

Balancing the Budget: Feature Selection and Tracking for Multi-Camera Visual-Inertial Odometry

Lintong Zhang, David Wisht, Marco Camurri, Maurice Fallon

Abstract—We present a multi-camera visual-inertial odometry system based on factor graph optimization which estimates motion by using all cameras simultaneously while retaining a fixed overall feature budget. We focus on motion tracking in challenging environments such as in narrow corridors and dark spaces with aggressive motions and abrupt lighting changes. These scenarios cause traditional monocular or stereo odometry to fail. While tracking motion across extra cameras should theoretically prevent failures, it causes additional complexity and computational burden. To overcome these challenges, we introduce two novel methods to improve multi-camera feature tracking. First, instead of tracking features separately in each camera, we track features continuously as they move from one camera to another. This increases accuracy and achieves a more compact factor graph representation. Second, we select a fixed budget of tracked features which are spread across the cameras to ensure that the limited computational budget is never exceeded. We have found that using a smaller set of informative features can maintain the same tracking accuracy while reducing back-end optimization time. Our proposed method was extensively tested using a hardware-synchronized device containing an IMU and four cameras (a front stereo pair and two lateral) in scenarios including an underground mine, large open spaces, and building interiors with narrow stairs and corridors. Compared to stereo-only state-of-the-art VIO methods, our approach reduces the drift rate (RPE) by up to 80 % in translation and 39 % in rotation.

I. INTRODUCTION

State estimation is a fundamental capability required for autonomous robot navigation in real-world scenarios. Motion tracking using cameras is very popular due to their low weight, small form factor, and hardware cost. Specifically, Visual-Inertial Odometry (VIO) methods, which fuse feature tracking from a camera with estimation from an Inertial Measurement Unit (IMU), have now become the standard for odometry on Micro Aerial Vehicles (MAVs) [1]. These systems can also be deployed on diverse platforms, such as handheld rigs, quadrupeds, and vehicles with applications ranging from indoor and outdoor navigation to underground exploration [2] (Fig. 1).

When performing VIO, cameras can track scene information at the frame rate while the IMU can provide high frequency motion estimates between keyframes. Although VIO frameworks are generally effective, they are prone to failure in degenerate scenarios such as over/under exposure, textureless surfaces, and when there is aggressive motion. By using a minimal setup composed of just a monocular camera, the system has a single point of failure. Even if IMU measurements can be integrated, the system quickly diverges within 1-2 seconds of a feature tracking failure.

The authors are with the Oxford Robotics Institute, Department of Engineering Science, University of Oxford, UK {lintong, davidw, mcamurri, mfallon}@robots.ox.ac.uk




Fig. 1. We tested our multi-camera odometry algorithm with data from a handheld sensor rig in New College, Oxford (*left*) and the ANYmal quadruped [3] during preparation for the DARPA SubT Challenge [2] (*right*).

Fusing information from additional cameras can greatly improve robustness and estimation accuracy at the cost of extra complexity and computational burden. The most obvious benefit of Multi-Camera VIO (MC-VIO) is redundancy, e.g., if one of the cameras suffers from a sudden drop in tracked features, the other cameras can still estimate motion. Furthermore, if the cameras are arranged to have overlapping Fields of View (FoV), a particular feature could potentially be tracked across the cameras for as long as it is visible in any camera. Multi-camera solutions have been presented in the past [4], [5], [6], however, no proposal has studied cross camera feature tracking.

In this paper, we explore how to effectively track features across all the cameras and how to select the best subset of features to keep the computational time bounded.

A. Contribution

The main contributions of this work are the following:

- A novel factor graph formulation that tightly fuses feature tracks from any number of stereo and monocular cameras, along with IMU measurements, in a single consistent optimization process;
- A simple but effective method to track features across cameras with overlapping FoVs to reduce duplicate landmark tracking and improve accuracy;
- A submatrix feature selection (SFS) scheme to select the best landmarks for optimization with a fixed feature budget. This bounds computational time and achieves the same accuracy compared to using all available features;
- Extensive experimental evaluation across a range of scenarios demonstrating superior robustness, particularly when VIO with an individual camera fails.

The proposed algorithm, VILENS-MC, builds upon our previous VILENS estimation system [7], [8], by fusing multiple cameras and improving front-end feature processing.

II. RELATED WORK

There has been extensive research into monocular and stereo camera VIO. The methods can be categorized into either optimization or filter based approaches. In recent years, optimization based methods such as VINS-Mono [9], ORB-SLAM3 [10], and OKVIS [11] have become popular due to their ability to optimize a trajectory of poses. This is in contrast with methods based on Extended Kalman Filters (EKF) which marginalize all previous states. However, some modern filtering methods have achieved competitive performance via stochastic cloning [12]. For a more detailed review, readers can refer to comparisons in [13] and [14].

MC-VIO has been less extensively studied because of hardware complexity and potentially high computational requirements. Recent work from Jaekel et al. [4] fused two pairs of stereo cameras and accounted for the uncertainty in the extrinsics in both their front- and back-end systems to improve performance. However, their experiments had limited real-world scenarios and no ground truth comparisons.

A VIO system by Liu et al. [15], with three pairs of stereo cameras, estimated poses by minimising photometric errors. This work was developed and tested in autonomous driving scenarios where the authors showed the benefit of a multi-camera system when driving at night. However, since a constant velocity motion model was used, their method is not suitable for handheld sensing or drones with highly dynamic motions.

A unique omnidirectional setup was proposed by Seok et al. [5] with four very large FoV cameras. The four overlapping image regions were treated as four stereo cameras, but the system did not fully take advantage of the camera setup to track features across the camera pairs.

