Consistent Covariance Pre-Integration for Invariant Filters with Delayed Measurements

Eren Allak[1], Alessandro Fornasier[1], and Stephan Weiss[1]

Abstract— Sensor fusion systems merging (multiple) delayed sensor signals through a statistical approach are challenging setups, particularly for resource constrained platforms. For statistical consistency, one would be required to keep an appropriate history, apply the correcting signal at the given time stamp in the past, and re-apply all information received until the present time. This *re-calculation* becomes impractical (the bottleneck being the re-propagation of the covariance matrices for estimator consistency) for platforms with multiple sensors/states and low compute power.

This work presents a novel approach for consistent covariance pre-integration allowing delayed sensor signals to be incorporated in a *statistically consistent fashion with very low complexity*. We leverage recent insights in Invariant Extended Kalman Filters (IEKF) and their log-linear, state independent error propagation together with insights from the scattering theory to mimic the re-calculation process as a medium through which we can propagate waves (covariance information in this case) in single operation steps. We support our findings in simulation and with real data.

I. INTRODUCTION

Sensor fusion for robot state estimation is a well studied field and very active since many years. Nonlinear optimization techniques and (Extended) Kalman Filters ((E)KF) are the primary tools for state of the art approaches with the latter being the preferred method on computationally constrained platforms. For such platforms, delayed sensor signals are particularly cumbersome in state of the art approaches. Non-linear optimization methods inherently have the setup to seamlessly include such delayed signals at the cost of generally be at the higher end of computational complexity. For filter based approaches, delayed signals trigger a high computational load or require approximations with negative effects on the estimator performance and consistency: The incoming signal needs to be applied in the past from where the Markov chain needs to be re-built (*re-computation step*) until the present time.

Naively re-computing the information leads to a noticeable bottleneck in the covariance re-propagation steps ($\mathbf{P}_{k+1} = \mathbf{F} * \mathbf{P}_k * \mathbf{F}^T + \mathbf{Q}$ with $\mathbf{P}$ being the state covariance matrix, $\mathbf{F}$ the Jacobian of the state dynamics, and $\mathbf{Q}$ the process noise matrix) requiring the multiplication of three square matrices each of the size of the state vector dimensions. Due to the non-linear nature of most robotic systems, existing approaches reducing the load of this re-computation make assumptions or simplifications harming the consistency of the estimator. Most relevant to this work, the authors in [1] presented an approach for a one-step covariance pre-integration until the next measurement based on an analogy from the scattering theory and the Star-Product (SP-EKF). While heavily reducing comutational cost (about 90% less computation cost according to [1]), the only assumption in that work is that the re-linearization of the state dynamics ($\mathbf{F}$) after a delayed update has little effect on the performance and consistency, and can thus be omitted. This assumption may hold in many scenarios. In this work, we show that specifically for initial motion estimation and badly initialized estimators, this assumption does not hold. More generally, in such cases, EKF based approaches heavily suffer from inconsistency and the SP-EKF extension for lower computational complexity on delayed measurements further deteriorate the consistency significantly by re-using heavily wrong linearization points. Motivated by the significant reduction in computational complexity of the SP-EKF, we extend this approach to the Invariant EKF (IEKF) model. The log-linearity and linearization point independence compared to the EKF are key to our proposed extension in this paper combining the Star-Product with the IEKF approach. Not depending on the linearization points removes the weakness the SP-EKF approach discussed above, rendering our SP-IEKF approach both efficient and consistent. More specifically, our contributions are:

- Theory on consistent covariance pre-integration and dependence on linearization points.
- Thorough analysis of (in)consistencies of covariance pre-integration in relation to EKF and IEKF.
- Use case scenario (simulation and real data) proposing a method for consistent *and* efficient fusion of delayed measurements in an SP-IEKF framework.

In the following Section II, we review the related work. Section III-A briefly recaps previous work on SP-EKF. New findings regarding the inconsistency of SP-EKF its analysis, are given in Section III-B. By removing the estimate-dependency of the linearization points we propose a novel, consistent SP-IEKF based on invariant filtering in Section III-C. Sections IV and V finally show the results in simulations and with real data respectively.

[1]Eren Allak, Alessandro Fornasier and Stephan Weiss are with the Control of Networked Systems Group, Universität Klagenfurt, Austria {eren.allak, alessandro.fornasier, stephan.weiss}@ieee.org

II. Related Work

Techniques for estimator complexity reduction is well discussed in the community. Often, approaches for filter-based delay compensation assume the delay time to be known (e.g. when using constant rate sensors and constant delay channels) which allows for efficient state cloning [2] or filter duplication [3]. If this assumption is removed, as in our case, such techniques become intractable. The overhead to compensate for unknown time delay is then much larger than the gain coming from these approaches.

Not for delay compensation, but for relative pose updates stochastic cloning was presented in [4] and was extended to IEKF setups in [5]. The latter approach combined with scattering theory makes it possible to combine very efficiently and consistently arbitrarily many *relative* and propagation measurements into a *single* measurement for filtering frameworks which constitutes another novelty of this work with high potential use cases.

Pre-integration techniques, to our knowledge first presented in [6], are other methods to overcome computational complexity. The pre-integration of inertial measurements leads to considerable complexity reduction and increased robustness in non-linear optimization [6], [7]. However, these measurements form relative motion constraints and are not applied to statistical filters. Furthermore, propagation measurements cannot be combined with update measurements into a single constraint rendering the handling of multiple delayed sensors cumbersome and expensive.

The mitigation of these shortcomings and the possibility to apply covariance pre-integration in filter based systems was presented in [1]. The work makes use of the scattering theory that was first used to describe the propagation of waves through media in physics [8]. Only much later, a striking connection to state estimation was discovered and is presented in [9], which is based on the work [10], [11]. Through this connection, a significant complexity reduction for the fusion of delayed measurements in multi-sensor filter based estimators was proposed in [1]. The reduction becomes particularly apparent with high-rate system propagation as the covariance propagation steps were identified as computational bottlenecks. The scattering theory provides the background for this method to be used on *linear* systems. Although a thorough analysis of the complexity reduction and accuracy was done in [1], the analysis of how this approach affects the *consistency* of the estimator in *non-linear* systems was ignored. This is particularly relevant when linearization points significantly differ from their true values.

The consistency analysis is at the core of this work and uses tools also used by the SLAM community. For example, the ANEES and the ideal EKF are used to cancel the effects of linearization and examine the consistency of EKF-SLAM [12], [13].

