

Ultrafast Square-Root Filter-based VINS

Yuxiang Peng, Chuchu Chen and Guoquan Huang

Abstract—In this paper, we strongly advocate square-root covariance (instead of information) filtering for Visual-Inertial Navigation Systems (VINS), in particular on resource-constrained edge devices, because of its superior efficiency and numerical stability. Although VINS have made tremendous progress in recent years, they still face resource stringency and numerical instability on embedded systems when imposing limited word length. To overcome these challenges, we develop an ultrafast and numerically-stable square-root filter (SRF)-based VINS algorithm (i.e., SR-VINS). The numerical stability of the proposed SR-VINS is inherited from the adoption of square-root covariance while the remarkable efficiency is largely enabled by the novel SRF update method that is based on our new permuted-QR (P-QR), which fully utilizes and properly maintains the upper triangular structure of the square-root covariance matrix. Furthermore, we choose a special ordering of the state variables which is amenable for (P-)QR operations in the SRF propagation and update and prevents unnecessary computation. The proposed SR-VINS is validated extensively through numerical studies, demonstrating that when the state-of-the-art (SOTA) filters have numerical difficulties, our SR-VINS has superior numerical stability, and remarkably, achieves efficient and robust performance on 32-bit single-precision float at a speed nearly twice as fast as the SOTA methods. We also conduct comprehensive real-world experiments to validate the efficiency, accuracy, and robustness of the proposed SR-VINS.

I. INTRODUCTION

Visual-Inertial Navigation Systems (VINS) that employ a single camera and an inertial measurement unit (IMU) to provide 3D motion tracking, have great potential in many applications such as AR/VR and robotics [1]–[4]. VINS state estimation algorithms can be categorized into covariance and information forms. In the former such as the extended Kalman filter (EKF) and its variants, the estimator keeps tracking the dense covariance matrix to update the estimate [5]–[11]. In contrast, the information estimators such as extended Information filters (EIF) [12] or optimization-based methods [13]–[18], maintain the information (Hessian) matrix and exploit its sparse structure in solving for estimates. However, both covariance and information filters face challenging numerical issues, in particular on resource-constrained edge platforms [19], [20], when limited word length (32-bit float, instead of 64-bit double) is available or it is required to achieve a potential speedup by leveraging SIMD (Single Instruction/Multiple Data) to vectorized matrix operations [21]. In the covariance form, the covariance matrix tends to lose its positive definiteness and cause the filter to diverge. In the information form, as the information matrix can easily become ill-conditioned (e.g., condition

number larger than 10^9 [19]), naively inverting it during optimization would lead to large numerical errors (see Chapter 3.5.1 in [22]).

There exist methods that use the square root of the information matrix instead of its full matrix to mitigate the numerical instability and were shown to be effective to some extent in VINS [20], [23]–[32]. For example, the method in [20] maintains an upper triangular square root of prior information and uses QR-decomposition to incorporate new measurements into the prior, then invert it to solve for the state update. While this estimator achieves the same accuracy with the half of the word length, it still has the concerning numerical issue with a relatively high condition number (10^5) over time, especially when paired with a high-precision IMU [33], [34], resulting in substantial numerical inaccuracies that challenge long-term operations.

On the contrary, VINS estimators in the covariance form tend to offer better numerical stability. For instance, in the EKF-based VINS, the only matrix that typically requires inversion, the innovation covariance S , usually possesses a good condition number [19]. By using the square-root covariance matrix, we can not only inherit the merits of the covariance form but also benefit from square-root properties. Surprisingly, this idea of square-root filter (SRF) remains largely unexplored in VINS, primarily due to its inefficiency. Looking into history, the SRF has undergone significant improvements over the decades. Back in the 1960s, the initial SRF formulation was proposed by Potter and played a significant role in the Apollo project’s success [35], [36], which has been extended to account for propagation (process) noise [37]. A key challenge is to improve its efficiency. Solutions include update methods using eigenvalue [38], Cholesky [39] and QR decomposition [40] within the SRF framework. However, compared with the conventional KF, these update methods in the SRF were shown to be less efficient, which is mainly because the triangular structure of the square-root covariance has been broken after the update. Agee [41] and Carlson [42] proposed update methods that maintain triangular structure and exhibit comparable efficiency to the KF. However, these methods are limited to sequential updates. In modern computers, batch updates involving vector operations are preferable, allowing for more level-3 BLAS [43] operations.

To address the aforementioned issues and fully utilize the benefit of the square-root covariance, in this work, we develop a novel (P-)QR-based SRF for VINS, termed SR-VINS. In particular, we propose a new permuted-QR (P-QR) decomposition that fully utilizes the upper-triangular structure during matrix factorization, which is theoretically shown to improve efficiency during batch updates. Additionally, when integrating into the sliding-window filtering framework, the proposed SR-VINS chooses a special ordering of the state variables which is amenable for (P-)QR operations

This work was partially supported by the University of Delaware (UD) College of Engineering, Delaware NASA/EPSCoR Seed Grant, NSF (IIS-1924897, SCH-2014264), Google ARCore, and Meta Reality Labs.

The authors are with the Robot Perception and Navigation Group (RPNG), University of Delaware, Newark, DE 19716, USA. Email: {yxpeng, ccchu, ghuang}@udel.edu

in the SRF propagation and update and prevents unnecessary computation. Specifically, our main contributions can be summarized as follows:

- We propose a novel permuted-QR (P-QR) decomposition that not only fully utilizes the upper-triangular structure during matrix factorization, but also helps maintain the upper-triangular structure of the square-root covariance. With that, we develop an efficient (P-QR)-based SRF update method, which is shown to be significantly faster than the existing methods if $m > \frac{2}{3}n$ (where m and n are measurements and states size).
- We are among the first to design the SRF-based VINS with online calibration within an efficient sliding-window filter framework, which achieves unprecedented and remarkable numerical stability even running on 32-bit. Our implementation demonstrates notable efficiency gain as it is almost two times faster than the state-of-the-art filters.
- We perform extensive numerical studies to highlight potential numerical challenges in VINS and underscore the advantages of our proposed SR-VINS. Real-world experiments validate the notable efficiency boost of the proposed method while maintaining accuracy.