With a focus on aerial robotics, Müller et al. [16] used two pairs of stereo cameras on an MAV (looking up and down). The images from their wide-angle cameras were split in half and fed into four separate stereo VIO systems running on a Field Programmable Gate Array (FPGA). The VIO outputs and an IMU were then fused together using an EKF in a loosely coupled fashion. However, the experimental results were limited to an indoor office building.

Kuo et al. [17] introduced a more general design for multi-camera SLAM, which involved an adaptive initialization scheme, keyframe selection, and map management. Their system was based on SVO [18] and they claimed their approach required minimal parameter tuning. However, their real-world experimental results showed little improvement in accuracy when using a multi-camera setup.

Meanwhile using a Multi-State Constraint Kalman Filter (MSCKF), Eickenhoff et al. [6] fused up to six cameras with multiple IMUs. The MSCKF's low computational requirements allowed for asynchronous camera keyframe processing and real-time operation. They presented detailed simulation results, however, real-world experiments were limited to a small lab environment.

In contrast, we introduce an MC-VIO framework to fuse an arbitrary number of stereo and monocular cameras, with a focus on multi-camera feature selection and tracking.

We are particularly motivated to address common challenges and failure cases in mono and stereo VIO, which have




Fig. 2. VILENS-MC is evaluated using a custom built handheld device, Handheld ALphasense Ouster (HALO). See details in VII-A.

not been a focus of recent study. Hence, our system is designed and tested in a variety of challenging environments, covering fast and abrupt motions, severe illumination changes, indoor/outdoor scenes, and dark underground environments.

III. PROBLEM STATEMENT

We aim to estimate the position, orientation, and velocity of a mobile platform equipped with an IMU and multiple hardware-synchronized cameras. Fig. 2 shows the multi-camera device which is used in Section VIII for the majority of the experimental results.

The relevant reference frames are as follows: the earth-fixed world frame \mathcal{W} , the platform-fixed base frame \mathcal{B} , the IMU frame \mathcal{I} , and n individual camera frames \mathcal{C}_c where $c \in \{1, \dots, n\}$.

Unless otherwise specified, the position ${}_{\mathcal{W}}\mathbf{p}_{\mathcal{WB}}$ and orientation $\mathbf{R}_{\mathcal{WB}}$ of the base (with ${}_{\mathcal{W}}\mathbf{T}_{\mathcal{WB}} \in \text{SE}(3)$ as the corresponding homogeneous transform) are expressed in world coordinates; the linear velocity is in the base frame ${}_{\mathcal{B}}\mathbf{v}_{\mathcal{WB}}$, and IMU biases ${}_{\mathcal{I}}\mathbf{b}^g$, ${}_{\mathcal{I}}\mathbf{b}^a$ are expressed in the IMU frame.

A. State and Measurements Definition

The state of the sensor rig at time t_i is defined as follows,

$$\mathbf{x}_i \triangleq [\mathbf{R}_i, \mathbf{p}_i, \mathbf{v}_i, \mathbf{b}_i^g, \mathbf{b}_i^a] \in \text{SO}(3) \times \mathbb{R}^{12} \quad (1)$$

where: \mathbf{R}_i is the orientation, \mathbf{p}_i is the position, \mathbf{v}_i is the linear velocity, and \mathbf{b}_i^g , \mathbf{b}_i^a are, respectively, the usual IMU gyroscope and accelerometer biases. In addition to the states, we track point landmarks \mathbf{m}_ℓ as triangulated visual features.

The objective is to estimate the optimized trajectory \mathcal{X}_k of all states \mathbf{x}_i and landmarks \mathbf{m}_ℓ visible up to the current time t_k within a fixed lag smoothing window.

The measurements from the set of multiple cameras \mathcal{C} and an IMU \mathcal{I} are hardware synchronized but received at different frequencies (cameras at 30 Hz and IMU at 400 Hz). We define \mathcal{Z}_k as the full set of measurements received within the smoothing window.

B. Maximum-a-Posteriori Estimation

We maximize the likelihood of the measurements \mathcal{Z}_k , given the history of states \mathcal{X}_k ,

$$\mathcal{X}_k^* = \arg \max_{\mathcal{X}_k} p(\mathcal{X}_k | \mathcal{Z}_k) \propto p(\mathcal{X}_0) p(\mathcal{Z}_k | \mathcal{X}_k) \quad (2)$$

The measurements are formulated as conditionally independent and corrupted by white Gaussian noise. Therefore, Eq. (2) can be expressed as the following least squares minimization [19],




Fig. 3. Sliding-window factor graph structure, showing prior, visual, and preintegrated IMU factors. Landmarks can be tracked across both stereo and mono cameras to create longer feature tracks and improve accuracy.

$$\mathcal{X}_k^* = \arg \min_{\mathcal{X}_k} \|\mathbf{r}_0\|_{\Sigma_0}^2 + \sum_{i \in K_k} \left(\|\mathbf{r}_{\mathcal{I}_{ij}}\|_{\Sigma_{\mathcal{I}_{ij}}}^2 + \sum_{\ell \in M_i} \|\mathbf{r}_{\mathbf{x}_i, \mathbf{m}_\ell}\|_{\Sigma_{\mathbf{x}_i, \mathbf{m}_\ell}}^2 \right) \quad (3)$$

where \mathcal{I}_{ij} are the IMU measurements between t_i and t_j and K_k are all the keyframes indices in the sliding window up to t_k . Each term is the residual associated to a factor type, weighted by the inverse of its covariance matrix. The residuals include a prior, IMU factors, and visual landmarks.

IV. FACTOR GRAPH FORMULATION

In this section, we describe the measurements, residuals, and covariances of the IMU and visual factors shown in Fig. 3.