Apart of the consistency analysis, this work further focuses on overcoming the linearization point issue in previous work. This is achieved by transforming the scattering theory based covariance pre-integration method into an IEKF framework. The invariant EKF is an alternative to EKF based on error variables which are invariant by the action of a group of transformations. It was first proposed in the context of invariant dynamical systems [14] until the framework of linear observed systems on groups [15], [16], based on group-affine dynamics, generalized the method and brought strong properties such as *log-linear error propagation* and EKF stability [15], or pre-integrability [16]. The idea of using a non-linear error variable for a non-linear system was also generalized as a tool for the design of EKF-like methods preserving the observability and reachability properties of non-linear systems [17], [18], [19], with a specific success encountered in SLAM applications [17], [20] but also in more unexpected fields such as legged robots state estimation [21]. This framework is not restricted to EKF and was then leveraged to design invariant smoothers [22] and invariant H-infinity filters [23].

In this work, the novel combination of the IEKF framework with the scattering theory based covariance pre-integration leads to a consistent filter based estimator with unprecedented efficiency in case of delayed measurements.

III. Consistent and Efficient Fusion of Delayed Measurements

A. Fusion of Delayed Measurements

A non-linear discrete-time system with delayed measurements and additive white noise $\mathbf{n}_u, \mathbf{n}_y$ is given as follows. The delayed measurement $\mathbf{y}_{k+1}^d$ arrives at current time t_{k+1}, but was measured before at t_d:

$$\mathcal{X}_{k+1} = f\left(\mathcal{X}_k, \mathbf{u}_k\right) + \mathbf{n}_{u,k} \qquad \mathbf{n}_u \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_u) \qquad (1)$$

$$\mathbf{y}_{k+1} = h\left(\mathcal{X}_{k+1}\right) + \mathbf{n}_{y,k+1} \qquad \mathbf{n}_y \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_y) \qquad (2)$$

$$\mathbf{y}_{k+1}^d = h\left(\mathcal{X}_d\right) + \mathbf{n}_{y,d} \qquad t_d < t_{k+1} \qquad (3)$$

The most accurate and consistent, but also most costly, way to fuse this measurement is to *recalculate* the filter. Therefore all recent estimates are discarded until t_d. Then all measurements, including the delayed one $\mathbf{y}_{k+1}^d$, are re-applied in the correct order. This approach causes a significant computation spike, the bulk of it being the covariance computations.

Previous work [1] achieves a similarly accurate result than the recomputation method in terms of RMSE with considerably lower complexity. Through the use of concatenation techniques of scattering theory, many propagation measurements are combined into a single step, effectively enabling covariance pre-integration. A quick recap shown below will cover most of the theory also relevant for this work. The generic definition of the star product Eq. (4) is used to combine measurements for propagation and updates by their respective generators Eq. (5) and Eq. (6) with $\mathbf{F}$ and $\mathbf{H}$ being the Jacobian of the state dynamics and measurements respectively. As a result, a scattering matrix (Eq. 7) is found that can be used to recompute covariances

P in one step to the current time, as shown in Eq. (8).

$$\mathcal{S} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \star \begin{bmatrix} A & B \\ C & D \end{bmatrix}$$
$$= \begin{bmatrix} A(I-bC)^{-1}a & B + Ab(I-Cb)^{-1}D \\ c + dC(I-bC)^{-1}a & d(I-Cb)^{-1}D \end{bmatrix} \quad (4)$$

$$M_{t,k} = \begin{bmatrix} \mathbf{F}_k & \boldsymbol{\Sigma}_u \\ \mathbf{0} & \mathbf{F}_k^T \end{bmatrix} \quad \mathbf{F}_k = \left. \frac{\partial f(\mathcal{X},u)}{\partial \mathcal{X}} \right|_{\mathcal{X}=\hat{\mathcal{X}}_k, u=u_k} \quad (5)$$

$$M_{m,k} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{H}_k^T \boldsymbol{\Sigma}_y^{-1} \mathbf{H}_k & \mathbf{I} \end{bmatrix} \quad \mathbf{H}_k = \left. \frac{\partial h(\mathcal{X})}{\partial \mathcal{X}} \right|_{\mathcal{X}=\hat{\mathcal{X}}_k} \quad (6)$$

$$\mathcal{S}_{k,N}^0 = M_k \star M_{k+1} \star \ldots \star M_N \quad (7)$$

$$\mathcal{S}_{k,N} = \begin{bmatrix} \mathbf{I} & \mathbf{P}_k \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \star \mathcal{S}_{k,N}^0 = \begin{bmatrix} \cdot & \mathcal{P}_{k,N} \\ \cdot & \cdot \end{bmatrix} \quad (8)$$

The scattering matrix $\mathcal{S}_{k,N}$ is the outcome of a Kalman Filtering problem with the initial covariance $\mathbf{P}_k$ and the measurements $\mathbf{y}_k \ldots \mathbf{y}_N$ and $\mathbf{u}_k \ldots \mathbf{u}_N$. $\mathcal{P}_{i,N}$ is the covariance after applying all measurements to the initial covariance $\mathbf{P}_k$. Note that $\mathcal{S}_{k,N}^0$ depends on M_k to M_N which in turn depend on all previous dynamic and measurements Jacobians $\mathbf{F}$ and $\mathbf{H}$ between current time and delayed measurement. The fact that these Jacobians are not re-calculated in this approach at the refined states after the update in the past is the core issue of the inconsistency introduced when used on non-linear systems as explained in the following.

B. Estimation with EKF and Inconsistencies of the SP-EKF

The scattering theory approach was initially only connected to the Kalman Filter. Applying the theory for non-linear problems required the use of the Extended Kalman Filter (EKF). Therefore, to understand the inconsistencies that may arise we have to look at the main concept that is at the core of the EKF, which is linearization. The EKF uses the current estimates as the linearization points (i.e. not the true values). As such, linearization and linearization points have a considerable impact on the performance of the EKF as wrong estimates directly lead to wrong linearization points due to the state dependence of the Jacobians. Mitigating the effect of wrong linearization points is posed as one of the main challenges of this work.

