SC-KF Mobile Robot Localization: A Stochastic-Cloning Kalman Filter for Processing Relative-State Measurements

Anastasios I. Mourikis†, Stergios I. Roumeliotis†, and Joel W. Burdick‡

***Abstract*— This paper presents a new method to optimally combine motion measurements provided by proprioceptive sensors, with relative-state estimates inferred from feature-based matching. Two key challenges arise in such pose tracking problems: (i) the displacement estimates relate the state of the robot at two different time instants, and (ii) the same exteroceptive measurements are often used for computing consecutive displacement estimates, a process which renders the errors in these correlated. We present a novel Stochastic Cloning-Kalman Filtering (SC-KF) estimation algorithm that successfully addresses these challenges, while still allowing for efficient calculation of the filter gains and covariances. The proposed algorithm is not intended to compete with Simultaneous Localization and Mapping (SLAM) approaches. Instead it can be merged with any EKF-based SLAM algorithm to increase its precision. In this respect, the SC-KF provides a robust framework for leveraging additional motion information extracted from dense point features that most SLAM algorithms do not treat as landmarks. Extensive experimental and simulation results are presented to verify the validity of the proposed method and to demonstrate that its performance is superior to that of alternative position tracking approaches.**

***Index Terms*— Stochastic Cloning, robot localization, relative-pose measurements, displacement estimates, state augmentation.**

I. Introduction

Accurate localization is a prerequisite for a robot to meaningfully interact with its environment. The most commonly available sensors for acquiring localization information are *proprioceptive* sensors, such as wheel encoders, gyroscopes, and accelerometers that provide information about the robot's motion. In Dead Reckoning (DR) [1], a robot's pose can be tracked from a starting point by integrating proprioceptive measurements over time. The limitation of DR is, however, that since no external reference signals are employed for correction, estimation errors accumulate over time, and the pose estimates drift from their real values. In order to improve localization accuracy, most algorithms fuse the proprioceptive measurements with data from *exteroceptive* sensors, such as cameras [2], [3], laser range finders [4], sonars [5], etc.

When an exteroceptive sensor provides information about the position of features with respect to the robot at two different time instants, it is possible (under necessary observability conditions) to create an *inferred* measurement of the robot's displacement. Examples of algorithms that process exteroceptive data to infer motion include laser scan matching [4], [6], [7], vision-based motion estimation techniques using stereoscopic [2], [3], and monocular [8] image sequences, and matching of sonar returns [5]. The inferred

†Dept. of Computer Science & Engineering, University of Minnesota, Minneapolis, MN 55455. Email: {mourikis|stergios}@cs.umn.edu. ‡Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125. Email: jwb@robotics.caltech.edu.

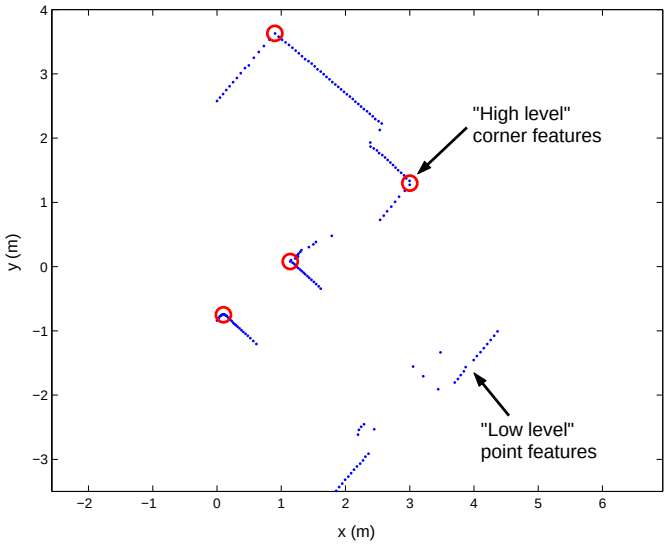

Fig. 1. Example of a planar laser scan and types of features observed. An algorithm has been employed to detect corners (intersections of line segments) in a laser scan. The extracted corner features can be used for performing SLAM, while all the remaining, "low-level", feature points, can be utilized in the SC-KF framework to improve the pose tracking accuracy.

relative-state measurements that are derived from these can be integrated over time to provide pose estimates [3], or combined with proprioceptive sensory input in order to benefit from both available sources of positioning information [9], [10]. This paper focuses on how to optimally implement the latter approach using an extended Kalman filter (EKF) [11]. This paper does not consider the case in which the feature measurements are used for SLAM. However, as discussed in Section VI, our approach is complementary to SLAM, and can be employed to increase its accuracy (cf. Fig. 1).

