone flies a rectangular pattern starting at (2,0.1) and then lands at a distance of 2.5 m. The event-camera is located at the origin of the coordinate system at $(x_\mathcal{W}, y_\mathcal{W}) = (0,0)$ and the optical axis of the 50 mm lens is aligned with the $x_\mathcal{W}$ direction. The bottom plot shows the roll and pitch angle measurements during the flight.

B. Closed-Loop Deployment

In this experiment we fly the small drone shown in Fig. 1 in closed-loop and the event-camera motion capture runs at 400 Hz leading to a system latency of 2.5 ms. The monocular event-camera motion capture system is the only source of state-estimation for the MPC controller [23] of the quadrotor. The drone is tasked to fly a rectangular pattern, hover and then land. The trajectory is shown in the top plot Fig. 4 and the roll and pitch angle measurements are shown at the bottom. When the drone is further away at $x_\mathcal{W} = 3$ m the position and orientation estimate become more noisy but overall we find that the system is able to safely fly the small drone, thereby demonstrating the performance of the developed system.

V. IMPLEMENTATION DETAILS

To ensure real-time operation of the event-camera motion capture system the delay in the processing pipeline must be kept to a minimum. This is achieved through an efficient, multi-threaded implementation in combination with filtering data early on in the pipeline. This section gives an overview of the most important concepts, specifically the data representation, the filtering and the multi-threading.

A. Event Data Representation

A single event as supplied by the camera is given as a four-tuple, consisting of an x and a y coordinate (both `uint16_t`), a polarity p which is either -1 or 1 (`int8_t`), and a timestamp t (`uint64_t`). Additionally, 3 bytes of padding are included for 16-byte alignment. The event camera supplies a stream of such raw events. For the following discussion of different data representations for blinking LED

detection, an event stream containing k events over a time period T coming from an event camera with image width W, height H and $N = W \times H$ pixels is considered.

1D Representations: The *event stream* is the most basic and raw representation of events and has recently gained some attention in combination with spiking neural networks [24]. For LED detection with classical CPU architectures however this representation is completely unsuitable. To extract any spatial information, the entire event stream must be searched for pixels with matching coordinates. Furthermore, having a memory layout where each event is stored serially is not efficiently using the cache: if we search for a given x coordinate, the remaining 14 bytes of the raw event representation are unused and just occupy cache space.

2D Representations: In the *event frame* representation the events are stored as a 2D grid by either summing the polarity or by counting the number of events. The accumulation is done for the time window of length t which represents the equivalent of the exposure time. The conversion from an event stream to an event frame is fast as can be done in linear time $\mathcal{O}(k)$ by iterating once over the stream. This representation is suitable for filtering out which pixels have a sufficiently high number of events to be candidates for a blinking LED, however it does not include any time information which would allow robust frequency detection.

A *time-surface* representation is also a 2D image, but each pixel in this 2D grid is assigned the value of the latest timestamp. For detecting blinking LEDs this representation is unsuitable as it contains no information related to periodic on-off transitions of a pixels.

3D Representations: In an *event volume*, events are stored as a 3D grid in a form that can be thought of as a stack of multiple event frames. This representation also includes time information and, given a sufficiently fine binning in the time domain, could be used to detect the frequency of a blinking LED. However, to accurately detect the frequency the binning would have to be very fine, yielding a huge memory footprint. For an accuracy of $5 \,\mu s$, a window length $t = 2$ ms, VGA resolution and `uint8_t` storage, the event volume would occupy 117 MB of memory. This size exceeds all levels of the processor cache, potentially affecting runtime adversely.

Signed Delta-Time Volume: To get past the shortcomings of those widely used event representation we propose a data representation that is ideally suited for the task of blinking LED detection: the *signed delta-time volume (SDTV)*. It is a 3D volume of size $W \times H \times D$ where D is the stack depth. For each pixel the time difference to the last event is stored and the polarity of the event is encoded in the sign of this time difference. This is possible because time must be monotonically increasing, so we can re-purpose the sign-bit for polarity encoding. The idea is illustrated in Fig. 5 for a single pixel stack of the SDTV.