A. Preintegrated IMU Factors

We follow the standard manner of IMU measurement preintegration from [1] to constrain the pose, velocity, and biases between two consecutive nodes of the graph, providing high frequency state updates between nodes. The residual has the following form:

$$\mathbf{r}_{\mathcal{I}_{ij}} = \left[\mathbf{r}_{\Delta \mathbf{R}_{ij}}^T, \mathbf{r}_{\Delta \mathbf{v}_{ij}}^T, \mathbf{r}_{\Delta \mathbf{p}_{ij}}^T, \mathbf{r}_{\mathbf{b}_{ij}^a}, \mathbf{r}_{\mathbf{b}_{ij}^g} \right] \quad (4)$$

For a detailed definition of the above residuals, see [1].

B. Stereo and Mono Landmark Factors

Next we define the visual factors. In the below equations, $\mathbf{m}_\ell \in \mathbb{R}^3$ is a visual landmark in Euclidean space, where $\pi : \text{SE}(3) \times \mathbb{R}^3 \mapsto \mathbb{R}^2$ is a projection function of a landmark into the image plane, given the platform pose \mathbf{T}_i (for simplicity, we omit the fixed transform between base and camera). The coordinates of \mathbf{m}_ℓ on the image plane (orange/green circles in Fig. 4) are denoted with $(u_\ell, v_\ell) \in \mathbb{R}^2$. The residual at state \mathbf{x}_i for landmark \mathbf{m}_ℓ is [7]:

$$\mathbf{r}_{\mathbf{x}_i, \mathbf{m}_\ell} = \begin{pmatrix} \pi_u^L(\mathbf{T}_i, \mathbf{m}_\ell) - u_{i,\ell}^L \\ \pi_u^R(\mathbf{T}_i, \mathbf{m}_\ell) - u_{i,\ell}^R \\ \pi_v(\mathbf{T}_i, \mathbf{m}_\ell) - v_{i,\ell} \end{pmatrix} \quad (5)$$

where (u^L, v) , (u^R, v) are the pixel locations of the detected landmark. $\Sigma_{\mathbf{m}}$ is computed using an uncertainty of 0.25 pixels. We also account for lens distortion using the covariance warping method from [8]. For monocular landmarks, only the first and the last rows of Eq. (5) are used.




Fig. 4. VILENS-MC takes advantage of any overlapping image regions (purple rectangles) in a multi-camera setup to track features across cameras. This increases feature track length and avoids tracking the same feature independently in different cameras. The arrows indicate features being tracked from one image to another.

For landmarks detected by a stereo camera pair, the landmark location is initialized using stereo triangulation. For landmarks detected in monocular cameras, we triangulate the feature location over the last N_{obs} frames using the Direct Linear Transform (DLT) algorithm from [20].

V. CROSS CAMERA FEATURE TRACKING

Most multi-camera systems do not take advantage of the overlapping FoVs of their setups, which produce common image regions that allow for cross-camera feature matching.

We propose a simple but effective method to continuously track features across different cameras which we call Cross Camera Feature Tracking (CCFT). CCFT can avoid the optimization of redundant landmarks by tracking the same feature in different cameras simultaneously.

As shown in Fig. 3, we continually add constraints between states \mathbf{x}_i and landmarks \mathbf{m}_ℓ even as they are tracked across cameras. For example, landmark \mathbf{m}_1 is first tracked in a stereo camera at $\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$ then tracked in a monocular camera at \mathbf{x}_3 .

A common approach to perform this feature matching would be via feature descriptors. However, this would be computationally expensive for real-time, high-frequency VIO on a mobile platform. For example, the combination of feature extraction and stereo matching in ORB-SLAM2 [21] takes ~ 24.8 ms which is significant for a 30 Hz image stream. Instead, we make use of the known camera extrinsics and the preintegrated IMU measurements between states to associate features geometrically.

When feature depth is available at the current image time t_k (e.g., from stereo camera triangulation), then we can directly project the visual landmark \mathbf{m}_ℓ from camera \mathbf{C}_i into camera \mathbf{C}_j via the extrinsic transformation between cameras $\mathbf{T}_{\mathbf{C}_j \mathbf{C}_i}$,

$$(u, v)_{\ell, \mathbf{C}_j, t_k} = \pi_{\mathbf{C}_j} (\mathbf{T}_{\mathbf{C}_j \mathbf{C}_i} \mathbf{C}_i \mathbf{m}_{\ell, t_k}) \quad (6)$$

Alternatively, if the feature depth is not directly available at the current timestamp (e.g., if the feature is only tracked in a monocular camera), then we instead use the preintegrated IMU measurements to estimate the landmark location at the current time,

$$(u, v)_{\ell, \mathbf{C}_j, t_k} = \pi_{\mathbf{C}_j} \left(\mathbf{T}_{\mathbf{C}_j \mathbf{C}_i} \hat{\mathbf{T}}_{t_k, t_{k-1}} \mathbf{C}_i \mathbf{m}_{\ell, t_{k-1}} \right) \quad (7)$$

where $\hat{\mathbf{T}}_{t_k, t_{k-1}}$ is the estimated transform between the camera poses at t_{k-1} and the current time t_k .

This process is completed for each feature in the overlapping image regions. The projected feature location is refined by matching it to the closest image feature in camera c_j , using a Euclidean distance metric. This is a highly effective method in practice, assuming there is a good extrinsic calibration. Any incorrect associations are handled by the optimizer using robust cost functions.

Fig. 4 shows an example of the feature matches using CCFT. The purple areas highlight the overlapping image regions where features can be tracked across cameras. The CCFT method typically reduces the number of landmarks added into the optimization back-end and gives improved estimation accuracy. This is discussed further in Section IX-A.