II. EFFICIENT SQUARE-ROOT FILTERING

In comparison to a canonical EKF (or its variants) tracking the dense covariance matrix \mathbf{P} , the SRF propagates and updates the corresponding *upper triangular* square-root matrix \mathbf{U} , i.e., $\mathbf{U}^\top \mathbf{U} = \mathbf{P}$, while its state estimates are computed in the same way as the EKF [44]. By doing so, in principle, it possesses some key features that are particularly compelling to visual-inertial estimation at the edge. For example, the SRF can represent a broader dynamic range and reduce numerical errors by using a reduced condition number (i.e., the square root of the condition number of \mathbf{P}), thus offering better numerical stability. Moreover, the SRF can be significantly more efficient in both computation and memory consumption because it can use lower precision without sacrificing accuracy. Additionally, it automatically ensures the symmetry and positive semi-definite of the corresponding covariance matrix. However, in practice, it is not easy to capitalize these benefits if blindly implementing the SRF, because the much-needed matrix triangulation operations required in the filter are computationally expensive. This is one of the reasons that the SRF has not been widely adopted in VINS, despite the aforementioned theoretical advantages. In this work, we, for the first time, fully take advantage of the upper-triangular structure of the SRF in VINS, by leveraging a new Permuted-QR (P-QR) algorithm.

A. Permuted-QR Decomposition

In contrast to the standard QR decomposition, for example, based on Givens rotation [45] or Householder [46], the proposed P-QR yields a lower, instead of upper, triangular matrix based on the following lemma:

Lemma 1: For a full-rank matrix $\mathbf{M}_{m \times n}$ ($m > n$), there exists the following lower-triangular P-QR decomposition:¹

$$\mathbf{M}_{m \times n} = \begin{bmatrix} \mathbf{A}_{(m-n) \times n} \\ \mathbf{B}_{n \times n} \end{bmatrix} \stackrel{\text{P-QR}}{=} \mathbf{Q}_{1_{m \times m}} \begin{bmatrix} \mathbf{0}_{(m-n) \times n} \\ \mathbf{F}_{n \times n} \end{bmatrix} \quad (1)$$

¹Although we here assume the matrix \mathbf{M} is of full column rank, our P-QR is applicable to rank-deficient matrices as the standard QR is.




Fig. 1: Evolution of the matrix structure when performing the proposed P-QR decomposition.

where \mathbf{Q}_1 is orthonormal and \mathbf{F} is lower triangular.

Proof: We employ an anti-diagonal permutation matrix, $\Pi = \text{adiag}([1 \cdots 1]_n)$, and a row permutation matrix Γ , to transform \mathbf{M} into \mathbf{M}' as:

$$\mathbf{M} := \Gamma \Gamma^\top \mathbf{M} \Pi \Pi^\top = \Gamma \Gamma^\top \mathbf{M}' \Pi^\top = \Gamma \mathbf{M}'' \Pi^\top \quad (2)$$

where Γ^\top is used to permute the rows of \mathbf{M}' to make \mathbf{M}'' as close to upper triangular as possible so that we can perform QR decomposition on it efficiently. This permutation can be found by searching the first non-zero element in each row and grouping them based on that to generate new row IDs that correspond to the permutation. The standard QR of \mathbf{M}'' yields the following orthonormal matrix \mathbf{Q}_2 and upper triangular matrix \mathbf{C} :

$$\mathbf{M}'' \stackrel{\text{QR}}{=} \mathbf{Q}_2 \begin{bmatrix} \mathbf{C} \\ \mathbf{0} \end{bmatrix} \quad (3)$$

Substitution of \mathbf{M}'' into (2) yields the following identities:

$$\mathbf{M} = \Gamma \mathbf{Q}_2 \begin{bmatrix} \mathbf{C} \\ \mathbf{0} \end{bmatrix} \Pi^\top = \underbrace{\Gamma \mathbf{Q}_2 \Pi'}_{\mathbf{Q}_1} \Pi'^\top \begin{bmatrix} \mathbf{C} \\ \mathbf{0} \end{bmatrix} \Pi^\top \quad (4)$$

$$= \mathbf{Q}_1 \Pi'^\top \begin{bmatrix} \mathbf{C} \\ \mathbf{0} \end{bmatrix} \Pi^\top = \mathbf{Q}_1 \begin{bmatrix} \mathbf{0} \\ \mathbf{F} \end{bmatrix} \quad (5)$$

where we have employed a new anti-diagonal permutation matrix $\Pi' = \text{adiag}([1 \cdots 1]_m)$ along with Π to permute the upper triangular \mathbf{C} into the lower triangular \mathbf{F} . ■

Figure 1 visualizes how the matrix structure evolves during the proposed P-QR decomposition. Note that the resulting lower triangular structure of \mathbf{F} will enable significant computation savings in the SRF update.

B. QR-based SRF Propagation

The proposed SRF leverages QR to propagate and update its state estimate and square-root matrix that are equivalent to those of the EKF. To see that, we know that the EKF propagates the state estimate and covariance as:

$$\hat{\mathbf{x}}_{k+1|k} = \Phi_k \hat{\mathbf{x}}_{k|k} \quad (6)$$

$$\mathbf{P}_{k+1|k} = \Phi_k \mathbf{P}_{k|k} \Phi_k^\top + \mathbf{W}_k \quad (7)$$

where Φ_k is the state transition matrix and \mathbf{W}_k is the system noise covariance. Detailed derivations are refer to our companion technical report [47]. The SRF propagates the state estimate in the same way as the EKF, while the square-root covariance matrix is propagated via the QR decomposition:

$$\begin{bmatrix} \mathbf{W}_k^{\frac{1}{2}} \\ \mathbf{U}_{k|k} \Phi_k^\top \end{bmatrix} \stackrel{\text{QR}}{=} \mathbf{Q}_k \begin{bmatrix} \mathbf{U}_{k+1|k} \\ \mathbf{0} \end{bmatrix} \quad (8)$$

where $\mathbf{W}_k = \mathbf{W}_k^{\frac{\top}{2}} \mathbf{W}_k^{\frac{1}{2}}$ via Cholesky factorization, and $\mathbf{U}_{k+1|k}$ is the upper-triangular covariance matrix. It is easy to verify that squaring (8) yields the covariance propagation (7), and thus they are equivalent.




Fig. 2: Visualization of the matrix structure during the SRF operations.

C. P-QR-based SRF Update

During the update, a canonical SRF does not exploit the special structure of the square-root matrix update and incurs more expensive operations. In contrast, we propose a novel square-root update equation that is significantly more efficient by leveraging the proposed P-QR decomposition.