Because of the fast blinking of the LEDs the time differences (in microseconds) between consecutive events are always within `int16_t` range, making storage compact. As most operations are done per pixel-stack, the memory layout
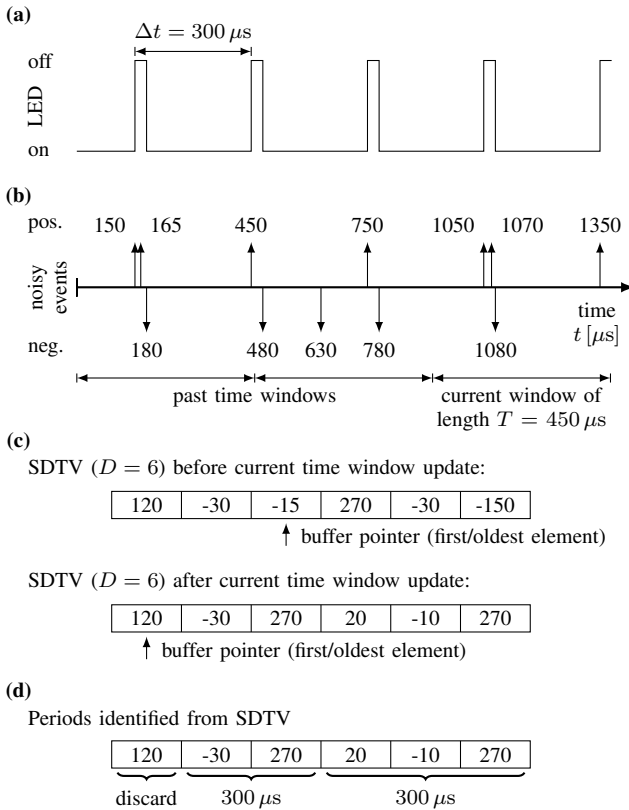
Fig. 5. Illustration on the construction of the *Signed Delta-Time Volume (SDTV)* from an event stream. **a)** The LED is blinking with a period of $300\,\mu s$ with a duty cycle of 10 %. **b)** A single pixel of the event camera records a noisy signal of this blinking LED. False double events (e.g. at $t = 150\,\mu s$, $165\,\mu s$) and spurious events (e.g. at $t = 630\,\mu s$) are included. **c)** Construction of the SDTV illustrated before processing the latest time window and after processing the time window. **d)** Periods robustly identified from the SDTV by summing up absolute time differences between negative $\rightarrow$ positive transitions (the first positive value is included). All events until the first positive $\rightarrow$ negative transition are discarded.

is such that the D-dimension is consecutive. Similar to the other representations, converting an event stream to SDTV is linear in the number of events. The signed delta-time volume is not computed per window of length T but updated as a cyclic buffer. This increases the accuracy of the frequency detection for LEDs blinking at a lower frequency than $f_{\max}$ since the amount of LED periods available for frequency identification is independent of the frequency.

The minimal depth D can be calculated based on the window length T and the frequency of the fastest LED $f_{\max}$ as $D = 2\,T\,f_{\max}$ because every LED should trigger two events (once on and once off) per period. Typical values of D are between 4 and 16, reducing the memory footprint by a factor of 25 to 100 compared to a event volume.

B. Filtering

When computing the signed delta-time volume representation from the event stream for a time window t of events, we also compute an event frame based on the event count of each pixel. Only pixels with more than $\beta \cdot 2t f_{\min}$ events are considered further where β is the probability that a transition triggers an event. We use $\beta = 0.8$ to purposely underestimate the detection probability.

For all selected pixels the SDTV is used to calculate mean, median and standard deviation of the period. The period is defined as the time between two on-events with at least one off-event in between as illustrated in Fig. 5d). In agreement with [9] we find that this is a robust measure. After rejecting pixels with a too-large standard deviation in the period, pixels are clustered together. Too small and too large clusters are rejected as the expected size of an LED is roughly known a priori. Clusters are then assigned to the individual LEDs by matching the measured average period in a cluster with the blinking frequencies of the LEDs. Each LED is tracked by a particle filter that gets the assigned clusters for each LED as an input.

C. Multi-Theading

In a real-time application like this, relying on generic multi-threading tools such as OpenMP can be problematic. For this reason, the threading is manually implemented to ensure optimal performance. Each thread in the pipeline shares its memory with the next thread in the pipeline. To ensure threads do not block each other, each thread allocates the required memory two times. During operation, the thread writes to one of its allocations while the other memory chunk is processed by the next pipeline step. Subsequently, the memory pointers are swapped and the newly filled batch processed.

The pipeline primarily consists of three threads. They
1) copy events from event camera driver into a buffer,
2) convert a linear event buffer into the optimized SDTV representation described in Section V-A, and
3) process the accumulated data to detect the LEDs, assign the LED clusters and solve the PnP (perspective-n-points) problem.