Two challenges arise when fusing proprioceptive and relative-pose[1] measurements in an EKF. Firstly, each displacement measurement relates the robot's state at two different time instants (i.e., the current time and previous time when exteroceptive measurements were recorded). However, the basic theory underlying the EKF requires that the measurements used for the state update be independent of any previous filter states. Thus, the "standard" formulation of the EKF, in which the filter's state comprises only the current state of the robot, is clearly not adequate for treating relative-state measurements.

A second challenge arises from the fact that when exte-

[1]Throughout this paper, the terms "displacement measurement" and "relative-pose measurement" are used interchangeably to describe a measurement of the robot's motion that is inferred from exteroceptive measurements. Depending on the type and number of available features, either all, or a subset, of the degrees of freedom of motion may be determined (cf. Section VII-A.2).

roceptive measurements are used to infer displacement, consecutive relative-state measurements will often be *correlated*. To understand the source of such correlations, consider, for example, the scenario in which a camera is employed to measure the pixel coordinates of the projections of the same landmarks at times t_{k-1}, t_k and t_{k+1}. The errors in the measurements at time t_k affect the displacement estimates for *both* time intervals $[t_{k-1}, t_k]$ and $[t_k, t_{k+1}]$, thereby rendering them correlated. Assuming that the measurements are uncorrelated (as is customarily done [2], [7], [10]), violates a basic assumption of EKF theory, leading to sub-optimal or incorrect estimates for the robot's state and covariance. This fact has been generally overlooked in the literature, and to the best of our knowledge, no prior work exists that directly addresses this issue.

In this paper we propose a direct approach to the problem of combining relative-pose measurements with proprioceptive measurements in order to improve the accuracy of DR. Our methodology augments the state vector of the Kalman filter to address the two aforementioned challenges. In particular, to properly account for the dependencies on the robot's state estimates at different time instants, we augment the Kalman filter state to include two instances (or "*clones*") of the state estimate—hence the name *Stochastic Cloning Kalman Filter (SC-KF)* [9]. Moreover, in order to appropriately treat the correlations between consecutive displacement estimates, we further augment the state to include the most recent exteroceptive measurements [11]. With these state augmentations the displacement measurements can be expressed as functions of the current filter state, and thus an EKF framework can be employed.

The following section reviews existing approaches for processing relative-state measurements, while Section III presents the structure of the correlations between consecutive measurements, and investigates their effect on displacement-only propagation of the robot state. Section IV describes in detail the SC-KF algorithm. Section V presents extensions of the SC-KF methodology, while Section VI discusses its relation to SLAM. In Section VII, it is shown that the attained position tracking accuracy is superior to that of existing approaches. Finally, the conclusions of this work are presented in Section VIII.

II. Related Approaches

Displacement measurements can be treated as average velocity measurements during the corresponding time interval. These average velocities can then be combined with velocity measurements obtained from the robot's proprioceptive sensors to improve their accuracy. However, this approach is only applicable if the relative-state measurements are made at a rate equal or higher to that of the proprioceptive sensors, which is rarely the case in practice. Alternatively, the robot's velocity estimate could be included in the state vector, and the average velocity estimates could then be used as instantaneous velocity pseudo-measurements in the EKF update step [12]. The shortcoming of this method is that treating an *average* velocity measurement as an *instantaneous* one can introduce significant errors when the rate of the displacement measurements is low. A different solution, proposed in [10], is to use the previous robot position estimates for converting the relative pose measurements to absolute position pseudo-measurements. However, since these pseudo-absolute measurements are correlated with the state, their covariance matrix has to be artificially inflated to guarantee consistency, thus resulting in suboptimal estimation (cf. Section VII-A).