In this work, we use the so called *ideal* EKF as a theoretical tool to evaluate the consistency by canceling effects of linearization. The only difference between the regular and the ideal EKF lies in the covariance computation: while the EKF uses the estimated states and noisy inputs for linearization, the ideal EKF uses the true values for linearization and inputs independent of the estimated states. These changes directly affect the covariance propagation and update steps. With the system of Eq. (1) and Eq. (2) in mind, covariance propagation


Fig. 1. With an wrong initial heading the estimator accumulates pre-computations of the SP-EKF with wrong linearization points (blue, red are true states). This leads to inconsistency when applied for covariance pre-integration. The advantage of efficiency is bought with inconsistency. The present work achieves both high efficiency *and* consistency in spite of wrong linearization points.

and update read as follows:

$$\mathbf{P}_{k+1|k} = \mathbf{F}_k \mathbf{P}_{k|k} \mathbf{F}_k^T + \boldsymbol{\Sigma}_u \quad (9)$$

$$\mathbf{P}_{k+1|k+1} = \mathbf{P}_{k+1|k} \mathbf{H}_{k+1}^T \left(\mathbf{H}_{k+1} \mathbf{P}_{k+1|k} \mathbf{H}_{k+1}^T + \boldsymbol{\Sigma}_y \right) \quad (10)$$

$$\mathbf{F} = \left. \frac{\partial f(\mathcal{X},u)}{\partial \mathcal{X}} \right|_{\mathcal{X}=\hat{\mathcal{X}}_{k+1|k}, u=u_k} \quad \mathbf{H} = \left. \frac{\partial h(\mathcal{X})}{\partial \mathcal{X}} \right|_{\mathcal{X}=\hat{\mathcal{X}}_{k+1|k}} \quad (11)$$

For the EKF, wrong estimates (or linearization points) can directly influence the covariance computations. Unfortunately, the SP-EKF suffers from the same problem as it can be seen in Eq. (5) and Eq. (6). Once computed with wrong estimates, the generators carry on the wrong information into the scattering matrix. Covariance pre-integration as suggested in [1] can get inconsistent at this point.

An illustrative example shown in Fig. 1 highlights the problem with wrong linearization points for the SP-EKF. A robot is estimating the whole trajectory only using relative pose measurements. Because of a wrong initial heading the SP-EKF builds up scattering matrices with wrong estimates leading to inconsistencies once covariance pre-integration is applied after a delayed global update. This inconsistency due to wrong linearization points do not happen for recalculation, because then all Jacobians are re-linearized at refined estimates.

Considering the recalculation to be the ideal method in terms of consistency, any deviation from it can be seen as sub-optimal. Since consistency is measured by the *Average Normalized Estimation Error Squared* (ANEES, [24]), a novel measure for inferior performance regarding consistency can be quantified by:

$$\Delta \text{ANEES}(A) = |\text{ANEES}_r - \text{ANEES}_A| \quad (12)$$

ANEES_r is the ANEES for the recalculation, while ANEES_A is the ANEES for any other algorithm A.

C. Consistent SP-IEKF

The estimate dependency was identified as the main source of inconsistencies for the SP-EKF. Formulating estimate-independent elements of scattering theory and a rigorous analysis of the consistency is hence at the core of this work.

1) Estimate Independence: The basic elements of scattering theory are the generators, which in turn are built out of Jacobians of the system equations and noise covariances, as it can be seen in (5) and (6). In general, following the EKF methodology will result in Jacobians being estimate-dependent and noise matrices being estimate-independent as shown later in an example.

Following an invariant filtering approach [15], the dependencies are often flipped: In that case, the Jacobians are estimate-independent and the noise covariances become estimate-dependent. More precisely, for group affine systems with left or right invariant observations, the Jacobians for the process and the observation are constant and estimate-independent. The limitation for the applicability is only minor since many systems comply with these requirements, including the presented setup with measurements of the 6D pose, velocity, GPS position, known landmarks, magnetometer measurements, and others [16]. Independence of the Jacobians to the linearization points comes with a set of strong properties regarding filter stability [15] and preservation of observability [17], [18] or reachability [19]. This can be seen as a strong argument in favor of the independence of the Jacobians rather than independence of the noise characteristics of linearization points. One of the contributions of the present paper is giving a convincing illustration of this idea in the context of covariance pre-integration, which we hope may inspire future other applications.

2) SP-IEKF for a Wheeled Robot with Odometry, Relative Pose and Landmark Updates: The presented method is demonstrated for a wheeled robot navigating with a wrong initial heading. Only once delayed known landmark measurements arrive, global information will correct the drift. In general, the robot will move by using odometry for propagation and relative pose measurements for updates. Note that drift will be accumulated without known landmark measurements, because position and heading are unobservable otherwise. When global, delayed information becomes available, it will correct a state in the past. To update the state estimate at the current time two options are possible, recalculation or using the scattering matrices as in SP-EKF. The scattering matrices (and generators) are build up and pre-computed by filters as the measurements for propagation and updates come in. These pre-computations are the reason for the efficiency of SP-EKF in covariance computations.

First, the invariant EKF is presented with Jacobians $\mathbf{F}_k$ and $\mathbf{H}_{k+1}$ that are estimate-independent and also shown are noise covariance matrices $\boldsymbol{\Sigma}_{u,SC}$ and $\tilde{\boldsymbol{\Sigma}}_y$, in turn, being estimate-dependent. Both are used to compute the generators. To update with relative poses we need to apply stochastic cloning[4][5]. The system and observation are:

$$\mathcal{X}_{k+1} = \mathcal{X}_k \cdot u_k \cdot n_u \qquad \mathcal{X}, u_k, n_u \in SE(2) \quad (13)$$
$$\mathcal{Y}_{k+1} = n_y \cdot \mathcal{X}_{k+1}^{-1} \cdot \mathcal{X}_{k-N} \quad \mathcal{Y}_{k+1}, n_y \in SE(2) \quad (14)$$
$$n_u = \exp(\lfloor \xi_{n,u} \rfloor^{\wedge}) \qquad \xi_{n,u} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_u) \quad (15)$$
$$n_y = \exp(\lfloor \xi_{n,y} \rfloor^{\wedge}) \qquad \xi_{n,y} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_y) \quad (16)$$

The state consists of the current estimate $\hat{\mathcal{X}}_k$ at time t_k and a clone of the state $\hat{\mathcal{X}}_{k,C}$ at time t_{k-N}. To integrate the cloning into the scattering matrices, a special generator is used with $\mathbf{F}_C = [\mathbf{0}_{3\times3}, \mathbf{I}_{3\times3}; \mathbf{0}_{3\times3}, \mathbf{I}_{3\times3}]$ and $\boldsymbol{\Sigma}_C = \mathbf{0}_{6\times6}$. Until a relative pose update, the estimator is propagated with odometry input u_k:

$$\hat{\mathcal{X}}_{k,SC} = [\hat{\mathcal{X}}_{k,C}; \hat{\mathcal{X}}_k] \quad (17)$$
$$\hat{\mathcal{X}}_{k+1|k} = \hat{\mathcal{X}}_{k|k} \cdot u_k$$
$$\hat{\mathcal{X}}_{k+1|k,C} = \hat{\mathcal{X}}_{k|k,C} \qquad \text{(Clone, no change)}$$
$$e_k = \mathcal{X}_k \cdot \hat{\mathcal{X}}_k^{-1} = \exp(\lfloor \xi_k \rfloor^{\wedge}) \,\text{(Right invariant)} \quad (18)$$
$$\xi_{k,SC} = [\xi_{k,C}; \xi_k] \quad (19)$$