This design makes it possible to run the pipeline at hight speeds on a modern laptop. Speeds exceeding $1\,\text{kHz}$ are possible, but due the slowest LED blinking at $1700\,\text{Hz}$ increasing the processing frequency beyond $800\,\text{Hz}$ might degrade robustness.

VI. BLINKING LED CIRCUIT

A. Choice of LEDs

For the accuracy of the event-based mocap system it is important that the center of the blinking LED can be detected easily by the event camera. To achieve this, the LED should be small, very bright and have a short switching time to produce a well-defined rising and falling edge. These requirements are ideally met by LEDs optimized for pulsed operation, such as infrared LEDs for data transmission.

In this work, the Osram SFH4350 [25] IR-LED 3 mm was selected as it features an extremely short switching time of $12\,\text{ns}$. Being designed for opto-electronics it also permits up to $1\,\text{A}$ pulsed forward-current for pulses shorter than $100\,\mu s$ if the duty cycle is below 2 %. The LED has its peak emission around $850\,\text{nm}$ with a spread of $50\,\text{nm}$ which is within the sensitivity of most CMOS-based imaging sensors.
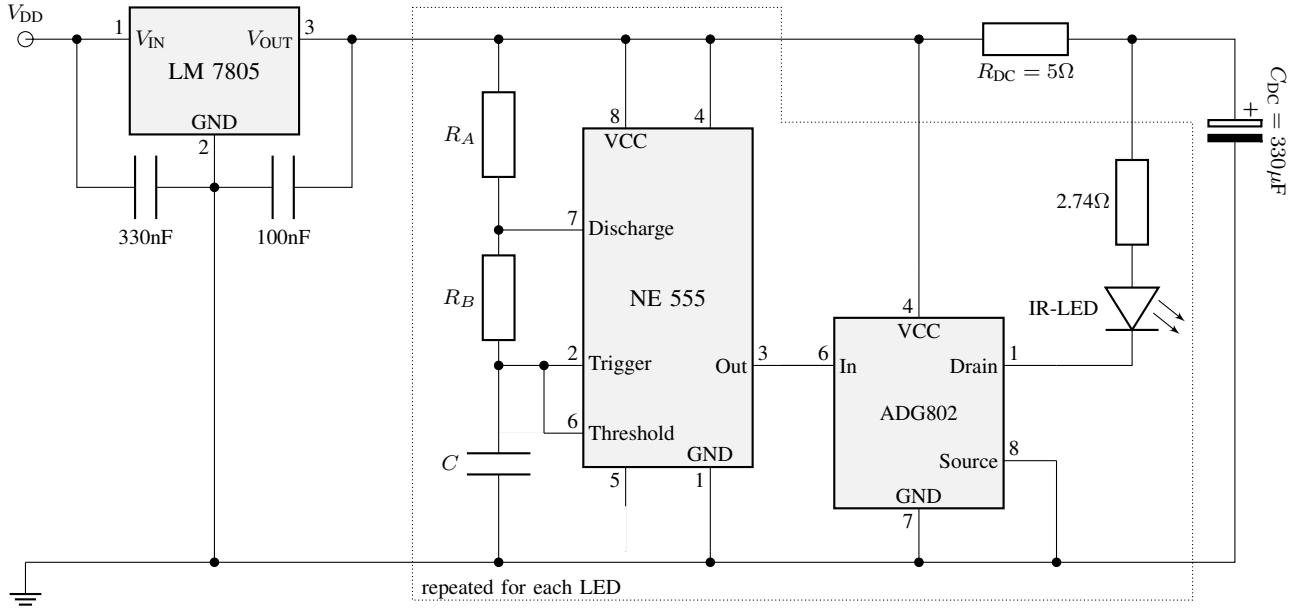
Fig. 6. Circuit diagram of the complete blinking LED circuit. For simplicity we only show one LED driver circuit (area enclosed in dotted line), which is then replicated multiple times to drive multiple LEDs. The input voltage range for the power supply is V_{DD} is 8-35 V. The resistor and capacitor values R_A, R_B, and C for the A-stable NE555 operation at different frequencies are given in Tab. I. For improved clarity the standard 100 nF ceramic bypass capacitor to stabilize VCC for the NE555 and ADG802 have been omitted in the circuit diagram above.