Lemma 2: It is equivalent to the EKF update if the SRF updates its square-root covariance and state estimate as:

$$\mathbf{U}_{k|k} = \mathbf{F}_k^{-\top} \mathbf{U}_{k|k-1} \quad (9)$$

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{U}_{k|k}^{\top} \mathbf{U}_{k|k} \mathbf{H}_k^{\top} \mathbf{R}_k^{-1} \mathbf{r}_k \quad (10)$$

where \mathbf{H}_k is the measurement Jacobian, \mathbf{R}_k is the noise covariance, and \mathbf{r}_k is the residual. Most importantly, \mathbf{F}_k is *lower triangular* (and thus $\mathbf{F}_k^{-\top}$ is *upper triangular*), which is obtained by the following P-QR:

$$\begin{bmatrix} \mathbf{R}_k^{-\frac{1}{2}} \mathbf{H}_k \mathbf{U}_{k|k-1}^{\top} \\ \mathbf{I} \end{bmatrix} \stackrel{\text{P-QR}}{=} \mathbf{Q} \begin{bmatrix} \mathbf{0} \\ \mathbf{F}_k \end{bmatrix} \quad (11)$$

Note that the SRF first updates the square-root covariance $\mathbf{U}_{k|k}$ (9) and then use it to update the state $\hat{\mathbf{x}}_{k|k}$ (10).

Proof: From the EKF update equation, we have:

$$\begin{aligned} \mathbf{P}_{k|k} &= \mathbf{P}_{k|k-1} - \mathbf{P}_{k|k-1} \mathbf{H}_k^{\top} (\mathbf{H}_k \mathbf{P}_{k|k-1} \mathbf{H}_k^{\top} + \mathbf{R}_k)^{-1} \mathbf{H}_k \mathbf{P}_{k|k-1} \\ &= \mathbf{U}_{k|k-1}^{\top} \left(\mathbf{I} - \mathbf{U}_{k|k-1} \mathbf{H}_k^{\top} (\mathbf{H}_k \mathbf{U}_{k|k-1}^{\top} \mathbf{U}_{k|k-1} \mathbf{H}_k^{\top} + \mathbf{R}_k)^{-1} \mathbf{H}_k \mathbf{U}_{k|k-1}^{\top} \right) \mathbf{U}_{k|k-1} \\ &= \mathbf{U}_{k|k-1}^{\top} \left(\underbrace{\mathbf{I} + \mathbf{U}_{k|k-1} \mathbf{H}_k^{\top} \mathbf{R}_k^{-1} \mathbf{H}_k \mathbf{U}_{k|k-1}^{\top}}_{\mathbf{F}_k^{\top} \mathbf{F}_k} \right)^{-1} \mathbf{U}_{k|k-1} \\ &= \mathbf{U}_{k|k-1}^{\top} \mathbf{F}_k^{-1} \mathbf{F}_k^{-\top} \mathbf{U}_{k|k-1} =: \mathbf{U}_{k|k}^{\top} \mathbf{U}_{k|k} \end{aligned}$$

where we have employed the matrix inversion lemma in deriving the third equality. ■

It is important to stress that the P-QR efficiently computes the lower-triangular matrix \mathbf{F}_k , which enables efficient update of the square-root covariance because both $\mathbf{F}_k^{-\top}$ and $\mathbf{U}_{k|k-1}$ are upper triangular. To see this, notice first that the left-hand side of (11) has an identity matrix at the bottom. Leveraging this structure allows for efficient QR decomposition because there is no need to zero out the elements below the diagonals of \mathbf{I} . When solving for $\mathbf{U}_{k|k}$, even though inverting \mathbf{F}_k^{\top} is needed, thanks to its upper triangular structure, we can solve it efficiently using back substitution on $\mathbf{F}_k^{\top} \mathbf{U}_{k|k} = \mathbf{U}_{k|k-1}$. With these structure benefits, we calculate the number of arithmetic operations required in the update with the assumptions: (i) measurements are uncorrelated, (ii) the terms that have orders smaller than 3 are ignored and (iii) the Householder algorithm is used to perform standard QR. Table I shows that our proposed SRF requires fewer operations than the most competitive Carlson update when $m > \frac{2}{3}n$ (m measurements, n states), which is often the case in VINS.

TABLE I: FLOPs of the different measurement update assuming uncorrelated measurements and ignoring the terms of order lower than 3 (m measurements, n states).

Methods	Potter [35]	Carlson [42]	Proposed
Flops	$6mn^2$	$\frac{7}{2}mn^2$	$3mn^2 + \frac{1}{3}n^3$

III. SQUARE-ROOT SLIDING-WINDOW FILTER (SR-SWF)-BASED VINS

In this section, we apply the proposed efficient SRF to the VINS problem and develop an ultrafast and numerically stable visual-inertial estimation algorithm, termed SR-VINS. The proposed SR-VINS is formulated in an efficient sliding-window filtering framework, while exploiting the sparse triangular structure of the system to perform (P-)QR.

In particular, we employ a specially-ordered state vector to avoid unnecessary computations in the ensuing SRF operations. At time t_k , the system state \mathbf{x}_k consists of the current navigation states \mathbf{x}_{I_k} , calibration parameters \mathbf{x}_{cb} , historical IMU pose clones \mathbf{x}_C , and SLAM features \mathbf{x}_f :

$$\mathbf{x}_k = [\mathbf{x}_{I_k}^{\top} \mathbf{x}_{cb}^{\top} \mathbf{x}_C^{\top} \mathbf{x}_f^{\top}]^{\top} \quad (12)$$

$$\mathbf{x}_{I_k} = [{}^I_k \bar{q}^{\top} \quad {}^G \mathbf{p}_{I_k}^{\top} \quad {}^G \mathbf{v}_{I_k}^{\top} \quad \mathbf{b}_g^{\top} \quad \mathbf{b}_a^{\top}]^{\top} \quad (13)$$

$$\mathbf{x}_{cb} = [t_d \quad {}^I_C \bar{q}^{\top} \quad {}^C \mathbf{p}_I^{\top} \quad \zeta^{\top}]^{\top} \quad (14)$$

$$\mathbf{x}_C = [\mathbf{x}_{T_k}^{\top} \dots \mathbf{x}_{T_{k-c}}^{\top}]^{\top}, \quad \mathbf{x}_f = [\mathbf{f}_1^{\top} \dots \mathbf{f}_g^{\top}]^{\top} \quad (15)$$

where ${}^I_C \bar{q}$ is the unit quaternion (${}^I_C \mathbf{R}$ in rotation matrix form) that represents the rotation from the global $\{G\}$ to the IMU frame $\{I\}$; ${}^G \mathbf{p}_I$, ${}^G \mathbf{v}_I$ are the IMU position, velocity in $\{G\}$, and \mathbf{f}_i is i 'th feature position; \mathbf{b}_g and \mathbf{b}_a are the gyroscope and accelerometer biases; $\mathbf{x}_{T_i} = [{}^I_i \bar{q}^{\top} \quad {}^G \mathbf{p}_{I_i}^{\top}]^{\top}$. t_d denote the time offset between camera and IMU, $\{{}^I_C \bar{q}, {}^C \mathbf{p}_I\}$ is the extrinsic between camera and IMU sensors and ζ is the camera intrinsic parameters.