Contrary to the preceding *ad-hoc* methods for processing displacement measurements, several existing approaches employ these measurements to impose constraints between consecutive robot poses. Algorithms that only use displacement measurements for propagating the robot's state estimate are often described as sensor-based odometry methods [2], [4]. In these algorithms, only the last two robot poses (the current and previous one) are ever considered. While our stochastic cloning approach (which was first introduced in [9]) also relies only upon the last two robot poses, tracking is achieved by *fusing* the displacement measurements with proprioceptive information. Therefore, our method can be seen as an "enhanced" form of odometry. On the other hand, several existing approaches maintain a state vector comprised of a history of robot poses, and use the displacement measurements to impose constraints between pairs of these poses. In [13], the robot's orientation errors are assumed to be temporally uncorrelated, which transforms the problem of optimizing the network of robot poses into a linear one, where only the robot positions are estimated. In [14]–[16] the full 3D robot pose of an autonomous underwater vehicle is estimated, while in [7], [17] displacement constraints are employed for estimating the pose history of a robot in 2D.

All of the approaches discussed so far do *not* properly account for the correlations that exist between consecutive displacement estimates, as they are assumed to be independent. However, as shown in Section III, this assumption does not generally hold. One could avoid such correlations by using each feature measurement in the computation of only one displacement estimate [14]. For example, half the measurements at each time step can be used to estimate the previous displacement, and the other half to estimate the next one. The drawback of this methodology is that incorporating only part of the available exteroceptive measurements when computing each relative-pose estimate results in less accurate displacement estimates. In our work, all available measurements are used to compute the relative-pose measurements at every time step, and the correlations introduced by this process are explicitly identified and accounted for.

Solutions to the well-known Simultaneous Localization and Mapping (SLAM) problem (cf. Section VI) "circumvent" the problem of treating the displacement measurements by including the features' positions in the state vector, and jointly estimating the robot's and features' state. While SLAM offers high localization accuracy, the computational complexity associated with the estimation of the positions of a large number of features may be prohibitive for some real-time applications (e.g., autonomous aircraft landing). Thus there exists a need for methods that enable *direct* processing of the displacement measurements, at a lower computational cost.

In this paper, we propose an algorithm for optimally fusing the potentially correlated relative displacement estimates with proprioceptive measurements. The SC-KF considers exteroceptive measurements in pairs of consecutive measurements that are first processed to create an inferred relative-pose measurement, and then fused with the proprioceptive measurements. The sole objective of the SC-KF is to estimate the robot's state, and therefore the states of features used for deriving the displacement measurements are not estimated. Hence the proposed algorithm can optimally fuse relative-pose measurements with the minimum computational overhead (Section IV-D). The proposed method can be used either as a stand-alone localization algorithm, or combined with SLAM in order to increase its localization accuracy (cf. Section VI).

III. Relative-pose Measurement Correlations

Before presenting the SC-KF algorithm, we first study the structure of the correlations between consecutive displacement estimates. Let z_k and z_{k+m} denote the vectors of exteroceptive measurements at times t_k and t_{k+m}, respectively, whose noise covariance matrices are R_k and R_{k+m}. These are measurements, for example, of range and bearing from a laser range finder, or of bearing from a camera. By processing these measurements (e.g., via laser scan matching), an estimate, $z_{k,k+m}$, for the change in the robot pose between times t_k and t_{k+m} is computed as a function (either closed-form or implicit) of z_k and z_{k+m}:

$$z_{k,k+m} \quad = \quad \xi(z_k, z_{k+m}) \tag{1}$$

Linearization of (1) relates the error in the displacement estimate, $\widetilde{z}_{k,k+m}$, to errors in the exteroceptive measurements:[2]

$$\widetilde{z}_{k,k+m} \quad \simeq \quad J^k_{k,k+m}\widetilde{z}_k + J^{k+m}_{k,k+m}\widetilde{z}_{k+m} + n_{k,k+m} \tag{2}$$