$$e_{k+1|k} = \mathcal{X}_{k+1} \cdot \hat{\mathcal{X}}_{k+1|k}^{-1} \quad (20)$$
$$= \mathcal{X}_k \cdot u_k \cdot n_u \cdot (\hat{\mathcal{X}}_{k|k} \cdot u_k)^{-1}$$
$$= \mathcal{X}_k \cdot \hat{\mathcal{X}}_{k|k}^{-1} \cdot \hat{\mathcal{X}}_{k|k} \cdot u_k \cdot n_u \cdot u_k^{-1} \cdot \hat{\mathcal{X}}_{k|k}^{-1}$$
$$= e_{k|k} \cdot \tilde{n}_u$$
$$\mathbf{F}_k = \mathbf{I}_{3\times3} \qquad \mathbf{F}_{k,SC} = \mathbf{I}_{6\times6} \quad (21)$$
$$\tilde{\boldsymbol{\Sigma}}_u = \text{Ad}(\hat{\mathcal{X}}_{k|k}u_k)\boldsymbol{\Sigma}_u\text{Ad}(\hat{\mathcal{X}}_{k|k}u_k)^T \quad (22)$$
$$\boldsymbol{\Sigma}_{u,SC} = [\mathbf{0}_{3\times3}, \mathbf{0}_{3\times3}; \mathbf{0}_{3\times3}, \tilde{\boldsymbol{\Sigma}}_u] \quad (23)$$

When the relative pose measurement $\mathcal{Y}_{k+1}$ arrives, the update is done similar to [5]:

$$\mathcal{Y}_{k+1} = n_y \cdot \mathcal{X}_{k+1}^{-1} \cdot \mathcal{X}_k \quad (24)$$
$$\mathcal{Y}_{k+1} \cdot \hat{\mathcal{X}}_{k,C}^{-1} = n_y \cdot \hat{\mathcal{X}}_{k+1|k}^{-1} \cdot \hat{\mathcal{X}}_{k+1|k} \cdot \mathcal{X}_{k+1}^{-1} \cdot e_{k,C}$$
$$\mathcal{Y}_{k+1} \cdot \hat{\mathcal{X}}_{k,C}^{-1} = n_y \cdot \hat{\mathcal{X}}_{k+1|k}^{-1} \cdot e_{k+1|k}^{-1} \cdot e_{k,C}$$
$$\hat{\mathcal{X}}_{k+1|k} \cdot \mathcal{Y}_{k+1} \cdot \hat{\mathcal{X}}_{k,C}^{-1} = \hat{\mathcal{X}}_{k+1|k} \cdot n_y \cdot \hat{\mathcal{X}}_{k+1|k}^{-1} \cdot e_{k+1|k}^{-1} \cdot e_{k,C}$$
$$\hat{\mathcal{X}}_{k+1|k} \cdot \mathcal{Y}_{k+1} \cdot \hat{\mathcal{X}}_{k,C}^{-1} = \tilde{n}_y \cdot e_{k+1|k}^{-1} \cdot e_{k,C} \quad (25)$$
$$\mathbf{H}_{k+1} = [\mathbf{I}_{3\times3}, -\mathbf{I}_{3\times3}] \quad (26)$$
$$\tilde{\boldsymbol{\Sigma}}_y = \text{Ad}(\hat{\mathcal{X}}_{k+1|k})\boldsymbol{\Sigma}_y\text{Ad}(\hat{\mathcal{X}}_{k+1|k})^T \quad (27)$$

The reason why $\mathbf{H}_{k+1}$ has a flipped sign compared to [5] is because $e = \mathcal{X}\hat{\mathcal{X}}^{-1}$ is used instead of $e = \hat{\mathcal{X}}\mathcal{X}^{-1}$. Both are related in the tangent space $\mathfrak{se}(2)$ by a flipped sign, since one is the inverse of the other. Furthermore the update with known landmarks is done as in [15].

The regular EKF Jacobians are also shown, without derivation for brevity:

$$\mathbf{F}_k = \begin{pmatrix} 1 & \mathbf{0}_{1\times2} \\ \mathbf{R}(\hat{\theta}_{k|k})\mathbf{u}_{k,x} & \mathbf{I}_{2\times2} \end{pmatrix} \quad (28)$$

$$\mathbf{H}_{k+1} = \begin{pmatrix} 1 & \mathbf{0}_{1\times2} & -1 & \mathbf{0}_{1\times2} \\ \mathbf{0}_{2\times1} & \mathbf{R}^T(\hat{\theta}_{k+1|k}) & \mathbf{h}_{(2:3,4)} & -\mathbf{R}^T(\hat{\theta}_{k+1|k}) \end{pmatrix} \quad (29)$$

$$\mathbf{h}_{(2:3,4)} = -\lfloor 1 \rfloor_{\times}\mathbf{R}^T(\hat{\theta}_{k+1|k})(\hat{\mathbf{x}}_{k,C} - \hat{\mathbf{x}}_{k+1|k})$$

As described earlier, they are estimate-dependent. Noise independent of the state is added following the same noise characteristic as for the IEKF (15) and (16), since noise values are small and up to the first order equal.
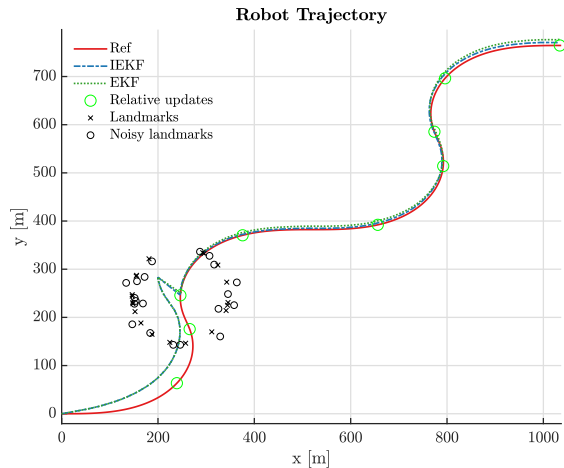
Fig. 2. The robot starts navigating with a wrong heading (blue and green) for half of the trajectory. Global information provided by landmarks (black) arrives delayed and is used to update the current estimate efficiently. The degrading effect of wrong linearization points increases as the heading error increases. The proposed method (blue) is almost not affected by this.
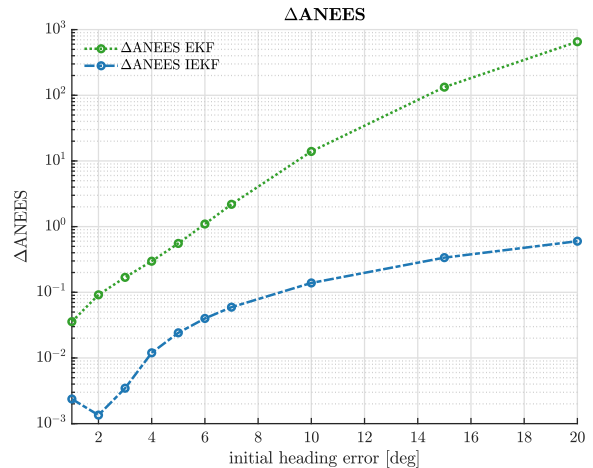

Fig. 3. The increased inconsistency due to covariance pre-integration through the Star-Product approach is shown for the EKF (green) and our IEKF based implementation (blue). A wrong initial heading leads to estimates, and thus linearization points, far from true values. The proposed SP-IEKF method is barely affected while the SP-EKF method reaches critical values (note the log scale on the y-axis!).