A. Propagation with Cloning and Marginalization

The IMU kinematics and its linearization are given by [48]:

$$\mathbf{x}_{I_{k+1}} = \mathbf{g}(\mathbf{x}_{I_k}, \mathbf{a}_k, \boldsymbol{\omega}_k, \mathbf{w}_k) \quad (16)$$

$$\tilde{\mathbf{x}}_{I_{k+1}|k} = \Phi_k \tilde{\mathbf{x}}_{I_k|k} + \mathbf{w}_k \quad (17)$$

where $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{W}_k)$ is assumed to be additive white Gaussian. To perform IMU propagation in the SRF, we augment the inertial state by padding the new state at the top via stochastic cloning, i.e., $[\mathbf{x}_{I_{k+1}}^{\top} \quad \mathbf{x}_{I_k}^{\top}]^{\top}$, and propagate the corresponding square-root covariance via the standard QR [see (8)]:

$$\begin{bmatrix} \mathbf{W}_k^{\frac{1}{2}} & \mathbf{0} \\ \mathbf{U}_k \Phi_k^{\top} & \mathbf{U}_k \end{bmatrix} \stackrel{\text{QR}}{=} \mathbf{Q} \begin{bmatrix} \mathbf{U}_{k+1} & \mathbf{U}_{k,1} \\ \mathbf{0} & \mathbf{U}_{k,2} \end{bmatrix} \quad (18)$$

where \mathbf{U}_k and \mathbf{U}_{k+1} are the square-root covariance corresponding to the \mathbf{x}_{I_k} and $\mathbf{x}_{I_{k+1}}$, respectively.

We are now to marginalize certain states such as the oldest state from the square-root covariance. As shown in Figure 2, we first remove the columns of $\begin{bmatrix} \mathbf{U}_{k+1} & \mathbf{U}_{k,1} \\ \mathbf{0} & \mathbf{U}_{k,2} \end{bmatrix}$ corresponding to the marginalized states and form \mathbf{U}' , and then perform QR of \mathbf{U}' to obtain the upper-triangular square-root covariance \mathbf{U}_R :

$$\mathbf{U}' \stackrel{\text{QR}}{=} \mathbf{Q}_R \begin{bmatrix} \mathbf{U}_R \\ \mathbf{0} \end{bmatrix} \quad (19)$$

It is important to note that thanks to our special state ordering by placing non-marginalized variables at the top (e.g., \mathbf{x}_I and \mathbf{x}_{cb}) and those to be marginalized (e.g., features) at the bottom, the resulting square-root covariance can be directly obtained in an upper-triangular form *without* performing QR (19), thus leading to significant computation savings.

B. Measurement Update

The camera bearing measurement of the feature at timestamp t_k can be related to the state and formulate the measurement function. Note features \mathbf{f}_i (See Eq. (15)) can vary in parameterizations, like global/anchor 3D position or inverse depth; we utilize feature models as in [48]:

$$\mathbf{z}_k = h_d(h_p(C^k \mathbf{p}_f), \zeta) + \mathbf{n}_k \quad (20)$$

where \mathbf{z}_k is the raw uv pixel measurement, \mathbf{n}_k is the measurement noise. h_d and h_p are the intrinsic distortion and projection functions, respectively. h_p can support any camera model (e.g., radial-tangential and equidistant [49]). $C^k \mathbf{p}_f$ can be related to any feature parameterizations using the corresponding transformation function (see [48] for more detail). Linearizing (20) yields (note that we here drop off the time index to simplify notations):

$$\mathbf{r} = \mathbf{H}_I \tilde{\mathbf{x}}_I + \mathbf{H}_{cb} \tilde{\mathbf{x}}_{cb} + \mathbf{H}_f \tilde{\mathbf{x}}_f + \mathbf{n} := \mathbf{H}_x \tilde{\mathbf{x}}_x + \mathbf{H}_f \tilde{\mathbf{x}}_f + \mathbf{n}_s \quad (21)$$

where \mathbf{H}_I , \mathbf{H}_{cb} , and \mathbf{H}_f are the measurement Jacobians related to IMU, calibration, and feature states. To balance efficiency and accuracy, short-tracked features are processed as MSCKF features, whereas long-lived SLAM features are included in the state vector [6].

If measurements are corresponding to the MSCKF feature \mathbf{f}_m we project the linearized measurement function onto the left nullspace \mathbf{N} of the feature Jacobian \mathbf{H}_f , to remove the feature dependency, as in the MSCKF [5]:

$$\mathbf{r}_m := \mathbf{N}^\top \mathbf{r} = \mathbf{N}^\top \mathbf{H}_x \tilde{\mathbf{x}} + \mathbf{N}^\top \mathbf{n} =: \mathbf{H}_x^m \tilde{\mathbf{x}}_x + \mathbf{n}_m \quad (22)$$

If measurements are corresponding to a SLAM feature \mathbf{f}_o which has been initialized, the residual (21) is re-written as:

$$\mathbf{r}_s = \mathbf{H}_x^s \tilde{\mathbf{x}}_x + \mathbf{H}_{f_o}^s \tilde{\mathbf{x}}_{f_o} + \mathbf{n}_s \quad (23)$$

1) *Delayed Initialization of SLAM Features:* We now initialize a new SLAM feature \mathbf{x}_{f_N} (instead of \mathbf{x}_{f_o}) given a set of measurements as in (21). We first perform P-QR to compress \mathbf{H}_f^s to obtain the lower triangular matrix \mathbf{H}'_{f_2} :

$$\mathbf{H}_f^s \stackrel{\text{P-QR}}{=} \begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \mathbf{H}'_{f_2} \end{bmatrix} \quad (24)$$