where the noise term $n_{k,k+m}$ arises from inaccuracies in the displacement estimation algorithm (e.g., errors due to feature matching [6]). We assume that the exteroceptive measurement errors, $\widetilde{z}_k$ and $\widetilde{z}_{k+m}$, and the noise term, $n_{k,k+m}$, are zero-mean and independent, an assumption which holds in most practical cases if proper sensor characterization is performed. In (2), $J^k_{k,k+m}$ and $J^{k+m}_{k,k+m}$ are the Jacobians of the function ξ with respect to z_k and z_{k+m}, respectively, i.e.,

$$J^k_{k,k+m} = \nabla_{z_k}\xi \quad \text{and} \quad J^{k+m}_{k,k+m} = \nabla_{z_{k+m}}\xi$$

Generally, not all feature measurements in the vector z_k are used to estimate displacement. For example, in laser scan matching there usually exists only partial overlap between consecutive scans and therefore not all laser returns are matched. As a result, if M_k denotes the number of feature measurements in $z_k^T = \begin{bmatrix} (z_k)_1^T \dots (z_k)_{M_k}^T \end{bmatrix}$, the i-th component of the Jacobian matrices $J^k_{k,k+m}$ and $J^{k+m}_{k,k+m}$ takes the form

$$\left(J^t_{k,k+m}\right)_i = \begin{cases} \nabla_{(z_t)_i}\xi, & i\text{th feature used} \\ & \text{to compute } z_{k,k+m} \\ 0, & \text{else} \end{cases} \tag{3}$$

for $i = 1 \dots M_k$ and $t = k, k+m$. Thus for some applications, the Jacobians may be significantly sparse.

Our goal is to compute the correlation between the displacement estimates for the time intervals $[t_{k-\ell}, t_k]$ and $[t_k, t_{k+m}]$, which is defined as $E\{\widetilde{z}_{k-\ell,k}\widetilde{z}^T_{k,k+m}\}$. For this purpose we employ (2), and the independence of exteroceptive measurement errors at different time-steps, to obtain:

$$\begin{aligned} E\{\widetilde{z}_{k-\ell,k}\widetilde{z}^T_{k,k+m}\} &= J^k_{k-\ell,k}E\{\widetilde{z}_k\widetilde{z}^T_k\}J^{k\ T}_{k,k+m} \\ &= J^k_{k-\ell,k}R_k J^{k\ T}_{k,k+m} \end{aligned} \tag{4}$$

Note that exteroceptive measurements typically consist of observations of a number of features detected in the robot's vicinity (e.g., distance and bearing to points on a wall, or the image coordinates of visual features). In such cases, the measurements of the individual features are mutually independent, and therefore the covariance matrix R_k is block diagonal. In light of (3), when R_k is block diagonal, expression (4) is equal to zero only if *different* features are used to estimate displacement in consecutive time intervals (i.e., if non-overlapping subsets of z_k are matched with $z_{k-\ell}$ and z_{k+m}, respectively). Clearly, this is not the case in general, and thus consecutive displacement estimates are in most cases not independent.

A. State Propagation Based Exclusively on Displacement Measurements

We now show how the preceding analysis can be employed in the simple setting where the robot state estimates are propagated using displacement measurements only. This is an important special case, which has been extensively studied in the literature (examples include visual odometry [2], [3], laser-based odometry [4], etc). Once the displacement estimate between t_k and t_{k+1} has been computed (cf. (1)), an estimate for the robot's pose at t_{k+1} is obtained by combining the previous pose estimate and the displacement measurement:

$$\widehat{X}_{k+1} \quad = \quad g(\widehat{X}_k, z_{k,k+1}) \ . \tag{5}$$

By linearizing this equation, the pose errors at t_{k+1} can be related to the errors in the previous state estimate and displacement measurement:

$$\widetilde{X}_{k+1} \quad \simeq \quad \Phi_k\widetilde{X}_k + \Gamma_k\widetilde{z}_{k,k+1} \tag{6}$$

where Φ_k and Γ_k represent the Jacobians of the state propagation function, $g(\widehat{X}_k, z_{k,k+1})$, with respect to the previous pose and the relative pose measurement, respectively:

$$\Phi_k = \nabla_{\widehat{X}_k}g, \qquad \Gamma_k = \nabla_{z_{k,k+1}}g \ . \tag{7}$$