VI. SUBMATRIX FEATURE SELECTION

In general, increasing the number of features tracked in a VIO system improves the estimation accuracy [15]. However, it also increases computation, eventually reaching a point where the algorithm can fail due to computational constraints. This is a particular problem in multi-camera systems, where more features can be tracked than the optimizer can handle. Thus it is important to track and optimize only the best features.

Specifically, our Submatrix Feature Selection (SFS) algorithm is based on [22], where the authors presented a feature selection algorithm to reduce the computational cost of active map-to-frame feature matching. We instead apply this algorithm to MC-VIO and implement changes to improve numerical stability. The aim is to maintain accuracy while reducing optimization time of MC-VIO.

In this section we first outline the construction of the feature Jacobian and covariance matrices. Then we describe the *Max-logDet* algorithm to select the most representative subset of tracked features to be optimized.

A. Construct Joint Feature Jacobian Matrix

The aim of our optimization is to minimize the cost function defined in Eq. (3). Since we are interested in selecting the best features in each frame k we focus on improving the conditioning of the landmark residual,

$$\mathcal{X}^* = \arg \min_{\mathcal{X}} \sum_{\ell \in \mathbf{M}} \|\mathbf{r}_{\mathbf{x}, \mathbf{m}_\ell}\|_{\Sigma_{\mathbf{m}_\ell}}^2 \quad (8)$$

$$= \arg \min_{\mathcal{X}} \sum_{\ell \in \mathbf{M}} \|\pi(\mathbf{T}, \mathbf{m}_\ell) - z_\ell\|_{\Sigma_{\mathbf{m}_\ell}}^2 \quad (9)$$

Stacking \mathbf{m}_ℓ , z_ℓ into column matrices \mathbf{M} , \mathbf{Z} we get,

$$\mathcal{X}^* = \arg \min_{\mathcal{X}} \|\Pi(\mathbf{T}, \mathbf{M}) - \mathbf{Z}\|_{\Sigma}^2 \quad (10)$$

Applying the method from [22], we perform a first order linearization around the initial value \mathbf{T}_g ,

$$\|\Pi(\mathbf{T}, \mathbf{M}) - \mathbf{Z}\|_{\Sigma}^2 \approx \|\Pi(\mathbf{T}_g, \mathbf{M}) + H_{\mathbf{T}}(\mathbf{T} \ominus \mathbf{T}_g) - \mathbf{Z}\|_{\Sigma}^2 \quad (11)$$

where $H_{\mathbf{T}}$ is the pose Jacobian linearized about the value \mathbf{T}_g , and \ominus is the Lie group difference operator between poses [8]. To minimize Eq. (11), we can use Gauss-Newton optimization to iteratively update the pose estimate,

$$\mathbf{T}_{g+1} = \mathbf{T}_g - H_{\mathbf{T}}^+ (\Pi(\mathbf{T}_g, \mathbf{M}) - \mathbf{Z}) \quad (12)$$

where $H_{\mathbf{T}}^+$ is the Moore-Penrose pseudo-inverse of $H_{\mathbf{T}}$. This method has converged (i.e., reached minimal error) when the update step is approximately zero, i.e.,

$$\epsilon_{\mathbf{T}} = H_{\mathbf{T}}^+ (\Pi(\mathbf{T}_g, \mathbf{m}) - \mathbf{Z}) \approx 0 \quad (13)$$

Performing another first order linearization around \mathbf{M}_g ,

$$\epsilon_{\mathbf{T}} \approx H_{\mathbf{T}}^+ (\Pi(\mathbf{T}_g, \mathbf{M}_g) - \mathbf{Z} + H_{\mathbf{M}}(\mathbf{M} - \mathbf{M}_g)) \quad (14)$$

$$\epsilon_{\mathbf{T}} = H_{\mathbf{T}}^+ (\epsilon_{\mathbf{Z}} - H_{\mathbf{M}} \epsilon_{\mathbf{M}}) \quad (15)$$

where $H_{\mathbf{M}}$ is the landmark-to-image projection Jacobian. We define the measurement error as $\epsilon_{\mathbf{Z}} = \Pi(\mathbf{T}_g, \mathbf{M}_g) - \mathbf{Z}$, and the landmark error as $\epsilon_{\mathbf{M}} = \mathbf{M} - \mathbf{M}_g$.

Without loss of generality, the measurement error can be modeled as $\epsilon_{z_\ell} \sim \mathcal{N}(0, \Sigma_{z_\ell})$ which is influenced by image processing parameters and noise.

However, there could exist biased landmark error, such that the mean of error distribution is non-zero. This could be caused by batch perturbation of the landmark positions by incorrect camera poses, or a slight scale error. Hence the landmark error is modeled as non-zero-mean i.i.d. Gaussian, $\epsilon_{\mathbf{m}_\ell} \sim \mathcal{N}(\mu_{\mathbf{m}_\ell}, \Sigma_{\mathbf{m}_\ell})$, such that the expected value of $\epsilon_{\mathbf{T}}$ is,

$$\mathbb{E}[\epsilon_{\mathbf{T}}] = \mathbb{E}[H_{\mathbf{T}}^+ (\epsilon_{\mathbf{Z}} - H_{\mathbf{M}} \epsilon_{\mathbf{M}})] \quad (16)$$

$$\mathbb{E}[\epsilon_{\mathbf{T}}] = H_{\mathbf{T}}^+ H_{\mathbf{M}} \mathbf{1}_n \mu_{\mathbf{M}} \quad (17)$$

$$H_{\mathbf{M}}^+ H_{\mathbf{T}} \mathbb{E}[\epsilon_{\mathbf{T}}] = \mathbf{1}_n \mu_{\mathbf{M}} \quad (18)$$

where $\mathbf{1}_n$ is a tall matrix made of identity matrices to transform the dimension of $\mu_{\mathbf{M}}$. $H_{\mathbf{M}}$ consists of n diagonal blocks made up of individual landmark Jacobians $H_{\mathbf{m}_\ell} \in \mathbb{R}^{2 \times 3}$, where n is the number of landmarks with known 3D position. $H_{\mathbf{T}}$ consists of pose Jacobians for each landmark $H_{\mathbf{T}, \mathbf{m}_\ell} \in \mathbb{R}^{2 \times 6}$.