B. Circuit Design

The design of the LED driver circuit is tightly coupled with the entire event-based motion capture system:

1) Faster blinking frequencies increase the reactiveness of the system as at least one full period must be detected for identifying an LED. For robustness, a more conservative approach to detect at least two periods is better. Consequently, if the LEDs blink at 1 kHz the overall system is limited to 500 Hz output rate.

2) If the LEDs blink too fast, measuring the signal becomes difficult. Typically event cameras perform very well at measuring signals with frequencies up to 2-3 kHz [26]–[28] and the events are timestamped with 1 μs time resolution.

3) At least four LEDs are neccessary to yield a unique solution to the PnP problem.

4) the individual frequencies should not alias into each other. This means that, ideally, all LEDs have blinking frequencies within a factor of two.

5) Due to the limitations of the LED a duty cycle of 2 % can not be exceeded.

To control the blinking LED either a microcontroller or an analog circuit can be used. Because the high LED forward current of 1 A necessitates an analog output stage, we opted for a fully analog design using NE555 [29] precision timers. To generate the signal for the LEDs, the NE555 is operated in A-stable mode (c.f. Sec. 8.3.2 [29]).

The current output of the precision timer is limited to 200 mA, but its performance significantly degrades if the output current exceeds 10 % of the maximum value (c.f.

TABLE I
Resistor and capacitor values (see Fig. 6) for the different LEDs. The calculated periods as well as the measured frequencies f and duty-cycles α are listed.

Part Specification			Calc. from Sec. 8.3.2 [29]			Measured	
R_A [kΩ]	R_B [kΩ]	C [nF]	t_{on} [μs]	t_{off} [μs]	f [kHz]	f_{meas} [kHz]	α_{meas} [%]
68.1	0.39	10	2.7	477	2.094	1.73	0.66
59.0	0.39	10	2.7	415	2.413	1.98	0.75
51.1	0.39	10	2.7	359	2.781	2.29	0.87
44.2	0.39	10	2.7	312	3.207	2.61	0.99
40.2	0.39	10	2.7	284	3.520	2.86	1.09

Figure 3 of [29]). Therefore an SPST (single pole, single throw) digital switch is used. We selected an ADG802 as it features close to 1 A pulsed current and has typical switching times around 55 ns [30]. While considerably slower than the LED, it is still fast enough for the given application.

The LEDs, precision timers and switches are all supplied with a single LM7805 voltage regulator. This is possible because the time-averaged load is well below the design limit of the voltage regulator. Each NE555 draws $I_{NE555} = 3\,\mathrm{mA}$ of supply current. The time averaged current $\bar{I}_{LED}$ for an LED pulsed with a duty cycle $\alpha = t_{on}/t_{period}$ with a pulse current $I_p = 1\,\mathrm{A}$ is given by

$$\bar{I}_{LED} = \alpha \cdot I_p \tag{2}$$

Assuming there is $N = 5$ LEDs and they are operated at an average duty cycle of 1 % this leads to a total, time averaged current of

$$\bar{I} = N \cdot \left(\bar{I}_{LED} + I_{NE555} \right) = 65\,\mathrm{mA} \tag{3}$$
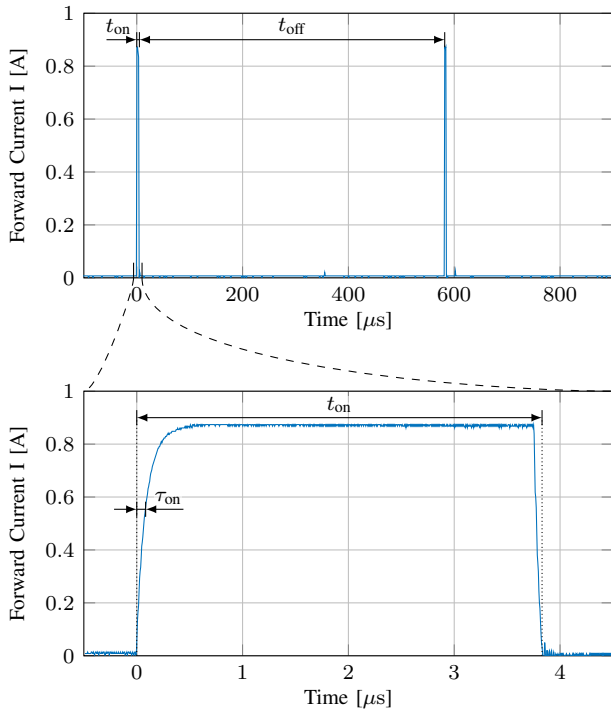
which is within specifications for an LM7805 [31] without