The covariance matrix of the pose estimates is propagated by:

$$\begin{aligned} P_{k+1} &= E\{\widetilde{X}_{k+1}\widetilde{X}^T_{k+1}\} \\ &= \Phi_k P_k \Phi_k^T + \Gamma_k R_{k,k+1}\Gamma_k^T \end{aligned}$$

[2]The "hat" symbol, $\widehat{\ }$, is used to denote the estimated value of a quantity, while the "tilde" symbol, $\widetilde{\ }$, is used to signify the error between the actual value of a quantity and its estimate. The relationship between a variable, x, its estimate, $\widehat{x}$, and the error $\widetilde{x}$, is $\widetilde{x} = x - \widehat{x}$.

$$+ \Phi_k E\{\widetilde{X}_k \widetilde{z}_{k,k+1}^T\}\Gamma_k^T + \Gamma_k E\{\widetilde{z}_{k,k+1}\widetilde{X}_k^T\}\Phi_k^T \quad (8)$$

where $R_{k,k+1}$ denotes the noise covariance of the displacement estimates. A common simplifying assumption in the literature (e.g., [2], [7]) is that the measurement noise, $\widetilde{z}_{k,k+1}$, and state error, $\widetilde{X}_k$, are uncorrelated, and thus the last two terms in (8) are set to zero. However, this assumption does not generally hold when correlations exist between consecutive displacement estimates. In particular, by linearizing the state propagation equation at t_k, we obtain (cf. (6)):

$$\begin{aligned} E\{\widetilde{z}_{k,k+1}\widetilde{X}_k^T\} &= E\left\{\widetilde{z}_{k,k+1}\left(\Phi_{k-1}\widetilde{X}_{k-1} + \Gamma_{k-1}\widetilde{z}_{k-1,k}\right)^T\right\} \\ &= E\{\widetilde{z}_{k,k+1}\widetilde{X}_{k-1}^T\}\Phi_{k-1}^T + E\{\widetilde{z}_{k,k+1}\widetilde{z}_{k-1,k}^T\}\Gamma_{k-1}^T \\ &= E\{\widetilde{z}_{k,k+1}\widetilde{z}_{k-1,k}^T\}\Gamma_{k-1}^T \ . \end{aligned} \quad (9)$$

Note that the error term $\widetilde{X}_{k-1}$ depends on the measurement errors of all exteroceptive measurements up to and including time t_{k-1}, while the error term $\widetilde{z}_{k,k+1}$ depends on the measurement errors at times t_k and t_{k+1} (cf. (2)). As a result, the errors $\widetilde{X}_{k-1}$ and $\widetilde{z}_{k,k+1}$ are independent. Therefore, by applying the zero-mean assumption for the error $\widetilde{z}_{k,k+1}$ we obtain $E\{\widetilde{z}_{k,k+1}\widetilde{X}_{k-1}^T\} = 0$. Employing the result of (4) and substituting (9) in (8), we obtain the following expression for the propagation of the pose covariance in the case of inferred displacement measurements:

$$\begin{aligned} P_{k+1} &= \Phi_k P_k \Phi_k^T + \Gamma_k R_{k,k+1}\Gamma_k^T \\ &+ \Phi_k \Gamma_{k-1} J_{k-1,k}^k R_k J_{k,k+1}^{k\ T}\Gamma_k^T \\ &+ \Gamma_k J_{k,k+1}^k R_k J_{k-1,k}^{k\ T}\Gamma_{k-1}^T\Phi_k^T \end{aligned} \quad (10)$$

In Algorithm 1, the steps necessary for propagating the robot's state estimate and its covariance using displacement measurements are outlined.