Multiplying (21) by $\begin{bmatrix} \mathbf{Q}_1 & \mathbf{Q}_2 \end{bmatrix}^\top$ yields:

$$\begin{bmatrix} \mathbf{Q}_1^\top \\ \mathbf{Q}_2^\top \end{bmatrix} \mathbf{r} = \begin{bmatrix} \mathbf{0} & \mathbf{H}'_x \\ \mathbf{H}'_{f_2} & \mathbf{H}'_{x_2} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{x}}_{f_N} \\ \tilde{\mathbf{x}}_x \end{bmatrix} + \begin{bmatrix} \mathbf{Q}_1^\top \\ \mathbf{Q}_2^\top \end{bmatrix} \mathbf{n}_s \quad (25)$$

$$\Rightarrow \begin{bmatrix} \mathbf{r}_i \\ \mathbf{r}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{H}'_{f_2} \end{bmatrix} \tilde{\mathbf{x}}_{f_N} + \begin{bmatrix} \mathbf{H}'_x \\ \mathbf{H}'_{x_2} \end{bmatrix} \tilde{\mathbf{x}}_x + \begin{bmatrix} \mathbf{n}_i \\ \mathbf{n}_2 \end{bmatrix} \quad (26)$$

We now can efficiently initialize the square-root covariance with the new feature being included in the state based on the bottom linear system \mathbf{r}'_2 of (26) as follows:

$$\mathbf{U}' = \begin{bmatrix} \mathbf{U} & \mathbf{U} \mathbf{H}'_{x_2}{}^{-\top} \mathbf{H}'_{f_2}{}^{-\top} \\ \mathbf{0} & \mathbf{R}'^{\frac{1}{2}} \mathbf{H}'_{f_2}{}^{-\top} \end{bmatrix} \quad (27)$$

where $\mathbf{R}' = \mathbb{E}[\mathbf{n}'_2 \mathbf{n}'_2{}^\top]$. The top linear system of (26) is used for SRF update as normal measurements.

2) *Outlier Rejection:* Mahalanobis distance test has to be used in practice in order to reject outliers, which is computed in the SRF: $d_m := \mathbf{r}^\top (\mathbf{H} \mathbf{U}^\top \mathbf{U} \mathbf{H}^\top + \mathbf{R})^{-1} \mathbf{r}$. As the measurements of the MSCKF features \mathbf{r}_m and SLAM feature initialization update \mathbf{r}_i are not related to features (i.e., $\mathbf{H} = [\mathbf{H}_x \ \mathbf{H}_f] = [\mathbf{H}_x \ \mathbf{0}]$), we can compute:

$$\mathbf{U} \mathbf{H}^\top = \begin{bmatrix} \mathbf{U}_1 & \mathbf{U}_2 \\ \mathbf{0} & \mathbf{U}_3 \end{bmatrix} \begin{bmatrix} \mathbf{H}_x^\top \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_1 \mathbf{H}_x^\top \\ \mathbf{0} \end{bmatrix} \quad (28)$$

Clearly, given the upper-triangular structure of \mathbf{U} and the unique structure of the measurement Jacobian, we only need to compute $\mathbf{U}_1 \mathbf{H}_x^\top$, instead of multiplying the measurement Jacobian \mathbf{H}_x^\top with the full \mathbf{U} . For the SLAM feature update measurement \mathbf{r}_s , the sparsity of the measurement Jacobian which only relates to the corresponding IMU pose and feature allows us to leverage the upper-triangular structure of \mathbf{U} to compute d_m more efficiently. Note that the computed $\mathbf{U}_1 \mathbf{H}_x^\top$ can be used in the update to avoid redundant computation.

3) *SRF Update:* We perform batch update using all the MSCKF feature measurements \mathbf{r}_m (22), SLAM feature initialization \mathbf{r}_i (26), and SLAM feature measurements \mathbf{r}_s (23). The stacked measurements are given by:

$$\begin{bmatrix} \mathbf{r}_m \\ \mathbf{r}_i \\ \mathbf{r}_s \end{bmatrix} = \begin{bmatrix} \mathbf{H}_x^m \\ \mathbf{H}_x^i \\ \mathbf{H}_x^s \end{bmatrix} \tilde{\mathbf{x}}_x + \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{H}_{f_o}^s \end{bmatrix} \tilde{\mathbf{x}}_{f_o} + \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \tilde{\mathbf{x}}_{f_N} + \begin{bmatrix} \mathbf{n}_m \\ \mathbf{n}_i \\ \mathbf{n}_s \end{bmatrix} \\ = [\mathbf{H}_x \ \mathbf{H}_{f_o} \ \mathbf{0}] \begin{bmatrix} \tilde{\mathbf{x}}_x^\top & \tilde{\mathbf{x}}_{f_o}^\top & \tilde{\mathbf{x}}_{f_N}^\top \end{bmatrix}^\top + \mathbf{n} \quad (29)$$

With this, we perform the SRF update as in (9) and (10). Note that the above measurement does not depend on the new SLAM feature and the corresponding Jacobian is zero, which can be leveraged to make the update even more efficient.

Algorithm 1 SR-VINS

Propagation and Cloning: Propagate the IMU state while cloning the latest IMU pose [Eq. (18)] (*skip QR*)

Marginalization: Marginalize oldest clone and lost tracked SLAM features [Eq. (19)] (**QR**)

Measurements Formulation: Using the tracked features to formulate measurements and prepare for updates.

- MSCKF features via nullspace projection [Eq. (22)]
- SLAM feature initialization [Eq. (24),(26),(27)]
- SLAM features re-observation [Eq. (23)]

SRF update:

- Stack meas. [Eq. (29)] and do SRF update [Eq.(11)] (**P-QR**).
-

At this point, we have presented the main steps of the proposed SR-VINS as summarized in Algorithm 1. Note that the order of the state variables [see (12)] is especially tailored to speed up marginalization. For example, \mathbf{x}_I and \mathbf{x}_{cb} are prioritized at the top as they would not be marginalized, while the clones \mathbf{x}_C are ordered from the latest to oldest

TABLE II: Simulation parameters and prior standard deviations for measurement perturbations.

Parameter	Value	Parameter	Value
Gyro. White Noise	2.0e-4	Gyro. Rand. Walk	2.0e-5
Accel. White Noise	5.0e-4	Accel. Rand. Walk	4.0e-4
Cam Freq. (Hz)	10	IMU Freq. (Hz)	400
Num. Clones	11	Tracked Feat.	100
Max. MSCKF Feat.	40	Max. SLAM Feat.	50

TABLE III: RMSE values for orientation (deg.) and position (m) based on 200 runs on UD-ARL with different estimators.

