

TCB-VIO: Tightly-Coupled Focal-Plane Binary-Enhanced Visual Inertial Odometry

Matthew Lisondra , Junseo Kim , Glenn Takashi Shimoda , Kourosh Zareinia , *Senior Member, IEEE*,
and Sajad Saeedi , *Member, IEEE*

Abstract—Vision algorithms can be executed directly on the image sensor when implemented on the next-generation sensors known as focal-plane sensor-processor arrays (FPSP)s, where every pixel has a processor. FPSPs greatly improve latency, reducing the problems associated with the bottleneck of data transfer from a vision sensor to a processor. FPSPs accelerate vision-based algorithms such as visual-inertial odometry (VIO). However, VIO frameworks suffer from spatial drift due to the vision-based pose estimation, whilst temporal drift arises from the inertial measurements. FPSPs circumvent the spatial drift by operating at a high frame rate to match the high-frequency output of the inertial measurements. In this letter, we present TCB-VIO, a tightly-coupled 6 degrees-of-freedom VIO by a Multi-State Constraint Kalman Filter (MSCKF), operating at a high frame-rate of 250 FPS and from IMU measurements obtained at 400 Hz. TCB-VIO outperforms state-of-the-art methods: ROVIO, VINS-Mono, and ORB-SLAM3.

Index Terms—Sensor fusion, visual-inertial SLAM.

I. INTRODUCTION

AGILE and mobile robotic systems often operate under strict power and processing constraints. As a result, there is an increasing demand for low-latency and low-power camera technologies. State-of-the-art algorithms utilizing conventional cameras typically achieve frame rates of 40–80 frames per second (FPS) [1], [2], [3]. These cameras rely on an architectural design where data is captured, digitized, and transferred to separate digital processing hardware for further computation. This traditional architecture, characterized by the modular separation of sensors and processors, can introduce significant latency and increase power consumption. On-sensor vision processing presents a novel paradigm that addresses these challenges by co-locating sensors and processors on the same chip. Focal-plane sensor processor arrays (FPSPs) [4] exemplify this approach,

Received 20 June 2025; accepted 19 September 2025. Date of publication 9 October 2025; date of current version 23 October 2025. This article was recommended for publication by Associate Editor G. Caron and Editor P. Vasseur upon evaluation of the reviewers' comments. This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). (*Matthew Lisondra and Junseo Kim contributed equally to this work.*) (*Corresponding author: Sajad Saeedi.*)

Matthew Lisondra and Glenn Takashi Shimoda are with the University of Toronto, Toronto, ON M4Y0E4, Canada.

Junseo Kim is with the Delft University of Technology, 2628 CD Delft, The Netherlands.

Kourosh Zareinia is with Toronto Metropolitan University, Toronto, ON M5B 2K3, Canada.

Sajad Saeedi is with University College London, WC1E 6EA London, U.K. (e-mail: s.saeedi@ucl.ac.uk).

This article has supplementary downloadable material available at <https://doi.org/10.1109/LRA.2025.3619774>, provided by the authors.

Digital Object Identifier 10.1109/LRA.2025.3619774

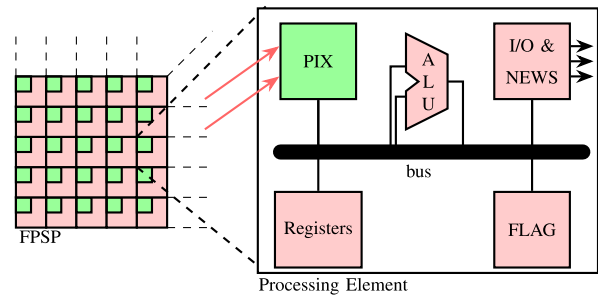


Fig. 1. In a focal-plane sensor-processor array, each pixel, i.e., processing element, combines a photosensor circuit (PIX) with on-pixel compute resources, including an arithmetic-logic unit (ALU), input/output (I/O) circuits, local communication links (NEWS), local memory (Registers), and activity control (FLAG). This architecture enables image processing to be performed directly on the sensor [5].

with each camera pixel equipped with a small processor capable of processing and sharing data with adjacent pixels (See Fig. 1). This tight integration of sensing and processing reduces power consumption and minimizes delays caused by communication overhead, making FPSPs highly suitable for real-world robotic applications [5].

Due to the limited space on FPSPs, per-pixel memory is highly constrained, and processing is typically performed in analog. Analog computation has inherent limitations, including reduced numerical precision, circuit inaccuracies, thermal effects, and noise leakage [6]. These constraints challenge implementing computer vision algorithms and can lead to noisy results. Integrating additional sensing modalities, such as inertial measurement units (IMU)s, using visual-inertial odometry (VIO) algorithms [7] can reduce the impact of the noise. However, adapting VIO algorithms to FPSPs requires a redesign of the processing pipeline to preserve the high-speed and low-power advantages of FPSPs.

Early VIO algorithms were filtering-based, using the Extended Kalman Filter (EKF) [7], [1], and the popular MSCKF [8], where landmark positions are marginalized from the state vector to reduce the computational complexity. Later, smoothing approaches emerged, optimizing states over a window [9] with tools like GTSAM [10], leading to variants such as VINS-Mono [2] and ORB-SLAM3 [3]. Despite improved accuracy, they can suffer from inconsistencies and linearization errors [11]. VIO methods can also be categorized as loosely coupled or tightly coupled [7]. Loosely coupled methods, such as MSF [12], process the two modalities independently and fuse them later. In contrast, tightly coupled algorithms, such as

OKVIS [9], jointly optimize visual and inertial data to achieve higher accuracy and robustness. For conventional cameras, however, the additional computation required by tightly coupled VIO often reduces update rates and increases the latency between image capture and state update [13]. FPSPs can mitigate this trade-off: their in-pixel parallel processing enables fast on-sensor computations, while their ultra-high frame rate shortens the delay between image captures. For FPSPs, several visual odometry algorithms have been proposed [14], [15]. To date, the only VIO algorithm is BIT-VIO [16], which is a loosely coupled method; no tightly coupled VIO approaches exist. This gap is addressed in this paper.

This paper presents the first tightly-coupled 6 degrees-of-freedom (DoF) VIO algorithm, designed for FPSPs. The algorithm coined, TCB-VIO, is a **T**ightly-**C**oupled VIO, utilizing **BIT**-enhanced features [14]. TCB-VIO processes high-speed, on-sensor binary edge images and feature maps using a novel binary-enhanced Kanade-Lucas-Tomasi (KLT) tracker. It is an extended and modified adaptation of the tightly coupled framework, OpenVINS [17]. The contributions of this paper are: 1) First high frame-rate tightly-coupled VIO for FPSPs achieving a high frame rate of 250 FPS. 2) The framework processes binary edge data and feature coordinates directly on-sensor, reducing computational cost. These outputs are then processed by the novel binary-enhanced KLT tracker at high frame rates. 3) Real-world evaluations comparing against ROVIO [1], VINS-Mono [2], and ORB-SLAM3 [3].

In the rest of the paper, Section II describes the background of FPSP. Section III explains the TCB-VIO algorithm. Section IV details our experimental results. Section V concludes the work.