Algorithm 1 Pose Estimation Based on Relative-Pose Measurements

Initialization:

- Initialize the robot covariance matrix when the first exteroceptive measurement is received

Propagation: For each exteroceptive measurement:

- compute the displacement measurement using (1) and its Jacobians with respect to the current and previous exteroceptive measurement using (3).
- propagate the robot state estimate using (5)
- compute the Jacobians of the pose propagation function using (7)
- propagate the robot pose covariance matrix via (10) (during the first iteration, use only the first two terms)
- compute and store the matrix product $\Gamma_k J_{k,k+1}^{k+1}$ that will be used in the next iteration

B. Investigation of the effects of correlations

Based on numerous experiments and simulation tests, we have observed that when the correlations between displacement measurements are accounted for, the covariance estimate is typically smaller than when the correlations are ignored. We attribute this result to the fact that the correlation between consecutive relative-pose estimates tends to be *negative*. An intuitive explanation for this observation can be given by means of a simple example, for 1-dimensional motion. Consider a robot moving on a straight line, and recording measurements, z_k, of the distance to a single feature on the same line. If at time t_k the error in the distance measurement is equal to $\epsilon_k > 0$, this error will contribute towards *underestimating* the robot's displacement during the interval $[t_{k-1}, t_k]$, but will contribute towards *overestimating* the displacement during the interval $[t_k, t_{k+1}]$. Therefore, the error ϵ_k has opposite effects on the two displacement estimates, rendering them negatively correlated.

In this 1D example, it is interesting to examine the time evolution of the covariance when the correlations are properly treated. Note that the robot's displacement can be computed as the difference of two consecutive distance measurements, i.e., $z_{k,k+1} = z_k - z_{k+1}$. If the covariance of the individual distance measurements is equal to $R_k = R_{k+1} = \sigma^2$, then the covariance of $z_{k,k+1}$ is equal to $R_{k,k+1} = 2\sigma^2$. Moreover, for this example it is easy to see that all the Jacobians in (10) are constant, and given by $J_{k,k+1}^k = 1$, $J_{k-1,k}^k = -1$, $\Phi_k = \Gamma_k = \Gamma_{k-1} = 1$. Substituting these values in (10), we obtain the following equation for covariance propagation in this case:

$$P_{k+1} \ = \ P_k + R_{k,k+1} - R_k - R_k \ = \ P_k \ . \quad (11)$$

We thus see that the covariance of the robot's position estimate remains *constant* during propagation when the correlations are properly treated. This occurs because the error in the measurement z_k effectively "cancels out". On the other hand, if the correlations between consecutive displacement measurements are ignored, we obtain

$$P_{k+1} \ = \ P_k + R_{k,k+1} \ = \ P_k + 2\sigma^2 \ . \quad (12)$$

In this case the position covariance increases linearly, a result that does not reflect the evolution of the true state uncertainty.

In the context of this 1D example, we next study the time evolution of the covariance when features come in and out of the robot's field of view (FOV). Assume that a uniform distribution of features, with density ρ, exists on the line, and that the robot's FOV is limited to $\ell_{\max}/2$ in each direction. If the robot moves by $\Delta\ell$ between the time instants the measurements are recorded, then the overlap in the FOV at consecutive time instants is $\ell_{\max} - \Delta\ell$. Within this region lie $M_k = \rho(\ell_{\max} - \Delta\ell)$ features, whose measurements are used for displacement estimation. The least-squares displacement estimate is given by:

$$z_{k,k+1} = \frac{1}{M_k}\sum_{i=1}^{M_k}((z_k)_i - (z_{k+1})_i) \quad (13)$$

where $(z_k)_i$ and $(z_{k+1})_i$ are the measurements to the i-th feature at times t_k and t_{k+1}, respectively. The covariance of $z_{k,k+1}$ is given by:

$$R_{k,k+1} = \frac{2\sigma^2}{M_k} = \frac{2\sigma^2}{\rho(\ell_{\max} - \Delta\ell)} \ . \quad (14)$$