IV. SIMULATIONS

The effects of the state-dependency on consistency in the previously presented theory will be analyzed in a simulation setup. Evaluating consistent and efficient fusion of delayed global information by different approaches will be the main focus. But also the applicability of covariance pre-integration for navigation with a wrong heading is demonstrated.

A. The Experimental Setup

A wheeled robot follows the trajectory (red) depicted in Fig. 2 in 2D. It propagates estimates on each time step by odometry and updates at every 100 steps with relative pose measurements. The total length of the trajectory is 900 steps. The initial heading is unknown to the robot, such that the heading error is within the confidence of the initial covariance of the estimator. Approximately halfway through the trajectory delayed global information arrives in the form of landmark measurements at time step 500. The global position of the landmarks in the map is known and they were measured at time step 300, thus the measurement delay is 200 steps. This corrects the initially wrong heading to a large extent. The Star-Product based handling of this delayed measurement is implemented in three estimators: The ideal EKF, the EKF, and the IEKF. From step 300 to 500, all state estimates were far from the real trajectory and were used to pre-compute the scattering matrices of the SP-EKF (see Fig. 1). The different effects on the consistency are shown with the ΔANEES which is of particular interest at step 500 just after the reception of the global signal. At this step, pre-computed Star-Product chains from step 300 to 500 are applied for efficient re-propagation of the covariances due to the delayed measurement. A ΔANEES > 0 at step 500 would imply an inconsistency introduced by the Star-Product chains. For each set of experimental parameters, Monte Carlo simulations with 100 runs were made.

B. Comparison of EKF and Invariant EKF for Increasing Heading Errors

As described earlier, the ideal EKF knows the true state. So the linearization is always done at the true state and is therefore also estimate-independent. This is reflected in the ΔANEES being very close to zero *for any initial heading error*. This means that there is no difference in the use of the recalculation or the covariance pre-integration of the SP-EKF for the ideal EKF.

For small initial heading errors all estimators can work consistently and perform comparably, since deviation of state estimates from the true trajectory is insignificant. Fig. 3 shows the ΔANEES for the EKF and IEKF implementations at step 500 just after the delayed measurement has been applied (with Star-Product and naive re-calculation for each implementation). The effect on the consistency increases with the initial heading error for the EKF, meaning that the application of the Star-Product to the EKF approach significantly deteriorates the estimator consistency over the naive re-calculation. At the same time, the IEKF is less affected by the application of the Star-Product, indicating that the covariance pre-integration of SP-IEKF is as good as the recalculation while being much more efficient (please note the log-scale on the y-axis in Fig. 3). The capability to navigate for a long duration with a wrong heading and to consistently fuse global information after a long delay is demonstrated only for the SP-IEKF.

To better show-case the ΔANEES and how small its value as the difference for the IEKF implementation when naively recalcuated (i.e. best but most costly case) versus SP-IEKF variant is, the ANEES of both methods for the initial wrong heading of 7° is shown in Fig. 4. The figure shows the time evolution of both ANEES' from the time of the global update on once by using recalculations and once by using covariance pre-integration of SP-IEKF. This underlines the claim that SP-IEKF is not much affected by wrong linearization points.

Fig. 4. The plot shows a minimal difference in consistency after the global update for recalculation in a naive IEKF and covariance pre-integration in our suggested efficient SP-IEKF implementation. The initial heading error is 7°.
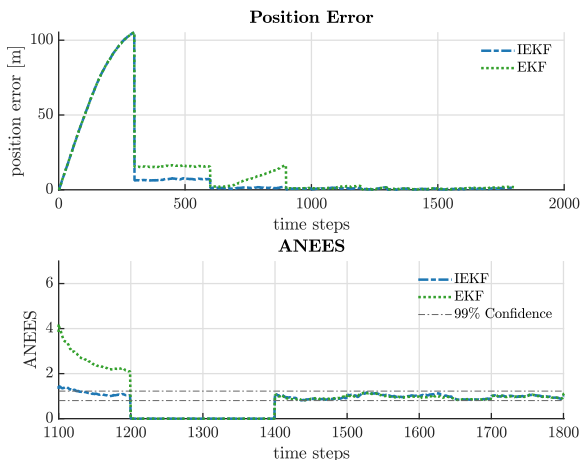

Fig. 5. A simulation is shown with a wrong initial heading of 10°. Global updates are received at step 500, 800, 1100, and 1400, each with a 200 step delay (thus corrections are applied at steps 300, 600, 900, and 1200). The error of the position estimate is plotted at the top. Only after four global updates the estimates become consistent at time step 1400 (see plot below). The SP-EKF remains inconsistent until then and the proposed method based on the IEKF should be preferred. From 1200 to 1400 there is no ANEES because the delayed update step directly propagates the covariance from step 1200 to 1400.

C. Global Information and Consistency

If global information is available at any time, the SP-EKF and the SP-IEKF perform similarly, since global updates will drive the estimator close to the true trajectory and force consistency. But specifically in the case of initial motion estimation and badly initialized estimators relying on the SP-EKF will lead to overconfident estimates as the example in Fig. 5 shows. Starting with an initial heading error of 10°, four global updates were required to reach consistency. While the SP-IEKF remains more consistent in the decisive initial phase of the trajectory.