Fig. 7. Current through an LED blinking at the lowest frequency of 1.73 kHz. The upper plot shows the entire period $t = t_{on} + t_{off}$, whereas the lower plot shows only the pulse through where the LED is on for $3.8\,\mu s$. From the plot, we get a switch-on time constant τ_{on} of the LED of $84\,ns$ (time to reach $63\,\%$ of the steay-state value).

any additional cooling (given the TO220 package and a 12 V supply). Therefore, a sufficiently large decoupling capacitor C_{DC} charged through R_{DC} is used to supply the LEDs with power, effectively shielding the LM7805 from all current spikes caused by the LEDs. Based on all the above considerations, the circuit shown in Fig. 6 has been designed and subsequently manufactured into a PCB with SMD version of the NE555 and the ADG802. The dotted line marks the components in the circuit that are replicated for each LED.

The resistor and capacitor values used in the NE555 timer circuit are listed in Tab. I. The values have been calculated such that the frequencies of all 5 LEDs follow the points 2 to 5. After building and manufacturing of the PCB, the measured frequencies are also listed in Tab. I. Note that the mismatch w.r.t the calculated values is about $20\,\%$. This mismatch is consistent across 5 identical copies of the board and of no concern for the practical applications as it is straightforward to measure the blinking frequency with an oscilloscope. Exemplarily, Fig. 7 shows the current through one LED.

VII. CONCLUSION

In this technical report we present that a low-latency, high-accuracy monocular motion-capture system with an event-camera. The experiments with static objects show that the system has millimeter accuracy and the closed-loop experiments with the small quadrotor demonstrate that it is is well-suited for real-time control in mobile robotics tasks.

Originally, we developed the system with a focus on confined environments where a commercial, multi-camera system can not be used. However, we now believe that monocular event-camera motion-capture systems are highly relevant in a broader sense: event-cameras are expensive sensors, but the cost of a single event camera is 10 to 100 times less than that of a full motion-capture system, making the presented approach appealing for low-cost applications where the limited tracking volume (only in front of the event-camera) is a great trade-off for the small form factor, portability and cost.

REFERENCES

[1] G. Terzakis and M. Lourakis, "A consistently fast and globally optimal solution to the perspective-n-point problem," in *European Conference on Computer Vision*. Springer International Publishing, 2020, pp. 478–494.

[2] G. Ducard and R. D'Andrea, "Autonomous quadrotor flight using a vision system and accommodating frames misalignment," in *IEEE International Symposium on Industrial Embedded Systems*, 08 2009, pp. 261 – 264.

[3] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, Aug 2023.

[4] L. Fröhlich, C. Küttel, E. Arcari, L. Hewing, M. Zeilinger, and A. Carron, "Contextual tuning of model predictive control for autonomous racing," in *International Conference on Intelligent Robots and Systems (IROS)*, 10 2022.

[5] M. Menolotto, D.-S. Komaris, S. Tedesco, B. O'Flynn, and M. Walsh, "Motion capture technology in industrial applications: A systematic review," *Sensors*, vol. 20, no. 19, 2020.

[6] E. Aljalbout, F. Frank, M. Karl, and P. van der Smagt, "On the role of the action space in robot manipulation learning and sim-to-real transfer," *IEEE Robotics and Automation Letters*, vol. 9, no. 6, pp. 5895–5902, 2024.

[7] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.

[8] X.-S. Gao, X.-R. Hou, J. Tang, and H.-F. Cheng, "Complete solution classification for the perspective-three-point problem," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 930–943, 2003.

[9] A. Censi, J. Strubel, C. Brandli, T. Delbruck, and D. Scaramuzza, "Low-latency localization by active led markers tracking using a dynamic vision sensor," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 891–898.

[10] G. Ebmer, A. Loch, M. N. Vu, R. Mecca, G. Haessig, C. Hartl-Nesic, M. Vincze, and A. Kugi, "Real-time 6-dof pose estimation by an event-based camera using active led markers," in *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2024, pp. 8122–8131.

[11] A. Loch, G. Haessig, and M. Vincze, "Event-based high-speed low-latency fiducial marker tracking," in *Fifteenth International Conference on Machine Vision (ICMV 2022)*, W. Osten, D. P. Nikolaev, and J. J. Zhou, Eds., vol. 12701, International Society for Optics and Photonics. SPIE, 2023, p. 1270113.