Methods	EKF	SRF	SRIF
double	0.957 / 0.146	0.957 / 0.146	0.957 / 0.146
float	0.960 / 0.146	0.959 / 0.146	1.045 / 0.174

for easy marginalization of the oldest one. The feature state, \mathbf{x}_f , is placed at the end, because: (i) SLAM features are marginalized frequently, (ii) this ordering makes the upper-triangular structure $[\mathbf{M}''$ in (3), see Figure 1] better preserved during P-QR when performing update, and (iii) it ensures \mathbf{U} is still upper-triangular after initializing a new SLAM feature. It is also important to note that as shown in Algorithm 1, the QR and P-QR decomposition is utilized in propagation, state marginalization, and update. However, these can be sped up by skipping QR in propagation and combined with marginalization. Due to space constraints, detailed derivations and functionalities such as online calibration can be found in our companion technical report [47].

IV. NUMERICAL STUDY

We use a 30-minute, 2.4km UD-ARL trajectory (see Figure 3) and the OpenVINS simulator [48] to produce realistic visual bearings and inertial measurements, as detailed in Table II. For a fair comparison across estimators, we build upon OpenVINS which utilizes EKF. We implemented float version of OpenVINS (EKF), the square-root inverse filter (SRIF), and the proposed SRF-based VINS (SR-VINS). In Figure 4, the top two plots illustrate the orientation and position errors across different estimators with both double and float. Meanwhile, the bottom plot depicts the condition number of the square root information matrix and the condition number of the \mathbf{F} matrix for SRF over time. Table III reports the average Root Mean Square Error (RMSE) for different estimators over 200 Monte-Carlo runs.

Given the covariance matrix \mathbf{P} , the square-root information matrix \mathbf{R} is given by: $\mathbf{R}^\top \mathbf{R} = \mathbf{P}^{-1}$. From the figure, we observe as the condition number of \mathbf{R} grows larger than $2e^5$, both orientation and position errors of SRIF(f) start showing a degraded performance compared to other filter design methods. This can also be seen in Table III, the float SRIF is inaccurate with large RMSE values. This is likely due to the numerical issue when performing inversion on ill-conditioned \mathbf{R} to solve for state update under limited machine precision (see Chapter 3.5.1 in [22]).

In contrast, the covariance-form estimators, both EKF and the proposed SRF, demonstrated consistent performance regardless of using double or float. This is evident from the comparable RMSE values in Table III, as well as the consistent error trends in Figure 4. When performing the SRF update, the inversion of \mathbf{F} is supposed to be the most




Fig. 3: Simulated 2.4km UD-ARL trajectory.




Fig. 4: **Top:** Orientation/position errors of different estimators performed on UD-ARL dataset. ‘d’ is for double; ‘f’ is for float. While most estimators perform similarly and are hard to distinguish from the plot, SRIF(f) shows a clear drop in accuracy over time. **Bottom:** The condition number of the square-root information matrix (black line) with that of the P-QR lower triangular matrix \mathbf{F} (red line, see Eq.(11)).

numerically challenging operation. We thus plot its condition number shown in Figure 4 (bottom). Its condition number is shown to be stable and close to 1, demonstrating the improved numerical stability of the proposed SRF. Intuitively speaking, $\mathbf{F}^{-\top}$ is the transition matrix between $\mathbf{U}_{k|k-1}$ and $\mathbf{U}_{k|k}$. Therefore, as long as the measurements used in the update are not extremely accurate compared with the propagated estimation, we would expect $\mathbf{F}^{-\top}$ to be close to an identity matrix and be well-conditioned, which is almost always the case in the VINS in practice.

V. REAL-WORLD EXPERIMENTS

We further evaluate the proposed SR-VINS on the EuRoC MAV dataset [50]. Only the left camera is used during the evaluation. The proposed system is built on top of OpenVINS [48]. We use the same default setup as OpenVINS [48], which extracts 200 sparse point features, keeps 11 clones, uses at most 50 SLAM features and 40 MSCKF features, performing camera-IMU extrinsic, time offset, and camera intrinsic calibration online. The proposed system is tested with both double and float versions, denoted as SRF(d) and

TABLE IV: Average Absolute Trajectory Error (ATE) in degrees/meters. ‘d’ and ‘f’ indicate the use of double and float. SRF(M) utilizes float and MSCKF features only for comparison with RVIO2.

Algo.	V101	V102	V103	V201	V202	V203	MH01	MH02	MH03	MH04	MH05
EKF(d)	0.70 / 0.06	1.67 / 0.06	2.88 / 0.07	0.95 / 0.10	1.38 / 0.06	1.28 / 0.14	1.74 / 0.10	0.91 / 0.17	1.14 / 0.12	0.95 / 0.25	1.03 / 0.41
EKF(f)	0.71 / 0.06	1.66 / 0.06	2.87 / 0.06	0.94 / 0.10	1.40 / 0.06	1.25 / 0.14	1.76 / 0.10	0.91 / 0.17	1.18 / 0.13	0.94 / 0.25	1.04 / 0.41
SRF(d)	0.68 / 0.05	1.68 / 0.06	2.88 / 0.06	0.99 / 0.11	1.40 / 0.06	1.28 / 0.14	1.75 / 0.10	0.93 / 0.18	1.18 / 0.12	0.95 / 0.24	1.04 / 0.39
SRF(f)	0.66 / 0.05	1.68 / 0.06	2.88 / 0.06	0.98 / 0.11	1.39 / 0.06	1.27 / 0.14	1.76 / 0.10	0.93 / 0.17	1.14 / 0.12	0.93 / 0.23	1.05 / 0.40
SRF(M)	0.63 / 0.08	1.75 / 0.06	1.76 / 0.08	0.74 / 0.10	1.36 / 0.08	1.19 / 0.16	1.56 / 0.15	0.95 / 0.22	1.02 / 0.17	1.12 / 0.25	0.93 / 0.39
RVIO2	0.88 / 0.09	2.27 / 0.10	2.02 / 0.10	2.19 / 0.13	1.90 / 0.11	1.50 / 0.15	2.60 / 0.17	1.00 / 0.15	1.08 / 0.19	1.10 / 0.24	0.95 / 0.32
VINS-Mono	0.82 / 0.07	2.74 / 0.10	5.15 / 0.15	2.13 / 0.09	2.57 / 0.13	3.43 / 0.29	0.78 / 0.20	0.86 / 0.18	1.82 / 0.23	2.51 / 0.41	0.94 / 0.29

TABLE V: Average estimator run time (ms) comparison (excluding feature tracking) on EuRocMAV dataset. SRF(M) means adopting the same clone size and using only MSCKF features similar to the default setup of RVIO2 (15 clones, track 200 features, all the features are processed as MSCKF features once they lose track or reach maximum clone size).