D. Inconsistencies of the SP-EKF after Reset

The direct way to show the inconsistencies caused by the SP-EKF covariance pre-integration is best shown by the ΔANEES. But there is also an alternative to show the added inconsistencies. As the IEKF outperforms the EKF in terms of consistecy (even with naive recalculations) for large heading errors, we could use the estimates and covariances of the IEKF after the delayed global update also for the EKF. By this procedure the following constraints for the comparison are ensured:

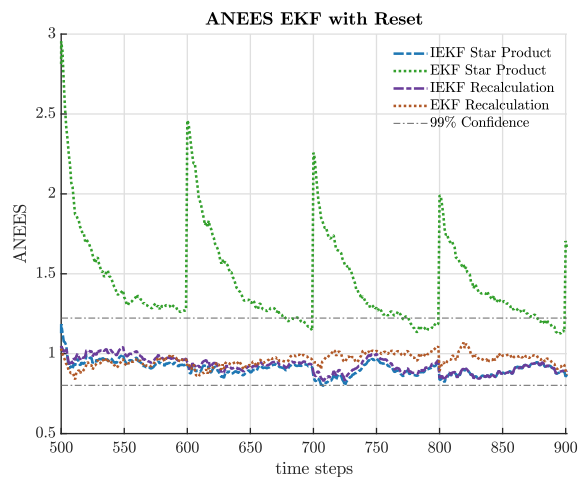- both estimators have the same consistent starting point after global information

Fig. 6. The time evolution of the inconsistencies is shown for the SP-EKF (green). After the delayed global update the estimates of the EKF were reset to consistent values. But the wrong linearization points in the covariance pre-integration for the SP-EKF cause the filter to get overconfident again (green). The spikes are caused by relative updates followed by propagations that increase uncertainty, diminishing the overconfidence. For re-calculation instead, the estimates stay consistent (orange). Indicating that the only source of inconsistency are wrong linearization points in the covariance pre-integration of the SP-EKF. Our approach, the SP-IEKF, is not affected by this due to the IEKF linearization-state independency.

- both estimators have done pre-computations of SP-EKF with the same wrong estimates
- the only difference is in covariance pre-integration with SP-EKF and estimate-dependency in $\mathbf{F}_k$ $\mathbf{H}_{k+1}$ (EKF) or $\tilde{\mathbf{\Sigma}}_{u,SC}$ $\tilde{\mathbf{\Sigma}}_y$ (IEKF)

At time step 300, after the delayed global update, the mean of the EKF is replaced by the mean of the IEKF and the IEKF covariance is transformed and replaces the covariance of the EKF. The transformation is necessary because the error variables of the IEKF (ξ_θ, ξ_x) and EKF ($\delta\theta, \delta x$) are slightly different, but interchangeable (see appendix, Eq. (40)).

Fig. 6 shows the time evolution of the ANEES of both estimators for recalculation and the use of SP-(I)EKF. At time step 500 the delayed global information was processed. As it can be seen for recalculation, both estimators stay consistent, performing similarly. But when covariance pre-integration by SP-EKF is used, inconsistencies become visible (dotted green line). These inconsistencies can only be caused by the estimate-dependency of $\mathbf{F}$ and $\mathbf{H}$ for the EKF and the spikes occur at relative updates, repeatedly making the estimator overconfident.

V. REAL DATA

For real data experiments we have used the UTIAS Multi-Robot Cooperative Localization and Mapping Dataset [25]. In each dataset, robots move to random waypoints in an indoor space while logging odometry data, and range-bearing observations to landmarks.

A. The Experimental Setup

In our experiments we have considered only a single wheeled robot, in particular Robot 1 from Dataset 1, following the red trajectory in Fig. 7. The traveled trajectory is approximately 20.5 meters in 375 seconds with odometry

data sampled at 6.25Hz, thus the total length of the trajectory is 2344 steps. As in the simulation scenario the initial heading is unknown to the robot and the estimates update every 50 time steps with relative pose measurements and every 700 time steps with global range-bearing measurements, delayed approximately by 300 time steps. Again, in this experiment we are comparing the consistency of the star product covariance pre-integration applied on two estimators: the EKF, and the IEKF in terms of ΔNEES.

Different adaptions have been made in order to correctly use the data provided by the dataset, in particular regarding the noise. First, as no noise values are provided in the data set, the motion noise statistics for the noisy odometry data have been set in such a way to obtain an average NEES equal to 1 in the case of propagation only (i.e. making the assumption of having a consistent integrator). As a second step, the range-bearing measurement noise statistics has been set to be Gaussian distributed $\sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_y)$ with $\mathbf{\Sigma}_y = diag\,(0.3m, 0.02rad)$.

To include range-bearing measurements into the IEKF framework, we have expressed the measurement as a function of the right invariant observations and we have derived the estimate-independent measurement Jacobians and the estimate-dependent noise matrix as follows, see also [15]:

$$\mathbf{y}_{k+1} = \begin{bmatrix} r_{k+1} \\ \alpha_{k+1} \end{bmatrix} = \phi\left(\mathcal{X}_{k+1} \diamond {}_G l_{k+1}\right) + \mathbf{n}_{y,k+1} \tag{30}$$

$$\mathbf{y}'_{k+1} = \phi^{-1}\left(\mathbf{y}_{k+1}\right) \simeq \hat{\mathcal{X}}_{k+1|k} \diamond {}_G l_{k+1} + \nabla_{\mathbf{n}_{y,k+1}} \mathbf{n}_{y,k+1} \tag{31}$$

$$\nabla_{\mathbf{n}_{y,k+1}} = \begin{bmatrix} \cos\left(\alpha_{k+1}\right) & \mathbf{h}_{(1:2,2)} \\ \sin\left(\alpha_{k+1}\right) \end{bmatrix} \tag{32}$$

$$\mathbf{h}_{(1:2,2)} = \lfloor 1 \rfloor_\times \mathbf{R}^T(\hat{\theta}_{k+1|k})({}_G l_{k+1} - \hat{\mathbf{x}}_{k+1|k})$$

$$\hat{\mathcal{X}}_{k+1|k} \diamond \mathbf{y}'_{k+1} = e_{k+1|k}^{-1} \diamond {}_G l_{k+1} + \hat{\mathcal{X}}_{k+1|k} \diamond \nabla_{\mathbf{n}_{y,k+1}} \mathbf{n}_{y,k+1} \tag{33}$$

$$\mathbf{H}_{k+1} = \begin{bmatrix} -\lfloor 1 \rfloor_\times {}_G l_{k+1} & -\mathbf{I} \end{bmatrix} \tag{34}$$

$$\tilde{\mathbf{\Sigma}}_{y'} = \mathbf{R}(\hat{\theta}_{k+1|k})\nabla_{\mathbf{n}_{y,k+1}} \mathbf{\Sigma}_y \nabla_{\mathbf{n}_{y,k+1}}^T \mathbf{R}^T(\hat{\theta}_{k+1|k}) \tag{35}$$