[12] M. S. Salah, M. Chehadah, M. A. Humais, M. Wahbah, A. Ayyad, R. Azzam, L. D. Seneviratne, and Y. H. Zweiri, "A neuromorphic vision-based measurement for robust relative localization in future space exploration missions," *IEEE Transactions on Instrumentation and Measurement*, vol. 73, pp. 1–12, 2022.

[13] A. Breitenmoser, L. Kneip, and R. Siegwart, "A monocular vision-based system for 6d relative robot localization," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 79–85.

[14] B. Benligiray, C. Topal, and C. Akinlar, "Stag: A stable fiducial marker system," 2019. [Online]. Available: https://arxiv.org/abs/1707.06292

[15] M. Kalaitzakis, B. Cain, S. Carroll, A. Ambrosi, C. Whitehead, and N. Vitzilaios, "Fiducial markers for pose estimation," *Journal of Intelligent & Robotic Systems*, vol. 101, no. 4, p. 71, Mar 2021. [Online]. Available: https://doi.org/10.1007/s10846-020-01307-9

[16] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.

[17] S. Garrido-Jurado, R. Muñoz-Salinas, F. Madrid-Cuevas, and M. Marín-Jiménez, "Automatic generation and detection of highly reliable fiducial markers under occlusion," *Pattern Recognition*, vol. 47, no. 6, pp. 2280–2292, 2014.

[18] M. Faessler, E. Mueggler, K. Schwabe, and D. Scaramuzza, "A monocular pose estimation system based on infrared LEDs," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2014, pp. 907–913.

[19] M. Muglikar, M. Gehrig, D. Gehrig, and D. Scaramuzza, "How to calibrate your event camera," in *IEEE Conf. Comput. Vis. Pattern Recog. Workshops (CVPRW)*, June 2021.

[20] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 4304–4311.

[21] V. C. Usenko, N. Demmel, and D. Cremers, "The double sphere camera model," *2018 International Conference on 3D Vision (3DV)*, pp. 552–560, 2018.

[22] V. Lepetit, F. Moreno-Noguer, and P. Fua, "Epnp: An accurate o(n) solution to the pnp problem," *Int. J. Comput. Vision*, vol. 81, no. 2, p. 155–166, Feb. 2009.

[23] P. Foehn, E. Kaufmann, A. Romero, R. Penicka, S. Sun, L. Bauersfeld, T. Laengle, G. Cioffi, Y. Song, A. Loquercio, and D. Scaramuzza, "Agilicious: Open-source and open-hardware agile quadrotor for vision-based flight," *Science Robotics*, vol. 7, no. 67, p. eabl6259, 2022.

[24] M. Gehrig, S. B. Shrestha, D. Mouritzen, and D. Scaramuzza, "Event-based angular velocity regression with spiking networks," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4195–4202.

[25] *OSRAM SFH 4350 Datasheet*, OSRAM, 2022, version 1.8. [Online]. Available: https://look.ams-osram.com/m/25305e9a40863944/original/SFH-4350.pdf

[26] Y. Wang, Y. Shen, K. Xu, M. Hassan, G. Zhao, C. Xu, and W. Hu, "Towards high-speed passive visible light communication with event cameras and digital micro-mirrors," in *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 704–717.

[27] P. N. McMahon-Crabtree, L. Kulesza, B. J. McReynolds, D. S. O'Keefe, A. Puttur, D. Maestas, C. P. Morath, and M. G. McHarg, "Event-based camera refractory period characterization and initial clock drift evaluation," in *Unconventional Imaging, Sensing, and Adaptive Optics 2023*, J. J. Dolne, M. F. Spencer, and S. R. Bose-Pillai, Eds., vol. 12693, International Society for Optics and Photonics. SPIE, 2023, p. 126930V.

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

[29] *xx555 Precision Timers*, Texas Instruments, 2014, sept. 1973, Rev. Sept. 2014. [Online]. Available: https://www.ti.com/lit/ds/symlink/ne555.pdf

[30] *ADG801/ADG802*, Analog Devices, 2012, rev. B. [Online]. Available: https://www.analog.com/media/en/technical-documentation/data-sheets/adg801_802.pdf

[31] *LM Family Fixed Voltage Regulators*, Texas Instruments, 2016, feb. 2000, Rev. Sept. 2016. [Online]. Available: https://www.ti.com/lit/ds/symlink/lm340.pdf