Algorithm	EKF	SRF	SRF(M)	RVIO2	VINS-Mono
Double	4.2	2.8	1.7	-	22.4
Float	3.0	2.2	1.1	1.8	-

SRF(f), respectively.²

We compare the proposed SR-VINS with the baseline Open-VINS (EKF(d)), which is originally in double, a float version of Open-VINS (EKF(f)) is also developed and evaluated. With default prior, float Open-VINS will experience negative diagonals in the covariance matrix and diverge in some sequences, thus its prior is tuned to make sure it runs on all the sequences. To make a fair comparison, all the versions of SRF and EKF use the same prior. We also compare with the open-sourced RVIO2 [27], which is a square-root inverse filter VIO based on robocentric state formulation, and VINS-Mono [14], which is an optimization-based sliding window VIO system. It is worth mentioning that EqVIO [8], [11] can also achieve impressive computational efficiency, however, its computation efficiency is gained by a much smaller state size, which is unable to fairly compare because its design principle is different from MSCKF-based VINS, thus, we do not include. The averaged Absolute Trajectory Error (ATE) values are reported in Table IV. Since RVIO2 only uses MSCKF features by default (i.e., no long-track SLAM features are maintained in the state vector, thus having a much smaller state size), we also report the float SRF performance with similar config (keeps 15 clones, tracks 200 features, all the features are processed as MSCKF features), denoted as SRF(M) in Table IV, for a fair comparison.

From Table IV, we can see that the performance of SRF(d), SRF(f), EKF(d), and EKF(f) are very similar as expected. The performance of double and float, EKF, and SRF are not exactly the same in the real world due to two reasons. First, the χ^2 test is adopted to reject outliers and robustify the estimator and might introduce randomness. For example, in certain cases, SRF(d) might reject measurements that pass χ^2 test in SRF(f) because of slight numerical differences, this will cause SRF in different versions to use different measurements and have different performance. Second, OpenVINS (EKF) performs a “sequential” update, which first processes MSCKF features and then SLAM fea-

tures for the consideration of efficiency, while SRF performs the update all at once. This also introduces differences in the state linearization points. Compared with RVIO2 and VINS-mono, SRF also achieves superior performance in almost all the sequences. Surprisingly, even SRF(M) achieves similar or even better performance than the other systems.

The efficiency of the estimators is also evaluated and reported in Table V. Clearly, SRF(f) is much faster than its baseline EKF(d), reducing the runtime almost by half. Regardless of being in double or float format, SRF consistently prevails over EKF. Remarkably, the double precision SRF even outperforms the float EKF. VINS-Mono runs the slowest as it performs iterative optimization. RVIO2 is also developed in float and shows excellent efficiency, but with a similar setup, SRF(M) in double prevails. Finally, SRF(M) achieves the best efficiency with 1.1 ms in estimator runs, which means it can run over 900Hz, especially suitable for running on a computation-constrained platform. The efficiency gain of SRF mainly comes from the proposed QR-based SRF update method, fully explored problem structure (state order, upper-triangular covariance, Jacobian structure, reusable computation).

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we have developed the first square-root filter (SRF)-based VINS (i.e., SR-VINS) which significantly improves both the numerical stability and efficiency. We strongly advocate that the SRF is ideal for VINS due to its ability to represent a broader dynamic range, guarantee the property of the covariance matrix, reduce the memory requirement for covariance, and improve numerical stability. However, it is not trivial to capitalize on these advantages because of the challenge of its update inefficiencies, especially in dealing with large measurement sizes. To overcome this issue and leverage the numerical advantage and the structure of the square-root covariance matrix, we have developed a novel permuted QR (P-QR)-based SRF update method. With this, we fully exploit the structure of the VINS problem to best utilize the upper triangular square-root covariance to gain speed boost. From our comprehensive numerical studies and real-world experiments, we have shown that the proposed SR-VINS can run robustly in float, gaining significant speedup (around 2 times faster than the SOTA filters), while exhibiting no accuracy loss, which makes it especially suitable for edging computing platforms. In the future, we are interested in further improving efficiency in visual tracking by leveraging the covariance matrix to reduce search space.

²All computational results were performed in a single thread on an Intel(R) Core(TM) i7-11800H @ 2.30GHz.