B. Comparison of EKF and IEKF for Increasing Initial Heading Errors

The results obtained for the IEKF and EKF through real data experiments and shown in Fig. 7 - 10 confirm the result obtained through simulations. In particular, Fig. 9 shows the NEES for both EKF and IEKF either with or without the covariance pre-integration with an initial heading error of $20°$. The improvement of the recalculation compared to the SP-EKF is striking, not so for the SP-IEKF, as expected – also showing in real data the independence of the SP-IEKF approach to the linearization points while still fully leveraging the complexity reduction of the Star-Product chains. Similarly, Fig. 10 shows the ΔNEES at a fixed time step for different initial heading errors, pointing out that the use of scattering theory for covariance pre-integration allows consistent and efficient fusion of delayed measurement in a IEKF framework, while it proportionally degrades the consistency when used in an EKF framework.
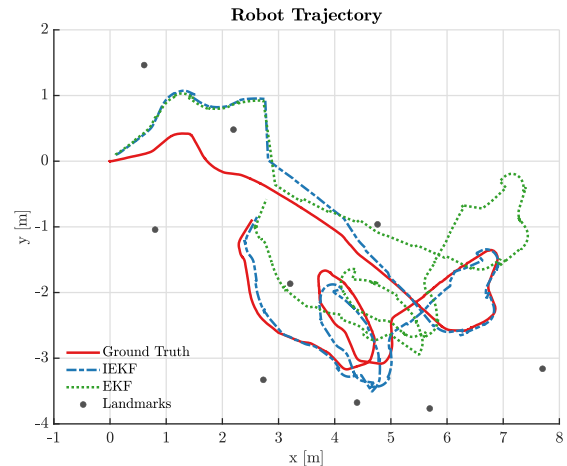

Fig. 7. Wheeled robot trajectory from the UTIAS Multi-Robot Cooperative Localization dataset (robot 1, dataset 1): ground truth (red) is shown while the dashed blue and the dotted green show respectively the SP-IEKF and SP-EKF estimates with an initial heading error of $20°$. The first of the 3 global updates arrives at step 700 with a delay of 300 steps.
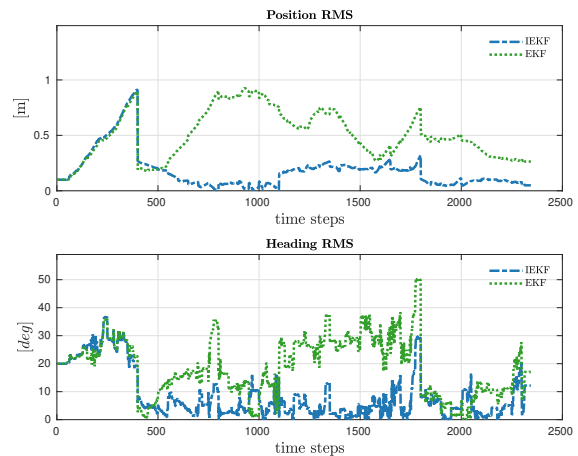

Fig. 8. Position and Heading RMS error for SP-IEKF (in dashed blue) and SP-EKF (in dotted green) for the real data depicted in Fig. 7.
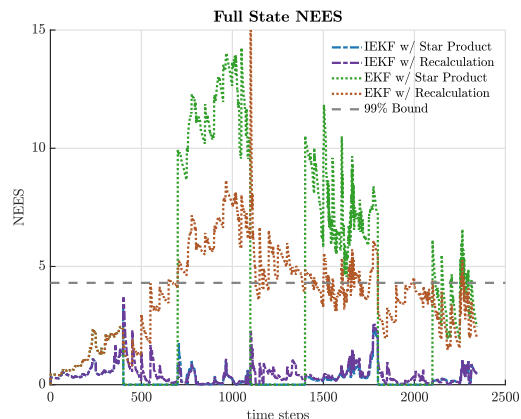

Fig. 9. Full state NEES for a single run of IEKF with recalculation (dashed purple), SP-IEKF (dashed blue), EKF with recalculations (dotted brown) and SP-EKF (dotted green) for initial heading error of $20°$. There is almost no difference in the NEES of recalculation over covariance pre-integration for the the IEKF, while there is significant difference of the NEES for the EKF.
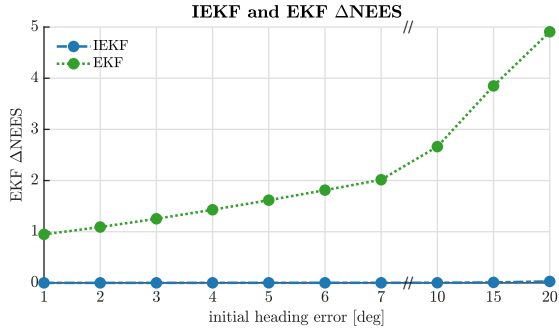
Fig. 10. ΔNEES for the SP-IEKF (in dashed blue) and the SP-EKF (in dotted green) for different initial heading errors computed at a fixed time step after the first global update and the use of covariance pre-integration. The initial heading error does not affect the consistency of the IEKF after the application of the covariance pre-integration, while it does for the EKF. Thereby confirming the same results as for the simulations.

VI. Conclusions

The assumption for the SP-EKF that the effect of the re-linearization is negligible was revisited in this work. As a result, the impact of wrong linearization points was described quantitatively, indicating the importance of estimate dependency. The efficient fusion of delayed measurement was extended by the use of the IEKF methodology to achieve consistent results. The validity of the proposed method was demonstrated in simulation and on a publicly available dataset for collaborative SLAM. By the use of scattering theory and IEKF for the covariance pre-integration, very efficient computations without loss of consistency can be achieved. Thus, allowing the consistent fusion of largely delayed measurements, like database lookups or enabling the processing of data-heavy measurements like point clouds.

Acknowledgement

Appendix

The $SE(2)$ Group elements and action and the $\lfloor \xi \rfloor^\wedge \in \mathfrak{se}(2)$ tangent space elements:

$$\xi = (\xi_\theta, \xi_x^T)^T \quad \lfloor \xi \rfloor^\wedge = \begin{pmatrix} \lfloor \xi_\theta \rfloor_\times & \xi_x \\ 0 & 0 \end{pmatrix} \qquad \xi \in \mathbb{R}^3$$

$$e = \exp(\lfloor \xi \rfloor^\wedge) = \begin{pmatrix} \mathbf{R}(\xi_\theta) & V(\xi_\theta)\xi_x \\ 0 & 1 \end{pmatrix} \qquad e \in SE(2)$$

$$V(\xi_\theta) = \frac{\sin \xi_\theta}{\xi_\theta}\mathbf{I} + \frac{1 - \cos \xi_\theta}{\xi_\theta}\lfloor 1 \rfloor_\times$$

$$\mathcal{X}_k = \begin{pmatrix} \mathbf{R}(\theta_k) & \mathbf{x}_k \\ 0 & 1 \end{pmatrix} \qquad \mathcal{X}_k \in SE(2)$$

$$\mathcal{X}_k \diamond \mathbf{p} = \mathbf{R}(\theta_k)\mathbf{p} + \mathbf{x}_k \qquad \mathbf{p} \in \mathbb{R}^2$$

$$\mathcal{X}_k^{-1} \diamond \mathbf{p} = \mathbf{R}(\theta_k)^T (\mathbf{p} - \mathbf{x}_k)$$