II. BACKGROUND: SCAMP-5 FPSP

A traditional camera consists of a 2D array of light-sensitive pixels. In contrast, FPSPs integrate a processor within each pixel on the same chip [4] (See Fig. 1). FPSPs are also known as processor-per-pixel arrays (PPA) or cellular-processor arrays (CPA). One example of FPSP is the SCAMP-5 system [18], which features a 256×256 -pixel FPSP. Each pixel, or Processing Element (PE), includes a photosensor circuit (PIX) and compute resources. These PEs operate in parallel, executing identical instructions on local data in a Single Instruction, Multiple Data (SIMD) architecture. Each PE can also communicate with its north, east, west, and south neighbors, shown as NEWS in Fig. 1. Each PE contains 7 analog registers, 13 1-bit registers, and an ALU composed of 176 transistors, which means they can perform logical and arithmetic operations on the sensor. Arithmetic operations are performed directly in the analog domain using analog registers. Avoiding digitization improves the computational speed but introduces computational noise [18]. The computational results can be outputted in various formats, e.g. events, binary frames, or analog frames. The SCAMP-5 FPSP has been used in some robotics applications. Examples are target tracking [19], agile drones [20], fast inference [21], [22], obstacle avoidance [23], and localization [24]. It is worth noting that the emergence of event cameras has led to a new generation of VIO algorithms operating at high effective frame rates [25]. Event cameras are promising for VIO in high-speed and low-light scenarios. However, unlike FPSPs, they lack the flexibility of user-programmable processing directly on the focal plane, which enables customized computation.

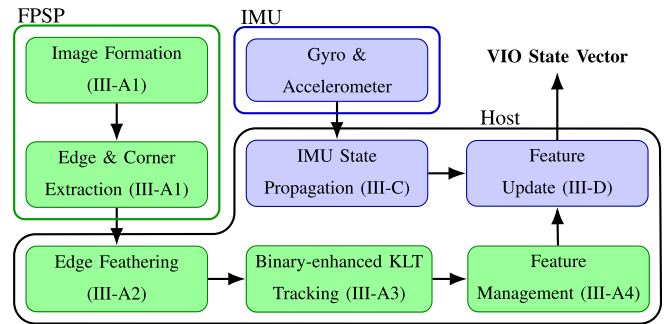


Fig. 2. TCB-VIO processing pipeline. Numbers inside each block correspond to the paper sections describing them. Green blocks denote visual processing: the FPSP performs on-sensor corner and edge extraction once an image is formed, and the extracted features are transferred to the host where a novel binary-enhanced KLT tracker ensures fast and robust tracking. Blue blocks denote inertial propagation and updates via MSCKF.

TABLE I
KEY PARAMETERS FOR TCB-VIO CONFIGURATION

Description (Symbol)	Value
Number of points extracted and tracked (N_p)	800
Max IMU clones in the sliding window, (N_c at Eq. 6)	15
Max SLAM-represented BIT feat in update ($N_{\text{SLAM, update}}$)	30
Max MSCKF-represented BIT feat in update ($N_{\text{MSCKF, update}}$)	60
Gaussian feathering variance for edge smoothing (σ_e)	2.5
Spatial window size in pixels (\mathbf{W}_i at Eq. (1))	21×21

III. PROPOSED METHOD

Fig. 2 shows an overview of the TCB-VIO pipeline. Edges and corner features are computed on the FPSP, then transferred to a host for further processing with a novel binary-enhanced Kanade-Lucas Tomasi (KLT) tracker. The results are fused with IMU measurements. The division of the workload between the FPSP and the host is dictated by the computational characteristics of each stage. Feature extraction is highly parallelizable across pixels and is well-suited to the SIMD architecture of the FPSP. In contrast, the data fusion and the optimization required for the tracker involve iterative and global computations that are not well-suited to the current FPSP architecture, and are therefore performed on the host.

A. Visual Processing

This section details visual processing steps, shown as green-filled blocks in Fig. 2.

1) *Image Formation and Feature Extraction*: The front-end visual processing of corner and binary features occurs directly on the SCAMP-5 FPSP. We use a modified version of FAST-16 implemented on the FPSP [26] for the front-end corner feature-part, where the processing of all pixels is done in parallel on FPSP [27], allowing for high frame-rate outputting. The same is done for the binary edge feature detection on the FPSP [14], where the FPSP analog registers were used, computed in parallel, like with the corner features. The rest of the processing steps are described as follows (See Fig. 3).

The set of N_p detected corners (see Table I for the size) and edge features acquired on the FPSP is stored in two binary feature maps: the corner map \mathbf{M}_c and the edge map \mathbf{M}_e . These maps are transmitted to the host device at a high frame rate of 250 FPS and are formally defined as $\mathbf{M}_c, \mathbf{M}_e \in \{0, 1\}^{256 \times 256}$,

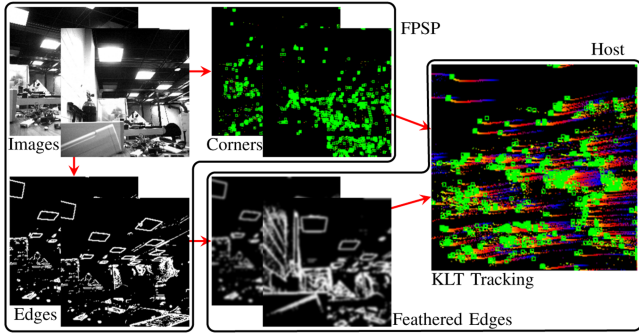


Fig. 3. Overview of the binary-enhanced KLT Tracking. The FPSP generates binary corners and edges. On the host, binary edges are feathered, and the KLT tracker operates in spatial windows centered on each corner feature.

where $M_c(i, j) = 1$ indicates that a corner feature is present at pixel location (i, j) , and $M_e(i, j) = 1$ denotes the presence of an edge feature at that pixel.

2) *Edge Feathering*: Once a binary edge image is received on the host, it is converted to an 8-bit grayscale image, where all edges take the value of 128. To enhance edge continuity and introduce smooth intensity transitions for improved feature tracking, a Gaussian-based feathering operation is applied to the grayscale image, generating $M_{\text{feathered}}$. The goal is to produce softly graded edge intensities that improve robustness in gradient-based feature processing. We empirically determine a suitable value of variance of the Gaussian kernel σ_e (see Table I for the value) that produces desirable feathering, preserving high intensity at the center of edges while introducing sufficient gradient smoothness for tracking stability.

3) *Binary-Enhanced KLT Tracking*: To enable efficient, high frame-rate visual tracking on binary data from the FPSP, we adapt the traditional KLT tracking algorithm [28], [29], [30] to operate on the feathered grayscale image. Unlike conventional KLT pipelines, which rely on dense intensity derivatives from natural images, our approach utilizes minimal yet structured gradients derived from binary edge maps to maintain visual tracking robustness at significantly reduced computational cost.

The binary-enhanced KLT utilizes both corners and feathered edges. Each corner pixel (x_i, y_i) for which $M_c(x_i, y_i) = 1$ is selected as a tracking point, forming the set $\mathbf{C} = \{(x_i, y_i)\}$.

While the corner locations come from the binary corner map M_c , the actual intensity information used for tracking is drawn from $M_{\text{feathered}}$, which encodes smooth gradients generated from the edge map M_e . For each corner, a spatial window \mathbf{W}_i (see Table I for the size) is centered on (x_i, y_i) , and optical flow is computed over these windows. The feathering process introduces local intensity uniqueness between adjacent edges, ensuring that similar binary edge structures become distinguishable for gradient-based motion estimation.