Thus, if one ignores the correlations between consecutive

displacement estimates, the covariance propagation equation is:

$$P_{k+1}^{\rm NC} = P_k^{\rm NC} + \frac{2\sigma^2}{\rho(\ell_{\max} - \Delta\ell)} \ . \qquad (15)$$

where the superscript NC denotes the fact that no correlations are treated. At the end of a path of length $\ell_{\rm total}$ (i.e., after $\ell_{\rm total}/\Delta\ell$ propagation steps), the estimated covariance of the robot position, starting from a zero initial value, will be given by:

$$P_{\rm final}^{\rm NC} = \frac{\ell_{\rm total}}{\Delta\ell} \frac{2\sigma^2}{\rho(\ell_{\max} - \Delta\ell)} \ . \qquad (16)$$

We now derive the corresponding covariance equations for the case that the correlations are properly incorporated. Since the robot moves by a distance $\Delta\ell$ between the time instants when the measurements are recorded, the number of features that are observed at three consecutive time instants (i.e., t_{k-1}, t_k, and t_{k+1}), is $\rho(\ell_{\max} - 2\Delta\ell)$. Employing this observation to evaluate the Jacobians in (10) yields the following expression for the propagation of the covariance:

$$P_{k+1} = P_k + \frac{2\sigma^2\Delta\ell}{\rho(\ell_{\max} - \Delta\ell)^2}, \quad \text{for} \quad \ell_{\max} > 2\Delta\ell \ . \qquad (17)$$

Note that if $\ell_{\max} < 2\Delta\ell$, no overlap exists between the FOV at times t_{k-1} and t_{k+1}, and thus no feature measurement is used twice for computing displacement estimates. In that case, expression (15) is exact. At the end of a path of length $\ell_{\rm total}$, the covariance of the robot position is:

$$P_{\rm final} = \frac{2\sigma^2\ell_{\rm total}}{\rho(\ell_{\max} - \Delta\ell)^2}, \quad \text{for} \quad \ell_{\max} > 2\Delta\ell \ . \qquad (18)$$

From (16) and (18) we see that for $\ell_{\max} > 2\Delta\ell$, the following relation holds:

$$\frac{P_{\rm final}^{\rm NC}}{P_{\rm final}} = \frac{\ell_{\max} - \Delta\ell}{\Delta\ell} > 1 \qquad (19)$$

This shows that when the correlations are ignored, the resulting covariance estimates are larger, similarly to what is observed in the experimental results.

Fig. 2 plots the variance in the robot's position at the end of a trajectory of length $\ell_{\rm total} = 100$ m, as a function of the size of the robot's displacement between consecutive measurements. The solid line corresponds to the case when the correlations between displacement measurements are accounted for (cf. (18)), while the dashed line corresponds to the case when these are ignored (cf. (16)). The parameters used to generate this plot are: the feature density is $\rho = 5$ features/m, the robot's FOV is $\ell_{\max} = 10$ m, and the standard deviation of each distance measurement is $\sigma = 0.2$ m. It is important to note that when the correlations between consecutive measurements are accounted for, the final uncertainty is a *monotonically increasing* function of the displacement between measurements, $\Delta\ell$. This agrees with intuition, which dictates that when measurements occur less frequently, the accuracy of the final state estimates deteriorates. However, when the correlations between displacement measurements are ignored, the covariance estimates do not have this property.

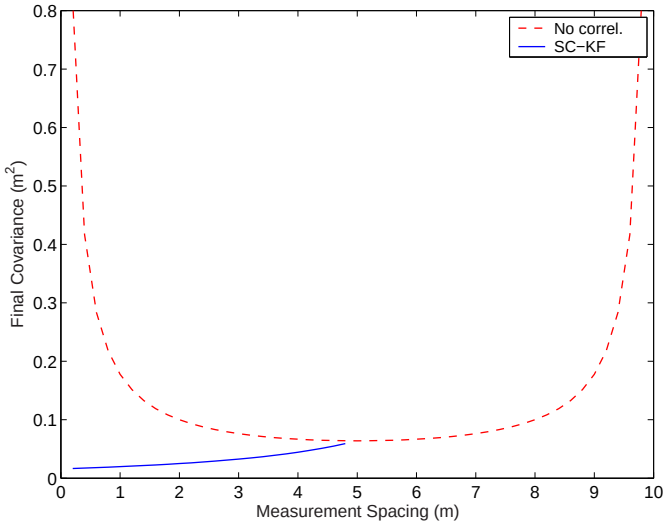

Fig. 2. The covariance estimates at the end of a 100 m trajectory using the expression of (10) (solid line), vs. when the correlations between consecutive displacement measurements are not accounted for (dashed line). Note that when the measurements occur more than 5 m apart, no correlations exist, and the two estimates are identical.