Relation between the EKF and IEKF error variables:

$$e = \mathcal{X} \cdot \hat{\mathcal{X}}^{-1} = \exp(\lfloor \xi \rfloor^\wedge)$$

$$= \begin{pmatrix} \mathbf{R}(\theta)\mathbf{R}^T(\hat{\theta}) & \mathbf{x} - \mathbf{R}(\theta)\mathbf{R}^T(\hat{\theta})\hat{\mathbf{x}} \\ \mathbf{0} & 1 \end{pmatrix} \tag{36}$$

$$\xi_\theta = \delta\theta = \theta - \hat{\theta} \leftarrow \mathbf{R}(\theta)\mathbf{R}^T(\hat{\theta}) \tag{37}$$

$$\xi_x \approx \mathbf{x} - \mathbf{R}(\theta)\mathbf{R}^T(\hat{\theta})\hat{\mathbf{x}} \qquad\qquad \xi_\theta \to 0 \tag{38}$$

$$= \delta x - \lfloor 1 \rfloor_\times \hat{\mathbf{x}}\delta\theta$$

$$\begin{pmatrix} \xi_\theta \\ \xi_x \end{pmatrix} = \begin{pmatrix} 1 & \mathbf{0}_{1\times 2} \\ -\lfloor 1 \rfloor_\times \hat{\mathbf{x}} & \mathbf{I}_{2\times 2} \end{pmatrix} \begin{pmatrix} \delta\theta \\ \delta x \end{pmatrix} = \mathbf{M} \begin{pmatrix} \delta\theta \\ \delta x \end{pmatrix} \tag{39}$$

$$\mathbf{P}_{\text{IEKF}} = \mathbf{M}\mathbf{P}_{\text{EKF}}\mathbf{M}^T \qquad \mathbf{P}_{\text{EKF}} = \mathbf{M}^{-1}\mathbf{P}_{\text{IEKF}}\mathbf{M}^{-T} \tag{40}$$

References

[1] E. Allak, R. Jung, and S. Weiss, "Covariance pre-integration for delayed measurements in multi-sensor fusion," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Nov 2019, pp. 6642–6649.

[2] R. Van Der Merwe, "Sigma-point kalman filters for probabilistic inference in dynamic state-space models," Ph.D. dissertation, Oregon Health & Science University, 2004, aAI3129163.

[3] T. D. Larsen, N. A. Andersen, O. Ravn, and N. K. Poulsen, "Incorporation of time delayed measurements in a discrete-time kalman filter," in *Decision and Control, 1998. Proceedings of the 37th IEEE Conference on*, vol. 4. IEEE, 1998, pp. 3972–3977.

[4] S. I. Roumeliotis and J. W. Burdick, "Stochastic cloning: A generalized framework for processing relative state measurements," in *Robotics and Automation, 2002. Proceedings. ICRA'02. IEEE International Conference on*, vol. 2. IEEE, 2002, pp. 1788–1795.

[5] A. Barrau and S. Bonnabel, "Invariant filtering for pose ekf-slam aided by an imu," in *2015 54th IEEE Conference on Decision and Control (CDC)*. IEEE, 2015, pp. 2133–2138.

[6] T. Lupton and S. Sukkarieh, "Efficient integration of inertial observations into visual slam without initialization," in *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*. IEEE, 2009, pp. 1547–1552.

[7] 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.

[8] R. Redheffer, "On the relation of transmission-line theory to scattering and transfer," *Journal of Mathematics and Physics*, vol. 41, no. 1-4, pp. 1–41, 1962.

[9] T. Kailath, A. H. Sayed, and B. Hassibi, *Linear estimation*. Prentice Hall, 2000, no. EPFL-BOOK-233814.

[10] B. Friedlander, T. Kailath, and L. Ljung, "Scattering theory and linear least squares estimation—ii. discrete-time problems," *Journal of the Franklin Institute*, vol. 301, no. 1-2, pp. 71–82, 1976.

[11] G. Verghese, B. Friedlander, and T. Kailath, "Scattering theory and linear least-squares estimation, part iii: The estimates," *IEEE Transactions on Automatic Control*, vol. 25, no. 4, pp. 794–802, 1980.

[12] T. Bailey, J. Nieto, J. Guivant, M. Stevens, and E. Nebot, "Consistency of the ekf-slam algorithm," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2006, pp. 3562–3568.

[13] S. J. Julier and J. K. Uhlmann, "A counter example to the theory of simultaneous localization and map building," in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, vol. 4, May 2001, pp. 4238–4243 vol.4.

[14] S. Bonnabel, P. Martin, and P. Rouchon, "Symmetry-preserving observers," *IEEE Transactions on Automatic Control*, vol. 53, no. 11, pp. 2514–2526, 2008.

[15] A. Barrau and S. Bonnabel, "The invariant extended kalman filter as a stable observer," *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1797–1812, April 2017.

[16] A. Barrau and S. Bonnabel, "Linear observed systems on groups," *Systems & Control Letters*, vol. 129, pp. 36–42, 2019.

[17] ——, "An ekf-slam algorithm with consistency properties," *arXiv preprint arXiv:1510.06263*, 2015.

[18] M. Brossard, A. Barrau, and S. Bonnabel, "Exploiting symmetries to design ekfs with consistency properties for navigation and slam," *IEEE Sensors Journal*, vol. 19, no. 4, pp. 1572–1579, 2018.

[19] A. Barrau and S. Bonnabel, "Extended kalman filtering with nonlinear equality constraints: a geometric approach," *IEEE Transactions on Automatic Control*, 2019.

[20] M. Brossard, S. Bonnabel, and A. Barrau, "Unscented kalman filter on lie groups for visual inertial odometry," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2018, pp. 649–655.

[21] R. Hartley, M. Ghaffari, R. M. Eustice, and J. W. Grizzle, "Contact-aided invariant extended kalman filtering for robot state estimation," *The International Journal of Robotics Research*, p. 0278364919894385, 2019.

[22] P. Chauchat, A. Barrau, and S. Bonnabel, "Invariant smoothing on lie groups," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1703–1710.

[23] M.-A. Lavoie, J. Arsenault, and J. R. Forbes, "An invariant extended h-infinity filter," in *IEEE Conference on Decision and Control (CDC)*. IEEE, 2019.

[24] X. R. Li, Z. Zhao, and V. P. Jilkov, "Practical measures and test for credibility of an estimator," in *Proc. Workshop on Estimation, Tracking, and Fusion—A Tribute to Yaakov Bar-Shalom*, 2001, pp. 481–495.

[25] K. Y. Leung, Y. Halpern, T. D. Barfoot, and H. H. Liu, "The utias multi-robot cooperative localization and mapping dataset," *The International Journal of Robotics Research*, vol. 30, no. 8, pp. 969–974, 2011.