REFERENCES

- [1] G. Huang, "Visual-inertial navigation: A concise review," in *Proc. International Conference on Robotics and Automation*, Montreal, Canada, May 2019.
- [2] C. Chen, Y. Yang, P. Geneva, W. Lee, and G. Huang, "Visual-inertial-aided online mav system identification," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Kyoto, Japan., 2022.
- [3] C. Chen, P. Geneva, Y. Peng, W. Lee, and G. Huang, "Monocular visual-inertial odometry with planar regularities," in *Proc. of the IEEE International Conference on Robotics and Automation*, London, UK., 2023.
- [4] Y. Peng, C. Chen, and G. Huang, "Quantized visual-inertial odometry," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [5] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint kalman filter for vision-aided inertial navigation," in *Proceedings 2007 IEEE international conference on robotics and automation*. IEEE, 2007, pp. 3565–3572.
- [6] M. Li and A. I. Mourikis, "High-precision, consistent ekf-based visual-inertial odometry," *The International Journal of Robotics Research*, vol. 32, no. 6, pp. 690–711, 2013.
- [7] P. Geneva, J. Maley, and G. Huang, "An efficient schmidt-ekf for 3D visual-inertial SLAM," in *Proc. Conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, CA, June 2019.
- [8] P. van Goor and R. Mahony, "An equivariant filter for visual inertial odometry," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14432–14438.
- [9] Y. Yang, C. Chen, W. Lee, and G. Huang, "Decoupled right invariant error states for consistent visual-inertial navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1627–1634, 2022.
- [10] C. Chen, Y. Yang, P. Geneva, and G. Huang, "FEJ2: A consistent visual-inertial state estimator design," in *International Conference on Robotics and Automation (ICRA)*, Philadelphia, USA, 2022.
- [11] P. van Goor and R. Mahony, "Eqvio: An equivariant filter for visual-inertial odometry," *IEEE Transactions on Robotics*, 2023.
- [12] G. Huang, A. I. Mourikis, and S. I. Roumeliotis, "An observability constrained sliding window filter for SLAM," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, CA, Sept. 2011, pp. 65–72.
- [13] 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.
- [14] 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.
- [15] L. Von Stumberg, V. Usenko, and D. Cremers, "Direct sparse visual-inertial odometry using dynamic marginalization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 2510–2517.
- [16] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.
- [17] C. Chen, P. Geneva, Y. Peng, W. Lee, and G. Huang, "Optimization-based vins: Consistency, marginalization, and fej," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [18] C. Chen, Y. Peng, and G. Huang, "Fast and consistent covariance recovery for sliding-window optimization-based vins," in *Proc. International Conference on Robotics and Automation*, Yokohama, Japan, May 2024.
- [19] D. G. Kottas and S. I. Roumeliotis, "An iterative kalman smoother for robust 3d localization on mobile and wearable devices," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 6336–6343.
- [20] K. Wu, A. M. Ahmed, G. A. Georgiou, and S. I. Roumeliotis, "A square root inverse filter for efficient vision-aided inertial navigation on mobile devices," in *Robotics: Science and Systems*, vol. 2. Rome, Italy, 2015, p. 2.
- [21] M. J. Flynn, "Very high-speed computing systems," *Proceedings of the IEEE*, vol. 54, no. 12, pp. 1901–1909, 1966.
- [22] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU press, 2013.
- [23] K. J. Wu and S. I. Roumeliotis, "Inverse schmidt estimators," *Multiple Autonomous Robotic System Laboratory, Department of Computer Science & Engineering, University of Minnesota, Tech. Rep. Number 2016-003*, 2016.
- [24] D. Caruso, A. Eudes, M. Sanfourche, D. Vissiere, and G. Le Besnerais, "An inverse square root filter for robust indoor/outdoor magneto-visual-inertial odometry," in *2017 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2017, pp. 1–8.
- [25] T. Ke, K. J. Wu, and S. I. Roumeliotis, "Rise-slam: A resource-aware inverse schmidt estimator for slam," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 354–361.
- [26] Z. Huai and G. Huang, "Markov parallel tracking and mapping for probabilistic slam," in *Proc. of the IEEE International Conference on Robotics and Automation*, Xi'an, China, 2021.
- [27] —, "Square-root robocentric visual-inertial odometry with online spatiotemporal calibration," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9961–9968, 2022.
- [28] F. Dellaert and M. Kaess, "Square root sam: Simultaneous localization and mapping via square root information smoothing," *The International Journal of Robotics Research*, vol. 25, no. 12, pp. 1181–1203, 2006.
- [29] M. Kaess, A. Ranganathan, and F. Dellaert, "isam: Fast incremental smoothing and mapping with efficient data association," in *Proceedings 2007 IEEE international conference on robotics and automation*. IEEE, 2007, pp. 1670–1677.
- [30] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "isam2: Incremental smoothing and mapping using the bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [31] N. Demmel, D. Schubert, C. Sommer, D. Cremers, and V. Usenko, "Square root marginalization for sliding-window bundle adjustment," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 13260–13268.
- [32] M. W. Givens and J. W. McMahan, "Square-root extended information filter for visual-inertial odometry for planetary landing," *Journal of Guidance, Control, and Dynamics*, vol. 46, no. 2, pp. 231–245, 2023.
- [33] P. Chauchat, A. Barrau, and S. Bonnabel, "Factor graph-based smoothing without matrix inversion for highly precise localization," *IEEE Transactions on Control Systems Technology*, vol. 29, no. 3, pp. 1219–1232, 2020.
- [34] P. Chauchat, S. Bonnabel, and A. Barrau, "Invariant smoothing with low process noise," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 4758–4763.
- [35] J. Potter and R. Stern, "Statistical filtering of space navigation measurements," in *Guidance and Control Conference*, 1963, p. 333.
- [36] R. H. Battin, "Astronautical guidance," 1964.
- [37] P. Dyer and S. McReynolds, "Extension of square-root filtering to include process noise," *Journal of Optimization Theory and Applications*, vol. 3, pp. 444–458, 1969.
- [38] J. Bellantoni and K. Dodge, "A square root formulation of the kalman-schmidt filter," *AIAA journal*, vol. 5, no. 7, pp. 1309–1314, 1967.
- [39] A. Andrews, "A square root formulation of the kalman covariance equations," *Aiaa Journal*, vol. 6, no. 6, pp. 1165–1166, 1968.
- [40] P. Kaminski, A. Bryson, and S. Schmidt, "Discrete square root filtering: A survey of current techniques," *IEEE Transactions on automatic control*, vol. 16, no. 6, pp. 727–736, 1971.
- [41] W. S. Agee and R. H. Turner, "Triangular decomposition of a positive definite matrix plus a symmetric dyad with application to kalman filtering," *White Sands Missile Range Tech. Rep.*, vol. 38, 1972.
- [42] N. A. Carlson, "Fast triangular formulation of the square root filter," *AIAA journal*, vol. 11, no. 9, pp. 1259–1265, 1973.
- [43] J. J. Dongarra, J. Du Croz, S. Hammarling, and I. S. Duff, "A set of level 3 basic linear algebra subprograms," *ACM Transactions on Mathematical Software (TOMS)*, vol. 16, no. 1, pp. 1–17, 1990.
- [44] P. S. Maybeck, *Stochastic models, estimation, and control*. Academic press, 1982.
- [45] W. Givens, "Numerical computation of the characteristic values of a real symmetric matrix," Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), Tech. Rep., 1954.
- [46] A. S. Householder, *The theory of matrices in numerical analysis*. Courier Corporation, 2013.
- [47] Y. Peng, C. Chen, and G. Huang, "Technical report: Ultrafast square-root filter-based vins," University of Delaware, Tech. Rep., 2024. [Online]. Available: https://udel.edu/~ghuang/papers/tr_srf.pdf
- [48] P. Geneva, K. Ekenhoff, W. Lee, Y. Yang, and G. Huang, "Openvins: A research platform for visual-inertial estimation," in *Proc. of the IEEE International Conference on Robotics and Automation*, Paris, France, 2020. [Online]. Available: https://github.com/rpng/open_vins
- [49] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [50] M. Burri, J. Nikolic, P. Gohl, T. Schneider, J. Rehder, S. Omari, M. W. Achtelik, and R. Siegwart, "The euroc micro aerial vehicle datasets," *The International Journal of Robotics Research*, vol. 35, no. 10, pp. 1157–1163, 2016.