To obtain a joint Jacobian matrix H_J , we add an additional row of zeros to $H_{\mathbf{T}, \mathbf{m}_\ell}$ and a row of $[0, 0, 1]$ to $H_{\mathbf{m}_\ell}$. This makes $H_{\mathbf{T}, \mathbf{m}_\ell}$ invertible without changing the structure of the least squares problem [23]. Performing block-wise multiplication results in the matrix H_J ,

$$H_J = [H_{\mathbf{m}_0}^{-1} H_{\mathbf{T}, \mathbf{m}_0} \quad \dots \quad H_{\mathbf{m}_{n-1}}^{-1} H_{\mathbf{T}, \mathbf{m}_{n-1}}]^T \quad (19)$$

from which the simplified pose covariance matrix follows,

$$\Sigma_{\mathbf{T}} = H_J^+ (H_J^+)^T = (H_J^T H_J)^{-1} \quad (20)$$

B. Submatrix Selection

Since the overall pose error depends on the spectral properties of H_J , we aim to find a submatrix H_{sub} (i.e., subset of features) that best preserves this distribution. This can be understood as finding row blocks in H_J to maximize the norm of the selected submatrix.

We adapt the stochastic sampling method presented in [22] for feature selection. We choose the *Max-logdet* metric since it best approximates the original full feature set in visual odometry [23], [24]. Additionally, the combination of greedy (deterministic) and stochastic sampling (randomized acceleration) gives the best approximation ratio of any polynomial time algorithm [25].

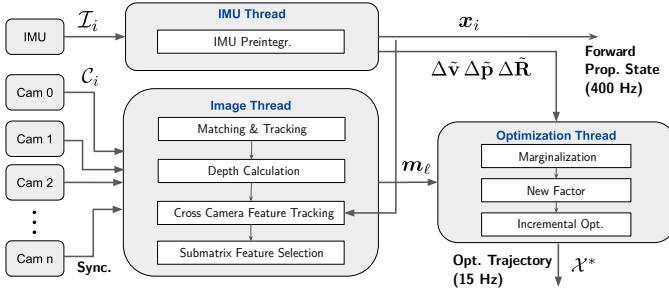


Fig. 5. Overview of the VILENS-MC architecture. The IMU and cameras are handled in separate threads by front-end measurement handlers. The back-end produces both a high frequency forward-propagated output and a lower frequency optimized output.

We note that since $H_{J_i} \in \mathbb{R}^{2 \times 6}$ is not a full rank matrix, the determinant of its square $H_{J_i}^T H_{J_i}$ would be zero. Hence in contrast to [22] add a small diagonal matrix λI to the log det calculation,

$$H_{J_i} = \arg \max (\log \det (H_{J_i}^T H_{J_i} + H_{sub}^T H_{sub} + \lambda I)) \quad (21)$$

This makes the squared matrix full rank with a non-zero determinant while preserving its spectral properties.

We apply this algorithm to each camera individually in our multi-camera setup. An ablation study, showing how this method maintains high accuracy and improves optimization time, is presented in Section IX-B.

VII. IMPLEMENTATION

An overview of the system architecture is shown in Fig. 5. Three parallel threads perform IMU preintegration, camera image processing, and optimization, respectively. The system outputs the optimized state estimate from the factor-graph at camera keyframe frequency (typically 15 Hz) for navigation or mapping, while a forward-propagated state at IMU frequency is also produced at 400 Hz if needed for high frequency tasks.

The factor graph is solved using fixed lag smoothing based on the efficient incremental optimization solver iSAM2 [26] from the GTSAM library [19]. For all of our experiments, we used a lag time of 3.5 s. To reduce the effect of outliers, the visual factors are added to the graph using the Dynamic Covariance Scaling (DCS) robust cost function [27].

A. Hardware and Calibration

The sensor used for our experiments is the Alphasense multi-camera development kit from Sevensense Robotics AG, shown in Fig. 2. An onboard FPGA synchronizes the IMU and four grayscale fisheye cameras – a frontal stereo pair with an 11 cm baseline and two lateral cameras. Each camera has a FoV of $126^\circ \times 92.4^\circ$ and a resolution of 720×540 px. This configuration produces an overlapping FoV between the front and side cameras of about 36° . The cameras and the embedded cellphone-grade IMU are operated at 30 Hz and 400 Hz, respectively. Camera-IMU extrinsic and intrinsic calibration was conducted offline using the Kalibr [28] toolbox. The Ouster OS0-128 Gen2 lidar shown in the figure was used only for ground truth generation and all sensors are hardware synchronized with an Intel NUC.



Fig. 6. Test datasets: *Top*: Newer College Dataset [29]. *Bottom Left*: Oxford Mathematical Institute. *Bottom Right*: A dark underground limestone mine.

B. Initialization

We initialize the IMU biases by averaging the first 1 s of data at system start up (assuming the IMU is stationary).

The scale initialization problem, which is often present in monocular visual odometry systems, is solved by a combination of preintegrated IMU measurements and depth from the stereo camera pair. The tracking of features across cameras also improves the speed of the initialization process.