For each feature, the displacement $\Delta \mathbf{u}$ is computed by minimizing the photometric error between consecutive frames:

$$\sum_{(x,y) \in \mathbf{W}_i} (I_t(x, y) - I_{t+1}(x + \Delta u_x, y + \Delta u_y))^2, \quad (1)$$

where I_t and I_{t+1} are the feathered edge images at times t and $t + 1$, respectively. The optimization is solved iteratively,

leading to $\Delta \mathbf{u}^k = \mathbf{H}^{-1} \mathbf{g}$, where [28]:

$$\mathbf{H} = \sum_{(x,y) \in \mathbf{W}_i} \nabla I(x, y) \nabla I(x, y)^\top, \quad (2)$$

$$\mathbf{g} = \sum_{(x,y) \in \mathbf{W}_i} \nabla I(x, y) (I_t(x, y) - I_{t+1}(x, y)). \quad (3)$$

The updated feature locations are given by $\mathbf{C}_{t+1} = \mathbf{C}_t + \Delta \mathbf{u}^k$. Tracking continues until convergence or the maximum iteration count: $\|\Delta \mathbf{u}^k\| < \epsilon, \forall (x, y) \in \mathbf{C}$. The binary-enhanced KLT tracker outputs a sequence of high-frequency 2D features with observations $\mathbf{z}_i^k \in \mathbb{R}^2$, where feature i is tracked across a sliding window of camera frames indexed by $k \in \{t_i, \dots, t_i - N_c\}$. These features serve as visual measurements, $\mathbf{z}_{\text{meas}}(\cdot)$, in the MSCKF backend to update the state (Section III-D). Beforehand, a decision is made whether to include them in the state vector, as discussed next.

4) *Feature Management*: The MSCKF uses features to constrain camera poses during state estimation [8]. These features are not explicitly maintained in the state vector and are therefore referred to as out-of-state (or MSCKF) features. Following Li and Mourikis [31], to improve robustness, a subset of features can be explicitly included in the filter state. These are called in-state (or SLAM) features, as they provide long-horizon geometric constraints. In TCB-VIO, the in-state features are stored in \mathbf{x}_{BIT} , described in Section III-B, and updated as described in Section III-D.

Despite the high frame rate of the FPSP, careful management of SLAM (in-state) features is required, since considering all detected features as in-state would significantly increase both the state dimension and the computational cost of updates. To balance accuracy and efficiency, only a small, reliable subset of features is promoted to the state vector. This classification is based on spatiotemporal observability and track length: features with sufficiently long and stable tracks are designated as in-state, whereas shorter-lived tracks remain out-of-state [17]. Features tracked for fewer than N_c frames, defined in Section III-B, remain out-of-state; those tracked for at least N_c are promoted to in-state, with $N_c = 15$ in our experiment (Table I).

B. Visual-Inertial State Vector

Section III-A described how features are extracted and processed. This section describes the inertial/visual state vectors, followed by inertial/visual fusion in the remaining sections. The notation used here is adopted from OpenVINS [17]. The state vector \mathbf{x}_i consists of all current estimates of the filter up to time t_i . It has four parts: the current navigation state \mathbf{x}_{NAV} , a set of $N_c \in \mathbb{N}$ IMU clones of a sliding window \mathbf{x}_{IMU} , a set of $N_r \in \mathbb{N}$ reprojected from 2D-to-3D mapped BIT-features \mathbf{x}_{BIT} (explained in more detail in Section III-A4) and, a set of camera calibration intrinsics and extrinsics parameters \mathbf{x}_{C} :

$$\mathbf{x}_i = [\mathbf{x}_{\text{NAV}}^\top \quad \mathbf{x}_{\text{IMU}}^\top \quad \mathbf{x}_{\text{BIT}}^\top \quad \mathbf{x}_{\text{C}}^\top]^\top. \quad (4)$$

The current IMU navigation state \mathbf{x}_{NAV} is defined as

$$\mathbf{x}_{\text{NAV}} = [I_i^{\mathbf{q}\top} \quad G^{\mathbf{p}}_{I_i} \quad G^{\mathbf{v}}_{I_i} \quad \mathbf{b}_{\omega, i}^\top \quad \mathbf{b}_{a, i}^\top]^\top, \quad (5)$$

where $I_i^{\mathbf{q}\top}$, $G^{\mathbf{p}}_{I_i}$, $G^{\mathbf{v}}_{I_i}$ represents the quaternion orientation, position, velocity from the IMU-frame (I_i) at time t_i to the global world-frame G , respectively. The $\mathbf{b}_{\omega, i}^\top$, $\mathbf{b}_{a, i}^\top$ are the gyroscopic

and accelerometer biases. The IMU clones is

$$\mathbf{x}_{IMU} = \left[\mathbf{x}_{IMU_{t_i}}^\top \cdots \mathbf{x}_{IMU_{t_i - N_c}}^\top \right]^\top, \quad (6)$$

where $\mathbf{x}_{IMU_{t_i}}^\top = [{}^G \mathbf{q}^\top \ G \mathbf{p}_{I_i}^\top]$. Note that the indexing is from t_i to $t_i - N_c$, forming a sliding window of states, from current to a past window of N_c prior. The extracted BIT features are projected in the global frame as 3D key features,

$$\mathbf{x}_{BIT} = \left[{}^G \mathbf{p}_{BIT,1}^\top \cdots {}^G \mathbf{p}_{BIT,N_r}^\top \right]^\top, \quad (7)$$

where ${}^G \mathbf{p}_{BIT,i} = (x_i, y_i, z_i)$ is in global frame G . The frame-to-frame tracking is processed at 250 FPS using a novel binary-enhanced Kanade-Lucas-Tomasi (KLT) tracker, in which the binary edge images and feature maps are computed on-sensor.

Finally, between the camera-frame C with respect to the IMU-frame I_i at time t_i , the camera intrinsics and extrinsics,

$$\mathbf{x}_C = \left[{}^I_i \mathbf{q}_i^\top \quad {}^C \mathbf{p}_{I_i}^\top \quad {}^G \Phi_{i0}^\top \right]^\top \quad (8)$$

are continually optimized, where ${}^I_i \mathbf{q}_i^\top, {}^C \mathbf{p}_{I_i}^\top$ represents the quaternion orientation and position from the IMU-to-camera frame. ${}^G \Phi_{i0}^\top$ is the camera intrinsic vector that updates in time, in the global frame G .

C. Propagation of State by IMU

The IMU is taken as a 6-axis input, where measurements in the local IMU frame are

$${}^I \boldsymbol{\omega}_{local}(t) = {}^I \boldsymbol{\omega}_T(t) + \mathbf{b}_G(t) + \mathbf{n}_G(t), \quad (9)$$

$${}^I \mathbf{a}_{local}(t) = {}^I \mathbf{a}_T(t) + {}^I_G \mathbf{R}(t) {}^G \mathbf{g} + \mathbf{b}_A(t) + \mathbf{n}_A(t), \quad (10)$$

where ${}^I \boldsymbol{\omega}_T(t), {}^I \mathbf{a}_T(t)$ are the true and un-noised gyroscope and accelerometer measurements. $\mathbf{b}_G(t), \mathbf{b}_A(t)$ represent the gyroscopic and accelerometer biases, respectively. ${}^I_G \mathbf{R}(t)$ is the rotation from global frame G to IMU-frame, and is used in transforming the approximate estimated gravity vector \mathbf{g} to the local IMU frame. $\mathbf{n}_G(t), \mathbf{n}_A(t)$ are white Gaussian noises.