C. Visual Feature Tracking

To get an even feature distribution across each image, images are split into 3×3 tiles before applying the FAST feature detector to each segment. These features are tracked between frames using the KLT feature tracker and across cameras using the CCFT method described above. Outliers are rejected using RANSAC-based methods. Thanks to the incremental optimization and multi-threading every second frame is used as a keyframe, achieving 15 Hz nominal output.

VIII. EXPERIMENTAL RESULTS

We evaluated our algorithm on a variety of challenging indoor and outdoor datasets, varying from narrow corridors to large open spaces. We compared our algorithm against state-of-the-art methods, evaluating both quantitative and qualitative performance. We also include ablation studies showing how our proposed contributions can reduce RPE and decrease computation time. Finally we demonstrate the versatility of this approach by applying VILENS-MC to a quadruped robot operating in an unlit mine.

A. Datasets

Fig. 6 gives an overview of the datasets used for evaluation. These include additional multi-camera experiments collected for the Newer College Dataset (NCD) [29], operation in an unlit underground Limestone Mine (MINE), and a large circuit around the Oxford Mathematical Institute (MATH).

These datasets were collected using the handheld multi-camera device described in Section VII-A. They were specifically chosen to test the limits of VIO systems and include challenges such as aggressive shaking of the device (up to




Fig. 7. Examples of challenging scenes where it was difficult to track features. Orange and green circles are features tracked by monocular and stereo cameras, respectively (some have been tracked across images). Note that stereo features were mostly absent or rejected in these examples.

5.5 rad/s), severe illumination changes, and loss of visual feature tracking in one or more cameras (see Fig. 7). All datasets were collected at a fast walking pace of ~ 1.5 m/s.

The specific experiments are:

- **NCD-Quad:** Moving between narrow corridors and a large open college quadrangle, with strong illumination changes caused by bright sunlight (244 m, 3 min).
- **NCD-Stairs:** Traversing a very narrow staircase, including a sequence where a door is opened directly in front of the camera (59 m, 2 min).
- **MINE:** A medium scale, dark underground environment with regular illumination changes due to onboard lighting (236 m, 3 min).
- **MATH:** A large scale outdoor environment with aggressive shaking of the device up to 5.5 rad/s (329 m, 4 min).

The multi-camera experiments which extend NCD (NCD-Quad, NCD-Stairs) will be publicly released¹.

To generate ground truth, ICP was used to align the current LIDAR scan to detailed prior maps, collected using a commercial laser mapping system. The high frequency motion estimate from the IMU was used to carefully remove lidar motion distortion [7]. For an in-depth discussion on ground truth generation the reader is referred to [29].

B. Quantitative Analysis

Table I shows the Relative Pose Error (RPE) over 10 m for OpenVINS [12], ORB-SLAM3 [10] and VILENS [8] (which our system builds upon). We refer to the configuration of

¹Available at <https://ori-drs.github.io/newer-college-dataset/>




Fig. 8. Top-down view of **MATH** dataset comparing the estimated trajectory of ORB-SLAM3, OpenVINS, and VILENS-MC against ground truth. A black circle marks the start of the trajectory, while colored crosses indicate the last pose of each trajectory. Textureless locations, where cameras were facing walls, are shown in Fig. 7, row 4.

10m Relative Pose Error (RPE) – Translation [m] / Rotation [°]				
	OpenVINS	ORB-SLAM3	VILENS-Stereo	VILENS-MC
Cameras	2	2	2	4
NCD-QUAD	1.01 / 1.26	0.23 / 0.91	0.30 / 1.24	0.31 / 0.82
NCD-STAIR	0.33 / 3.05	0.20* / 3.04	0.43 / 5.85	0.16 / 2.56
MINE	0.98* / 2.10	Fail	0.41* / 3.57	0.20 / 1.65
MATH	0.65* / 1.69	Fail	0.56* / 1.57	0.26 / 1.03

TABLE I
PERFORMANCE COMPARISON BETWEEN DIFFERENT ALGORITHMS
* = ONE OR MORE INSTABILITIES OCCURRED

VILENS which uses only stereo and IMU input as *VILENS-Stereo*. Note that the authors were not aware of any open-source MC-VIO algorithm for comparison.

Compared to our baseline (VILENS-Stereo), we reduce mean RPE by 45 % / 50 % in translation/rotation, while tracking fewer features. VILENS-MC also outperforms state-of-the-art stereo VIO algorithms by up to 80 % / 39 % in translation/rotation.

In benign VIO conditions, such as those found in the NCD-Quad dataset, the performance of VILENS-MC was similar to state-of-the-art systems. In particular, ORB-SLAM3 appears to outperform the other algorithms by taking advantage of its local map matching SLAM system to correct for failure events when the camera revisited a previous location.

However, in the MINE and MATH datasets, where there are few or no stereo features, VILENS-MC performs well – even when the other stereo VIO systems fail or have significant drift. This highlights a key benefit of the tight fusion of multiple cameras – natural robustness to scene degradation.

C. Degeneracy and Failure Analysis

1) **MATH:** A top-down view of the trajectories in the MATH dataset is shown in Fig. 8. In challenging scenarios where stereo camera feature tracking fails (marked in Fig. 8 as “textureless




Fig. 9. ORB-SLAM3 lost tracking and failed to estimate odometry for 4 s in NCD-Stair (scenario shown in Fig. 7 Row 3). Note that elevation (z axis) is properly tracked afterwards.