The IMU state evolves over time with notations adopted from [32] and kinematic propagation model from [17], [8]:

$${}^I_G \mathbf{q}(t) = \frac{1}{2} \begin{bmatrix} -[{}^I \boldsymbol{\omega}_T(t) \times] & {}^I \boldsymbol{\omega}_T(t) \\ -{}^I \boldsymbol{\omega}_T(t)^\top & 0 \end{bmatrix} \begin{bmatrix} {}^I \boldsymbol{\omega}_T(t) \\ {}^I_t \mathbf{q} \end{bmatrix}, \quad (11)$$

$${}^G \mathbf{p}_I(t) = {}^G \mathbf{v}_I(t), \quad {}^G \dot{\mathbf{v}}_I(t) = {}^I_t \mathbf{R}^\top {}^I \mathbf{a}_T(t), \quad (12)$$

$$\dot{\mathbf{b}}_G(t) = \mathbf{n}_G(t), \quad \dot{\mathbf{b}}_A(t) = \mathbf{n}_A(t). \quad (13)$$

The true un-noised gyroscope and accelerometer IMU measurements $\boldsymbol{\omega}_T(t), \mathbf{a}_T(t)$ can be obtained by rearranging Eq. (9)-(10), respectively with raw input IMU measurements $\boldsymbol{\omega}_{local}(t), \mathbf{a}_{local}(t)$ in the local-frame.

The solutions of the kinematic equations start with [17]

$${}^I_{G^{i+1}} \mathbf{R} = \exp \left(- \int_{t_i}^{t_{i+1}} {}^I \boldsymbol{\omega}_T(Z) dZ \right) {}^I_i \mathbf{R} \quad (14)$$

Here, the rotation from global frame G to IMU-frame from the last estimate ${}^I_i \mathbf{R}$ is updated incrementally. Then the position is acquired from the last position, velocity ${}^G \mathbf{p}_{I_i}, {}^G \mathbf{v}$, respectively, involving a double-integral on the last acceleration $\mathbf{a}_T(Z)$:

$${}^G \mathbf{p}_{I_{i+1}} = {}^G \mathbf{p}_{I_i} + {}^G \mathbf{v}_{I_i}(t_{i+1} - t_i) - \frac{1}{2} {}^G \mathbf{g}(t_{i+1} - t_i)^2$$

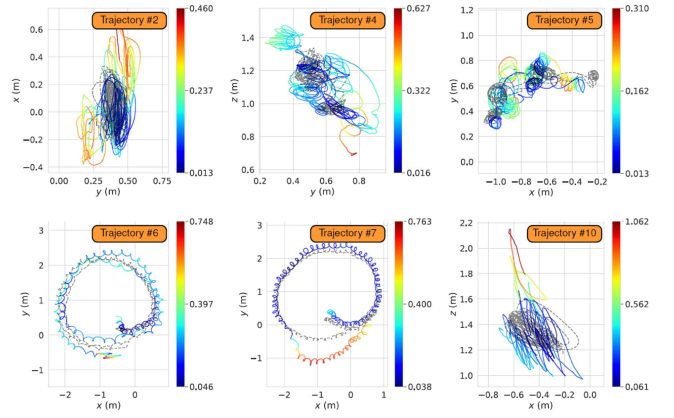


Fig. 4. Overview of representative testing trajectories used in our evaluation, aligned with the performance metrics reported in Table II. All trajectories were executed under fast and hostile motions, as reflected by the high angular velocities (10–33 rad/s) reported in Table II. Ground truth is shown in gray, with error mapping done for TCB-VIO.

$$+ {}^I_i \mathbf{R}^\top \int_{t_i}^{t_{i+1}} \int_{t_i}^s {}^I_i \mathbf{R}^\top {}^I \mathbf{a}_T(Z) dZ ds. \quad (15)$$

The velocity is then propagated incrementally at t_{i+1} integrating the last acceleration $\mathbf{a}_T(Z)$:

$${}^G \mathbf{v}_{I_{i+1}} = {}^G \mathbf{v}_{I_i} + {}^I_i \mathbf{R}^\top \int_{t_i}^{t_{i+1}} {}^I_i \mathbf{R}^\top {}^I \mathbf{a}_T(Z) dZ - {}^G \mathbf{g}(t_{i+1} - t_i). \quad (16)$$

The biases are also propagated by integrating the random walk noises $\mathbf{n}_G, \mathbf{n}_A$. They are zero-mean Gaussians:

$$\mathbf{b}_{G_{i+1}} = \mathbf{b}_{G_i} + \int_{t_i}^{t_{i+1}} \mathbf{n}_G(Z) dZ, \quad \mathbf{b}_{A_{i+1}} = \mathbf{b}_{A_i} + \int_{t_i}^{t_{i+1}} \mathbf{n}_A(Z) dZ. \quad (17)$$

The MSCKF backend of TCB-VIO linearizes the nonlinear IMU kinematic model and incrementally succeeds the state along with propagating the covariance, as explained in [8].

D. Update of 3D Point Features

At each measurement time-step t_m , a tracked corner feature, described in Section III-A3, is reprojected as a 3D bearing landmark, based on the following measurement model $\mathbf{h}(\cdot)$:

$$\begin{aligned} \mathbf{z}(t_m) &= \mathbf{h}(\mathbf{x}(t_m)) + \mathbf{n}(t_m) \\ &= \mathbf{T}_d(\mathbf{T}_p(\mathbf{T}_t(\mathbf{T}_r(\boldsymbol{\lambda}, \mathbf{K}), \mathbf{x}_{CG}), \zeta) + \mathbf{n}(t_m), \end{aligned} \quad (18)$$

where $\mathbf{n}(t_m)$ is zero-mean Gaussian noise, and $\mathbf{h}(\cdot)$ is a chain of functions following the standard projection pipeline [17]:

- $\mathbf{T}_r(\cdot)$ converts the landmark parameterization $\boldsymbol{\lambda}$ (inverse-depth or 3D position, with any context such as an anchor frame/extrinsics denoted by \mathbf{K}) into global coordinates.
- $\mathbf{T}_t(\cdot, \mathbf{x}_{CG})$ transforms the global point to the current camera frame using the camera pose in the global frame G , \mathbf{x}_{CG} .
- $\mathbf{T}_p(\cdot)$ orthogonally projects onto the normalized image plane.
- $\mathbf{T}_d(\cdot, \zeta)$ maps normalized coordinates to distorted pixel coordinates using intrinsic distortion parameters ζ .

TABLE II

ATE (M), FIRST BLOCK, AND RTE (M), SECOND BLOCK, FOR 10 TRAJECTORIES. EACH CELL SHOWS RMSE/MEDIAN. LOWEST OF ALL ERRORS ARE IN **BOLD**; LOWEST OF TCB-VIO, ROVIO, VINS-MONO ARE UNDERLINED. "F" = FAIL TO INITIALIZE, "PF" = PARTIAL FAILURE. NUMBERS AFTER EACH TRAJECTORY ID INDICATE THE ROUNDED AVERAGE ANGULAR VELOCITY, LENGTH, AND DURATION.