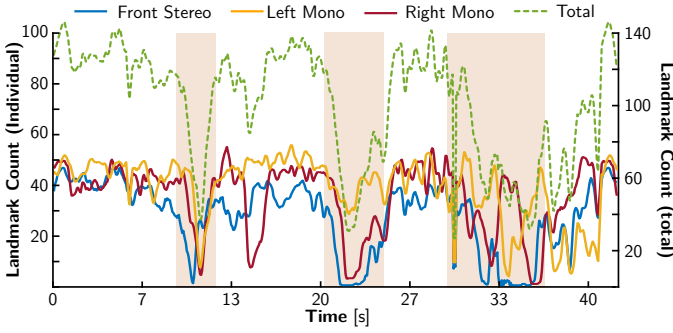


Fig. 10. A 40 s section from the **MINE** dataset showing landmark count per camera, blue line is the front stereo cameras, orange is the right, yellow is the left monocular camera, and green is the combined number of landmarks.

areas”), VIO algorithms may continue to operate but would rely on the IMU only, causing the RPE to quickly increase.

2) *NCD-Stair*: Fig. 9 shows an extract of the estimates from the NCD-Stair dataset. Note how ORB-SLAM3 suffers from significant drift when no features were detected in the stereo camera for 4 s until it re-initializes (orange area). Even though the average RPE over the entire NCD-Stair sequence remains low, this would be undesirable for an autonomous robot.

3) *MINE*: Fig. 10, shows the number of tracked landmarks during a 40 s interval in the MINE dataset. During this time, several instances occurred where feature tracking failed in either one or two cameras. These occurred when the device entered a dark room (around 10 s), when the device approached a wall and turned on the spot (around 20 s, shown in Fig. 7 row 3), and when the device entered a dark corner where only one camera could observe features at a time. By tracking features in multiple cameras, VILENS-MC maintained enough features for accurate estimation throughout these events.

IX. ABLATION STUDY

A. Cross Camera Feature Tracking

In this section, we test the effect of the CCFT method (Section V). This method enables continuous feature tracks over longer periods and adds additional constraints to landmarks.

Table II summarizes the effect this method has on mean RPE (over 10 m). It shows that on average CCFT decreases both translational and rotational RPE on all datasets except for

	Relative Pose Error – Translation [m] / Rotation [°]				
	NO-CCFT	WITH-CCFT	NO-SFS (90 feat.)	NO-SFS (150 feat.)	WITH-SFS (90 feat.)
NCD-QUAD	0.27 / 0.77	0.34 / 0.88	0.37 / 1.02	0.34 / 0.88	0.31 / 0.82
NCD-STAIR	0.17 / 3.60	0.16 / 2.85	0.28 / 5.06	0.16 / 2.85	0.18 / 3.03
MINE	0.25 / 1.91	0.21 / 1.64	0.34 / 2.70	0.21 / 1.71	0.20 / 1.75
MATH	0.42 / 1.01	0.28 / 1.01	0.33 / 1.16	0.28 / 1.01	0.26 / 1.03

TABLE II
ABLATION STUDIES: CROSS CAMERA FEATURE TRACKING (CCFT) AND SUBMATRIX FEATURE SELECTION (SFS)

translation in NCD-QUAD. We attribute this to the system’s susceptibility to extrinsic multi-camera calibration errors.

Additionally, CCFT can reduce the number of landmarks in the optimization by $\sim 10\%$. Since our hardware setup (ref. Section VII-A) had only a small overlapping FoV between frontal and lateral cameras, we expect to see a greater benefit if there was a larger overlapping FoV.

B. Submatrix Feature Selection

This ablation study demonstrates the benefit of the SFS method. The baseline approach simply spreads features evenly across the image based on their FAST score and mutual pixel distance.

A comparison between the baseline (NO-SFS) and proposed (WITH-SFS) methods is shown in Table II. The first two columns show NO-SFS with 90 and 150 features tracked and optimized across all cameras, while the WITH-SFS method selected only 90 out of 150 tracked features to optimize.

By selecting only the most representative features, WITH-SFS outperformed the NO-SFS method by up to 36 %, even though both methods optimized the same number of features.

Crucially, WITH-SFS offers lower computation without sacrificing accuracy. Comparing NO-SFS (with 150 features) to WITH-SFS (with only 90 features) we can see that SFS achieves similar accuracy with a 20 % reduction in optimization time (53.7 ms to 43.0 ms). The standard deviation of the computation was also reduced by 24.3 %. This demonstrates that SFS can achieve similar accuracy with lower and more consistent computational requirements.

X. DEMONSTRATION ON LEGGED ROBOTS

To demonstrate the versatility of our approach we tested VILENS-MC on the ANYmal C quadruped robot which has been equipped with an Alphasense for the DARPA Subterranean Challenge [2]. In the trajectory shown in Fig. 11 the robot autonomously explored the same environment as the MINE dataset. VILENS-MC achieved almost the same accuracy as for the handheld dataset (0.27 m, 2.65° RPE), indicating that performance was similar across the different platforms.

XI. CONCLUSION

We presented a novel factor graph formulation for state estimation that tightly fuses an arbitrary number of cameras. The joint optimization of IMU measurements with monocular and camera constraints enables graceful handling of degenerate scenarios without requiring hard switching between cameras. This also simplifies the initialization of the individual monocular cameras. We have demonstrated comparable tracking




Fig. 11. State estimate on an AnyMal C quadruped robot autonomously exploring an underground mine over ~ 10 min. Fig. 6 shows the robot's on-board lighting which caused strong illumination changes in the images.

performance to state-of-the-art visual-inertial odometry systems in benign conditions and better performance in extreme situations, such as the loss of tracking in one or more of the cameras, aggressive motions, or abrupt light changes.

We proposed two algorithmic components specific to multi-camera odometry: cross camera feature tracking and submatrix feature selection. Ablation studies showed that these components improve accuracy and reduce optimization time.