Traj.	1: 11 rad/s, 25 m, 19 s	2: 11 rad/s, 33 m, 26 s	3: 8 rad/s, 24 m, 29 s	4: 10 rad/s, 20 m, 24 s	5: 9 rad/s, 21 m, 26 s	6: 9 rad/s, 29 m, 24 s	7: 9 rad/s, 31 m, 31 s	8: 13 rad/s, 26 m, 16 s	9: 14 rad/s, 28 m, 13 s	10: 15 rad/s, 25 m, 12 s	
TCB-VIO	<u>0.199/0.202</u>	<u>0.223/0.156</u>	0.158/0.113	<u>0.218/0.192</u>	<u>0.140/0.121</u>	<u>0.255/0.234</u>	<u>0.315/0.168</u>	<u>0.210/0.136</u>	<u>0.601/0.469</u>	<u>0.379/0.013</u>	A
Ablation 1	0.354/0.207	10.056/5.666	<u>0.107/0.073</u>	F/F	PF/PF	F/F	PF/PF	F/F	PF/PF	F/F	T
Ablation 2	PF/PF	PF/PF	F/F	PF/PF	5.778/3.386	PF/PF	1.375/1.152	PF/PF	F/F	F/F	E
ROVIO	0.342/0.220	0.693/0.433	0.309/0.167	2.548/1.854	9.241/6.126	3.832/3.378	PF/PF	0.648/0.629	6.501/5.415	1.405/1.245	
VINS-Mono	0.113/0.083	0.563/0.416	2.798/1.203	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	
ORB-SLAM3	<u>0.091/0.053</u>	<u>0.151/0.096</u>	PF/PF	PF/PF	PF/PF	F/F	F/F	F/F	F/F	F/F	
TCB-VIO	<u>0.021/0.008</u>	<u>0.019/0.010</u>	<u>0.010/0.003</u>	<u>0.009/0.004</u>	<u>0.009/0.005</u>	<u>0.016/0.009</u>	<u>0.011/0.008</u>	<u>0.024/0.013</u>	<u>0.015/0.011</u>	<u>0.028/0.009</u>	R
Ablation 1	0.052/0.036	0.085/0.051	0.035/0.024	F/F	PF/PF	F/F	PF/PF	F/F	PF/PF	F/F	T
Ablation 2	PF/PF	PF/PF	F/F	PF/PF	0.013/0.007	PF/PF	0.013/0.010	PF/PF	F/F	F/F	E
ROVIO	0.085/0.062	0.078/0.067	0.054/0.040	0.051/0.042	0.079/0.056	0.076/0.058	PF/PF	0.089/0.083	0.149/0.118	0.112/0.105	
VINS-Mono	0.171/0.149	0.178/0.140	0.202/0.093	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	PF/PF	
ORB-SLAM3	0.091/0.072	0.098/0.071	PF/PF	F/F	PF/PF	F/F	F/F	F/F	F/F	F/F	

The predicted feature measurement is $\mathbf{z}_{\text{pred}}(t_m) = \mathbf{h}(\mathbf{x}(t_m))$. For the update, the residual is $r(t_m) = \mathbf{z}_{\text{meas}}(t_m) - \mathbf{h}(\mathbf{x}(t_m))$. For MSCKF features, residuals accumulated over the clone window are projected into the left null-space of the landmark Jacobian. This operation removes the landmarks analytically and allows for updating only the cloned poses.

IV. EXPERIMENTAL RESULTS

This section details the experiments along with discussions.

A. Experimental Setup

TCB-VIO operates with the SCAMP-5 FPSP, processing at 250. IMU measurements, at 400 Hz, and grayscale images for benchmarking are sourced from an Intel D435i RealSense. These two sensors are rigidly attached to a fixture. (See [33]) Benchmarking methods are ROVIO [1], VINS-Mono [2], and ORB-SLAM3 [3] using grayscale and IMU data. TCB-VIO cannot be evaluated on standard benchmark datasets due to on-sensor computation. For indoor experiments, the ground truth is provided by a Vicon motion capture system running at 300 Hz. To synchronize and calibrate the FPSP and IMU, Kalibr [34] was used. For metrics, the absolute trajectory error (ATE) and relative trajectory error (RTE) are stated as a median and root mean squared error (RMSE) [35]. 13 trajectories were used for benchmarking (10 indoors and 3 outdoors). The trajectories are created by rapidly shaking the fixture in all directions to simulate violent motions (See Fig. 4). Trajectories #1–3 are mostly rotational. Trajectories #4–7 are mostly translational. Trajectories #8–10 are also mostly rotational, but more challenging than trajectories #1–3, based on their average angular velocities (see Table II, top row). Trajectories #11–13 (see Table V) are similar to trajectories #1–3, but occurring outdoors. The length and duration of each trajectory are also provided in the tables.

B. Parameters of TCB-VIO

Table I lists the key parameters used in the experiments. An exhaustive search over a wide range tuned these. Graphs of the parameter search are on the project webpage [33]. While N_r , the size of \mathbf{x}_{BIT} in Eq. (7), may grow over time, the number of landmarks used for updating SLAM and MSCKF is limited, denoted with $N_{\text{SLAM, update}}$ and $N_{\text{MSCKF, update}}$.

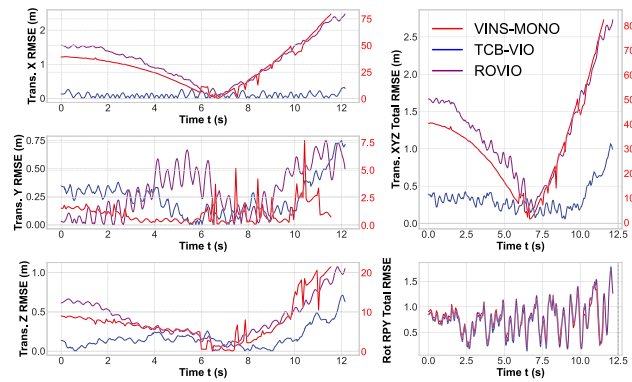


Fig. 5. Estimated RMSE for trajectory #10: (left) translation for X, Y, and Z axes, (top right) total translation, and (bottom right) total rotation. VINS-Mono has its own axes, in red, due to large error ranges. While rotational RMSE values for all three methods are close, TCB-VIO maintains the lowest overall translational RMSE.

C. Comparisons With ROVIO and VINS-Mono

As shown in Table II, TCB-VIO demonstrates superior robustness against fast and hostile motions, achieving the lowest ATE and RTE across all experiments, compared with ROVIO and VINS-Mono. VINS-Mono suffers from partial failures. ROVIO exhibits reasonable robustness, avoiding large errors throughout all experiments. However, its ATE and RTE values are still up to 4 to 10 times higher than those of TCB-VIO. Error values for partial failures are not reported, as they are at least 10× larger than the lowest error.

As an example, Fig. 5 shows the RMSE of trajectory #10 along all axes. TCB-VIO achieves the lowest total translational RMSE (Fig. 5 top-right). ROVIO diverges over time, especially in the X and Z directions, with errors exceeding 2.7 meters. VINS-Mono has its own axes for RMSE in red on the right of subplots because of its much larger error ranges, particularly in the X direction, resulting in higher errors compared to TCB-VIO, with errors exceeding 30 meters in the X. These trends align with the quantitative data in Table II, further emphasizing TCB-VIO’s robustness during fast and hostile motions. In terms of rotational RMSE (Fig. 5 bottom-right), all methods perform similarly, but TCB-VIO maintains the lowest overall error. This highlights its ability to handle rapid rotational dynamics effectively, unlike ROVIO and VINS-Mono, which experience occasional spikes due to tracking loss.

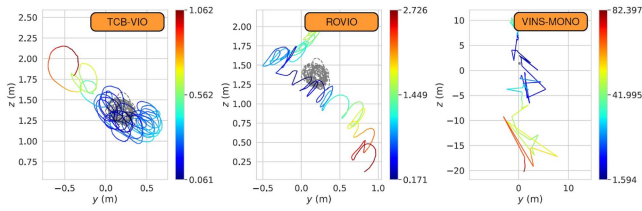


Fig. 6. Error mapping for trajectory #10, comparing TCB-VIO (left), ROVIO (middle), and VINS-Mono (right). TCB-VIO maintains low error, relative to ROVIO and VINS-MONO, closely following the ground-truth path (in gray) with minimal deviations. ROVIO shows significant drift, while VINS-Mono diverges. Note that the colorbar scales differ across the graphs.

Fig. 6 shows the tracking error, color-coded, for trajectory #10. Note that the scale of the error bars is different. The ground truth is shown in gray. TCB-VIO demonstrates lower error mapping relative to ROVIO and VINS-Mono algorithms. In contrast, both ROVIO and VINS-MONO exhibit relatively large error clusters, particularly at sharp turns and high dynamic regions, as evident from the red and orange segments on the map. Notably, VINS-MONO diverged to such an extent that it jittered and failed to maintain local consistency.

D. Trajectory Failures Present in ORB-SLAM3

As shown in Table II, while the metrics of TCB-VIO and ORB-SLAM3 for trajectories #1 and #2 are close, ORB-SLAM3 fails in the rest of the trajectories. In the table, an “F” indicates a failure to initialize, and “PF” is a partial failure. Fig. 7-(top row) corresponds to trajectory #8, where ORB-SLAM3 fails. This type of failure is characterized by ORB-SLAM3’s inability to initialize in certain regions, specifically those marked in pink. In the short non-pink regions, ORB-SLAM3 performs well, demonstrating its potential under favorable conditions. However, during pink intervals, ORB-SLAM3 crashes, losing tracking completely. In contrast, TCB-VIO, while deviating slightly from the ground truth in some areas, maintains reasonable tracking accuracy and does not suffer from the catastrophic failures observed in ORB-SLAM3.

Fig. 7-(bottom row) illustrates a PF-type failure, as seen in trajectory #3 from Table II. In this particular trajectory, compared to trajectories #1-2, the scene was designed to contain fewer trackable features, making successful tracking possible primarily in high-exposure conditions where the binarized image remained stable. Under these constraints, ORB-SLAM3 repeatedly attempts to reinitialize within the pink regions but diverges rapidly, leading to large ATE and RTE values. Although the pink regions are less prominent compared to F-type failures, their impact is still significant due to the rapid loss of tracking. The limited feature availability, coupled with faster motions, caused ORB-SLAM3’s failure to recover. In contrast, TCB-VIO demonstrates its robustness by maintaining a consistent trajectory, mitigating errors, and maintaining reliable performance throughout the trajectory, despite the challenging conditions.

E. Discussions

This section discusses five aspects of TCB-VIO’s performance and design, including its 1) time complexity, 2) computational load, 3) the impact of in-state (SLAM) features, 4) outdoor performance, and 5) ablation studies.

TABLE III
RUNTIME COMPARISON ON FPSP/CPU/GPU

Process	FPSP (SCAMP-5)	CPU (i7- 12650H)	GPU (RTX 3070)
Edge Detection	9.0 μ s	147.818 μ s	152.63 μ s
Corner Detection	292.0 μ s	76.671 μ s	20.23 μ s

TABLE IV
AVERAGE PER-FRAME ENERGY CONSUMPTION ON FPSP/CPU/GPU

Process	FPSP (SCAMP-5)	CPU (i7- 12650H)	GPU (RTX 3070)
Edge Detection	0.019 mJ	994.23 mJ	6.46 mJ
Corner Detection	1.15 mJ	987.22 mJ	1.65 mJ

1) *Time Complexity Analysis*: To quantify the computational benefits of running TCB-VIO on the FPSP, the execution times of edge and corner detection on the FPSP were compared against CPU and GPU. For the CPU/GPU measurements, 1000 images were captured by the SCAMP-5 FPSP and processed to perform edge and corner detection. Average per-frame results were reported using system utility tools. For FPSP, execution times were calculated from the number of instructions in the code and the execution time per instruction.

As shown in Table III, for edge detection, the SCAMP-5 FPSP achieves 16.4 \times speedup over the CPU and 17.0 \times over the GPU. For feature extraction (1000 keypoints), SCAMP-5 FPSP requires 292 μ s, while CPU and GPU require 76.67 μ s and 20.23 μ s. On SCAMP, the corner detection runtime is constant regardless of keypoint count since all 256×256 processing elements execute neighbour comparisons in parallel. Thus, the cost per frame is fixed. In contrast, CPU runtimes scale with keypoints. SCAMP-5 may appear slower in Table III, but this reflects its fixed cost model rather than inefficiency. SCAMP-5 executes operations at 5-10 MHz, far below conventional processors (CPU \approx 2.7 GHz, GPU \approx 400 MHz). This frequency gap explains SCAMP’s longer absolute runtimes despite its parallel architecture. Future FPSP designs with higher clock rates would narrow this gap while preserving parallelism. Note that the intensity image transfer time, significant for CPU/GPU, was not included. By avoiding intensity image transfer, host processor saturation is prevented.

2) *Computational Load Analysis*: The efficiency of TCB-VIO was evaluated in terms of energy consumption. Edge and corner detection were executed directly on-sensor by the SCAMP-5 FPSP using its dedicated processing elements, resulting in per-frame energy consumption of only tens of millijoules (mJ). In contrast, substantially higher CPU/GPU utilization was required when the same operations were offloaded to conventional processors. Per-frame energy consumption was measured for edge and corner detection on the FPSP, CPU, and GPU. For CPU/GPU evaluation, 1000 images were captured by the SCAMP-5 FPSP while the sensor was moved along a known trajectory. For the FPSP evaluation, the same trajectory was repeated twice, once with edge detection and once with corner detection enabled. Energy measurements were obtained using a USB power meter for the SCAMP-5 FPSP, while Intel RAPL and Nvidia SMI were employed to monitor CPU and GPU energy consumption, respectively.