Overall, we have presented a robust MC-VIO system that is capable of handling various challenging real-world environments while consistently achieving accurate state estimation. Future work would consider including camera extrinsics into the optimization.

ACKNOWLEDGMENT

This research was part funded by the EU H2020 Project THING and the Innovate UK-funded ORCA Robotics Hub (EP/R026173/1). Fallon was supported by a Royal Society University Research Fellowship and Wisth by a Google DeepMind studentship. This research has been conducted as part of the ANYmal research community.

REFERENCES

- [1] C. Forster, L. Carlone, F. Dellaert, and D. Scaramuzza, "On-Manifold Preintegration for Real-Time Visual-Inertial Odometry," *IEEE Transactions on Robotics*, vol. 33, no. 1, pp. 1–21, 2017.
- [2] M. Tranzatto, *et al.*, "CERBERUS: Deployment at the DARPA Subterranean Challenge Urban Circuit," *Field Robotics*, 2021, accepted, to appear.
- [3] M. Hutter, *et al.*, "ANYmal – A Highly Mobile and Dynamic Quadrupedal Robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2016, pp. 38–44.
- [4] J. Jaekel, J. G. Mangelson, S. Scherer, and M. Kaess, "A robust multi-stereo visual-inertial odometry pipeline," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [5] H. Seok and J. Lim, "ROVINS: Robust Omnidirectional Visual Inertial Navigation System," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6225–6232, 2020.

- [6] K. Eickenhoff, P. Geneva, and G. Huang, "MIMC-VINS: A Versatile and Resilient Multi-IMU Multi-Camera Visual-Inertial Navigation System," *IEEE Transactions on Robotics*, pp. 1–20, 2021.
- [7] D. Wisth, M. Camurri, S. Das, and M. Fallon, "Unified multi-modal landmark tracking for tightly coupled lidar-visual-inertial odometry," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1004–1011, 2021.
- [8] D. Wisth, M. Camurri, and M. Fallon, "VILENS: Visual, inertial, lidar, and leg odometry for all-terrain legged robots," *arXiv preprint arXiv:2107.07243*, 2021.
- [9] T. Qin, P. Li, and S. Shen, "VINS-Mono: A Robust and Versatile Monocular Visual-Inertial State Estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [10] 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 multi-map SLAM," *IEEE Transactions on Robotics*, pp. 1–17, 2021.
- [11] S. Leutenegger, S. Lynen, M. Bosse, R. Siegwart, and P. Furgale, "Keyframe-based visual-inertial odometry using nonlinear optimization," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 314–334, 2015.
- [12] P. Geneva, K. Eickenhoff, W. Lee, Y. Yang, and G. Huang, "OpenVINS: A Research Platform for Visual-Inertial Estimation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 4666–4672.
- [13] J. Delmerico and D. Scaramuzza, "A benchmark comparison of monocular visual-inertial odometry algorithms for flying robots," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 2502–2509.
- [14] M. He, C. Zhu, Q. Huang, B. Ren, and J. Liu, "A review of monocular visual odometry," *The Visual Computer*, vol. 36, no. 5, pp. 1053–1065, 2020.
- [15] P. Liu, M. Geppert, L. Heng, T. Sattler, A. Geiger, and M. Pollefeys, "Towards Robust Visual Odometry with a Multi-Camera System," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [16] M. G. Müller, *et al.*, "Robust visual-inertial state estimation with multiple odometries and efficient mapping on an MAV with ultra-wide FOV stereo vision," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3701–3708, 2018.
- [17] J. Kuo, M. Muglikar, Z. Zhang, and D. Scaramuzza, "Redesigning SLAM for Arbitrary Multi-Camera Systems," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 2116–2122.
- [18] C. Forster, Z. Zhang, M. Gassner, M. Werlberger, and D. Scaramuzza, "SVO: Semidirect visual odometry for monocular and multicamera systems," *IEEE Transactions on Robotics*, vol. 33, no. 2, pp. 249–265, 2017.
- [19] F. Dellaert and M. Kaess, "Factor Graphs for Robot Perception," *Foundations and Trends in Robotics*, vol. 6, pp. 1–139, 2017.
- [20] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.
- [21] R. Mur-Artal and J. D. Tardós, "ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [22] Y. Zhao and P. A. Vela, "Good Feature Matching: Toward Accurate, Robust VO/VSLAM with Low Latency," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 657–675, 2020.
- [23] —, "Good Feature Selection for Least Squares Pose Optimization in VO/VSLAM," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 1183–1189.
- [24] L. Carlone and S. Karaman, "Attention and Anticipation in Fast Visual-Inertial Navigation," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 1–20, 2019.
- [25] B. Mirzasoleiman, A. Badanidiyuru, A. Karbasi, J. Vondrák, and A. Krause, "Lazier than lazy greedy," *Proceedings of the National Conference on Artificial Intelligence*, vol. 3, pp. 1812–1818, 2015.
- [26] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping with fluid relinearization and incremental variable reordering," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 3281–3288.
- [27] K. MacTavish and T. D. Barfoot, "At All Costs: A Comparison of Robust Cost Functions for Camera Correspondence Outliers," in *Conference on Computer and Robot Vision*, 2015, pp. 62–69.
- [28] J. Rehder, J. Nikolic, T. Schneider, T. Hinzmann, and R. Siegwart, "Extending kalibr: Calibrating the extrinsics of multiple imus and of individual axes," in *IEEE International Conference on Robotics and Automation*, 2016, pp. 4304–4311.

- [29] M. Ramezani, Y. Wang, M. Camurri, D. Wisth, M. Mattamala, and M. Fallon, "The Newer College Dataset: Handheld LiDAR, Inertial and Vision with Ground Truth," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.