Table IV reports the per-frame energy consumption for the Sobel edge detection (kernel size $k = 3$) and BIT corner detection, averaged on 1000 256×256 images. For edge detection, SCAMP-5 consumes only \sim 0.019 mJ per frame,

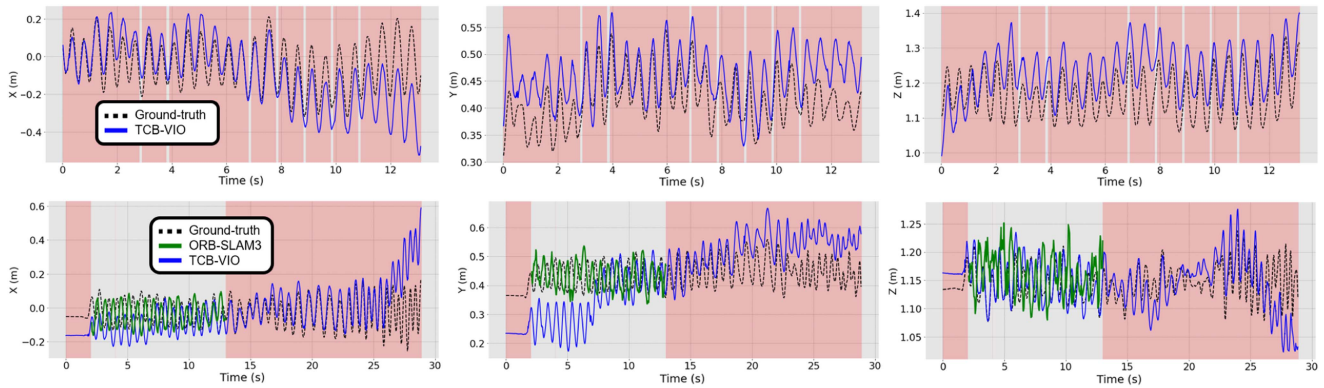


Fig. 7. (Top row) Trajectory #8 (F-type failure) showing failure during pink regions. TCB-VIO tracks consistently despite occasional deviations from ground-truth. (Bottom row) Trajectory #3 (PF-type failure) illustrating ORB-SLAM3 reinitializing but quickly diverging, highlighted in pink, leading to ATE and RTE errors.

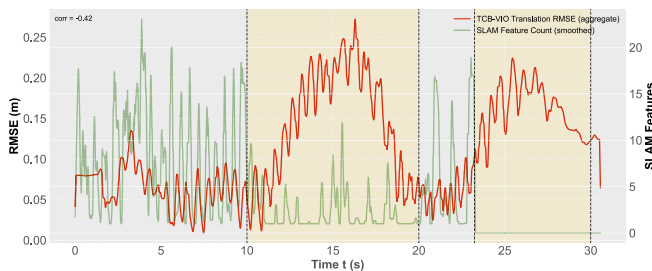


Fig. 8. Correlation between in-state feature and translational RMSE for trajectory #2. When feature count is <10 (yellow regions), TCB-VIO has higher translational RMSE.

whereas the CPU/GPU requires 994.23 mJ/6.46 mJ per frame. For corner detection, the FPSP consumes 1.15 mJ per frame, compared to 987.22 mJ/1.65 mJ on the CPU/GPU. This computational load analysis confirms that SCAMP-5 reduces energy consumption relative to a CPU/GPU-based pipeline, underscoring the lightweight and embedded nature of on-sensor SIMD computation.

3) *Impact of In-State Features*: TCB-VIO consistently achieves low ATE and RTE, but deviations from ground truth occur under specific conditions. Fig. 8 shows RMSE translation error (deviation) and in-state feature count, for trajectory #2, overlaid. In yellow regions, due to poor texture, feature count is low and RMSE is high, the converse is true in green regions. The negative correlation ($\text{corr} = -0.42$) shows that having persistent features helps constrain long-term drift. This trend was seen in other trajectories as well.

4) *Outdoor Trajectories*: In addition to the 10 indoor experiments, three additional experiments were conducted outdoors. Lacking ground truth in the outdoor environment, the performance was evaluated by starting and ending the trajectories from a single fixture, to ensure the start and end poses were identical. The VIO performance was then measured by the endpoint drift (distance) between the estimated start and end poses. As indicated in Table V, TCB-VIO exhibits the lowest error compared to the other algorithms. Videos of these experiments are on the project website [33].

5) *Ablation Studies*: To evaluate the design of the algorithm, two ablation studies were conducted, as summarized in Table II. The corresponding ATE and RTE results are reported in the

TABLE V

START-TO-END ERROR (M) FOR OUTDOOR EXPERIMENTS, PERFORMED IN SUN/SHADE CONTRAST. LOWEST ERROR IS IN **BOLD**. “F” AND “PF” ARE FOR ‘FAIL TO INITIALIZE’ AND ‘PARTIAL FAILURE’. THE LAST COLUMN SHOWS THE AVERAGE ANGULAR VELOCITY AND TRAJECTORY LENGTH/DURATION. THE VELOCITY AND LENGTH HAVE BEEN COMPUTED USING THE TCB-VIO POSE ESTIMATES.

Traj.	TCB-VIO	ROVIO	VINS-Mono	ORB-SLAM3	Description of Trajectory
11	1.713	21.6 (PF)	F	F	7 rad/s, 21 m, 20.3 s
12	0.130	11.8 (PF)	695.5 (PF)	F	9 rad/s, 47 m, 49.9 s
13	2.736	5.1	F	F	8 rad/s, 31 m, 15.9 s

rows labelled Ablation 1 and Ablation 2. In Ablation 1, edge images were transmitted to the host, and Shi–Tomasi features (extracted from the edges) were used for KLT tracking instead of BIT features. The only difference from the TCB-VIO pipeline is the type of features. In Ablation 2, the feathering block (Section III-A2) was removed from the TCB-VIO pipeline.

Table II shows that across all trajectories, TCB-VIO achieves lower ATE and RTE than the ablation studies, with the exception of trajectory #3, where its performance is nearly identical to Ablation 1. For instance, in trajectory #1, the median ATE decreases from 0.207 m (Ablation 1) to 0.202 m (TCB-VIO). More importantly, in challenging sequences such as trajectory #9, Ablation 1 fails to track, while Ablation 2 fails to initialize. These results highlight that both the choice of features and the inclusion of the feathering process are essential for the effectiveness of the TCB-VIO pipeline.

The established benchmark algorithms are known to perform well under standard and low-speed conditions. Our evaluation is specifically designed to highlight performance gains beyond the operational limits of conventional methods, which is the main contribution of TCB-VIO.

V. CONCLUSION

In this work, we presented TCB-VIO, a tightly-coupled 6-DoF VIO framework using the on-sensor compute capabilities of SCAMP5 FPSP. By integrating visual and inertial measurements at 250 FPS and 400 Hz, respectively, and processing binary edge images and feature maps with a novel binary-enhanced Kanade-Lucas-Tomasi (KLT) tracker, TCB-VIO achieves robust and accurate tracking under fast motions. The evaluations demonstrated that TCB-VIO outperforms state-of-the-art

methods, including ROVIO, VINS-Mono, and ORB-SLAM3 in terms of both translational and rotational accuracy. While ROVIO exhibits drift over time, and VINS-Mono fails to maintain tracking under extreme conditions, TCB-VIO maintains accurate trajectory estimation. This highlights the importance of high frame-rate synchronization between visual and inertial data, as well as the computational advantages provided by FPSPs. The proposed framework addresses the trade-off between tracking performance and latency by combining high-frequency visual-inertial data processing with efficient on-sensor computation. This work demonstrates the potential of FPSPs in advancing the capabilities of mobile robotic systems, paving the way for more reliable visual-inertial odometry. The proposed method has partial on-sensor computation due to hardware limitations. Future hardware with more flexible compute units could enable all computations to be migrated entirely to the FPSP. In particular, while the feathering process is highly parallelizable, it demands significant bandwidth to transfer the results to the host. An optimal solution would shift this computational load from the host to the FPSP. Another approach is to utilize lower-dimensional and more computationally efficient features for tracking; the corner and edge features in the binary-enhanced KLT algorithm in this work show promise for further investigation.

ACKNOWLEDGMENT

Authors would like to thank Piotr Dudek, Stephen J. Carey, and Jianing Chen at the University of Manchester for providing access to SCAMP-5, and Riku Murai for the feature source code.

REFERENCES

- [1] M. Bloesch, M. Burri, S. Omari, M. Hutter, and R. Siegwart, "Iterated extended Kalman filter based visual-inertial odometry using direct photometric feedback," *Int. J. Robot. Res.*, vol. 36, no. 10, pp. 1053–1072, 2017.
- [2] T. Qin, P. Li, and S. Shen, "VINS-Mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Trans. Robot.*, vol. 34, no. 4, pp. 1004–1020, Aug. 2018.
- [3] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. M. Montiel, and J. D. Tardós, "ORB-SLAM3: An accurate open-source library for visual, visual-inertial, and multimap SLAM," *IEEE Trans. Robot.*, vol. 37, no. 6, pp. 1874–1890, Dec. 2021.
- [4] A. Zarandy, *Focal-Plane Sensor-Processor Chips*. New York, NY, USA: Springer, 2011.
- [5] P. Dudek et al., "Sensor-level computer vision with pixel processor arrays for agile robots," *Sci. Robot.*, vol. 7, Jun. 2022, Art. no. eabl7755.
- [6] R. S. Amant et al., "General-purpose code acceleration with limited-precision analog computation," in *Proc. Int. Symp. Comput. Archit*, 2014, pp. 505–516.
- [7] D. Scaramuzza and Z. Zhang, *Visual-Inertial Odometry of Aerial Robots*. Berlin, Germany: Springer, 2020, pp. 1–9.
- [8] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint Kalman filter for vision-aided inertial navigation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2007, pp. 3565–3572.
- [9] S. Leutenegger, P. Furgale, V. Rabaud, M. Chli, K. Konolige, and R. Siegwart, "Keyframe-based visual-inertial SLAM using nonlinear optimization," *Robotics, Sci. Syst. (RSS)*, June 24–28, 2013.
- [10] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," *Georgia Inst. Technol.*, Atlanta, GA, USA, Tech. Rep. GT-RIM-CP&R-2012-002 2012.
- [11] J. A. Hesch, D. G. Kottas, S. L. Bowman, and S. I. Roumeliotis, "Camera-IMU-based localization: Observability analysis and consistency improvement," *Int. J. Robot. Res.*, vol. 33, no. 1, pp. 182–201, 2014.
- [12] S. Lynen, M. W. Achtelik, S. Weiss, M. Chli, and R. Siegwart, "A robust and modular multi-sensor fusion approach applied to MAV navigation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2013, pp. 3923–3929.
- [13] J. Delmerico and D. Scaramuzza, "A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 2502–2509.
- [14] R. Murai, S. Saeedi, and P. H. J. Kelly, "High-frame rate homography and visual odometry by tracking binary features from the focal plane," *Auton. Robots*, vol. 47, pp. 1579–1592, Jul. 2023.
- [15] L. Bose, J. Chen, S. J. Carey, P. Dudek, and W. Mayol-Cuevas, "Visual odometry for pixel processor arrays," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 4614–4622.
- [16] M. Lisondra, J. Kim, R. Murai, K. Zareinia, and S. Saeedi, "Visual inertial odometry using focal plane binary features (BIT-VIO)," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 1661–1668.
- [17] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVINS: A research platform for visual-inertial estimation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 4666–4672.
- [18] S. J. Carey, A. Lopich, D. R. Barr, B. Wang, and P. Dudek, "A 100,000 FPS vision sensor with embedded 535GOPS/W 256 × 256 SIMD processor array," in *Proc. Symp. VLSI Circuits*, 2013, pp. 182–183.
- [19] C. Greatwood et al., "Tracking control of a UAV with a parallel visual processor," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 4248–4254.
- [20] A. McConville et al., "Visual odometry using pixel processor arrays for unmanned aerial systems in GPS denied environments," *Front. Robot. AI*, vol. 7, 2020, Art. no. 126.
- [21] Y. Liu, J. Chen, L. Bose, P. Dudek, and W. Mayol-Cuevas, "Direct servo control from in-sensor CNN inference with a pixel processor array," *2021 IEEE International Conference on Robotics and Automation (ICRA) workshop: On and Near-sensor Vision Processing, from Photons to Applications*, 2021, arXiv:2106.07561.
- [22] H. M. So, L. Bose, P. Dudek, and G. Wetzstein, "PixelRNN: In-pixel recurrent neural networks for end-to-end-optimized perception with neural sensors," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 25233–25244.
- [23] J. Chen, Y. Liu, S. J. Carey, and P. Dudek, "Proximity estimation using vision features computed on sensor," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 2689–2695.
- [24] H. Castillo-Elizalde, Y. Liu, L. Bose, and W. Mayol-Cuevas, "Weighted node mapping and localisation on a pixel processor array," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 6702–6708.
- [25] A. R. Vidal, H. Rebecq, T. Horstschaefer, and D. Scaramuzza, "Ultimate SLAM? Combining events, images, and IMU for robust visual SLAM in HDR and high-speed scenarios," *IEEE Robot. Automat. Lett.*, vol. 3, no. 2, pp. 994–1001, Apr. 2018.
- [26] J. Chen, S. J. Carey, and P. Dudek, "Feature extraction using a portable vision system," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. Workshop*, 2017, vol. 2, p. 1–2.
- [27] J. Chen, S. J. Carey, and P. Dudek, "Scamp5D vision system and development framework," in *Proc. 12th Int. Conf. Distrib. Smart Cameras*, 2018, pp. 1–2.
- [28] B. D. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision," in *Proc. 7th Int. Joint Conf. Artif. Intell.*, 1981, vol. 2, pp. 674–679.
- [29] C. Tomasi and T. Kanade, "Detection and tracking of point," *Int. J. Comput. Vis.*, vol. 9, no. 137–154, 1991, Art. no. 3.
- [30] J. Shi et al., "Good features to track," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 1994, pp. 593–600.
- [31] M. Li and A. I. Mourikis, "Optimization-based estimator design for vision-aided inertial navigation," *Robot., Sci. Syst.*, vol. 8, 2013, pp. 241–248.
- [32] N. Trawny and S. I. Roumeliotis, "Indirect Kalman filter for 3D attitude estimation," Dept. Comput. Sci. Eng., Univ. Minnesota, Minneapolis, MN, USA, Tech. Rep. TR-2005-002. 2005.
- [33] M. Lisondra, "TCB-VIO," 2025. Accessed: Apr. 4, 2025. [Online]. Available: <https://sites.google.com/view/tcb-vio>
- [34] P. Furgale, J. Rehder, and R. Siegwart, "Unified temporal and spatial calibration for multi-sensor systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Tokyo, Japan, 2013, pp. 1280–1286.
- [35] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of RGB-D SLAM systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2012, pp. 573–580.