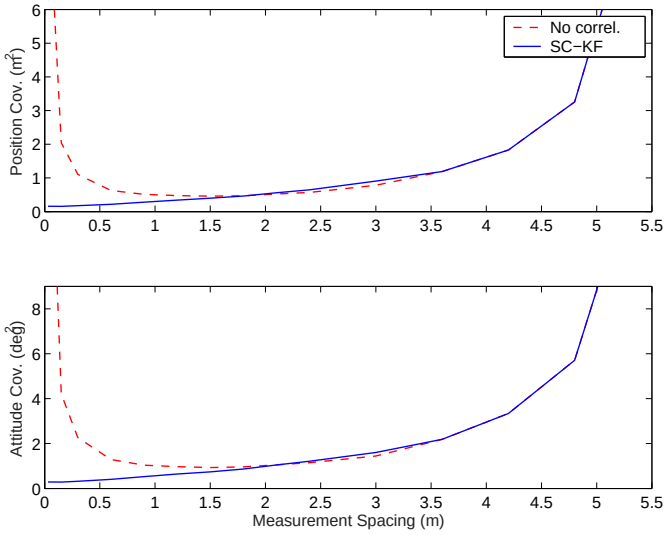

Fig. 3. The covariance estimates at the end of a 100 m trajectory, for a robot performing visual odometry with a stereo pair. The top plot shows the position uncertainty, while the bottom one the attitude uncertainty. When correlations are properly treated (solid lines), the covariance is a monotonic function of the measurement spacing. This is not the case when correlations are ignored (dashed lines).

Fig. 2 shows that for $\Delta\ell < \frac{\ell_{\max}}{2} = 5$ m, as measurements are recorded more frequently, the covariance estimates become *larger*. This behavior is clearly incorrect, and arises due to the fact that the dependency between consecutive displacement estimates is ignored.

The preceding analysis substantiates, at least in the simple case of a robot moving in 1D, that the use of expression (10) for covariance propagation results in considerably more accurate covariance estimates. Unfortunately, for robots moving in 2D [4] and 3D [2], the covariance propagation equations are time-varying (the Jacobians appearing in (10) depend on the robot state and the positions of the features relative to the robot). As a result, an analogous closed-form analysis for

general trajectories and arbitrary feature placement appears to be intractable. However, simulation experiments indicate that the conclusions drawn from the analytical expressions for the 1D case also apply to the more practical scenarios of robots moving in 2D and 3D. For example, Fig. 3 shows the position and attitude covariance at the end of a 100 m trajectory for a robot performing visual odometry with a stereo pair of cameras [2]. The plotted lines represent the traces of the submatrices of the covariance matrix corresponding respectively to position (top subplot) and attitude (bottom subplot). These plots once again show that the covariance is a monotonically increasing function of measurement spacing when the exact expression of (10) is employed, while an artificial "valley" appears when the correlation terms in (10) are ignored.

IV. Filtering with correlated relative-state measurements

We now describe the formulation of an EKF estimator that can fuse proprioceptive and relative-pose measurements, while properly accounting for the correlations in the latter.

To reiterate the challenge posed in Section I, displacement measurements relate *two* robot states, and therefore the *joint pdf* of these states must be available in the filter. For this reason, we augment the EKF (error) state vector[3] to include two copies of the robot's error state (cloning) [9]. The first copy of the error vector, $\widetilde{X}_k$, represents the pose error at the instant when the latest exteroceptive measurement was recorded, while the second copy, $\widetilde{X}_{k+i}$, represents the error in the robot's current state. In the propagation phase of the filter, only the current (evolving) state is propagated, while the previous state remains unchanged. Consequently, the robot states related by each displacement estimate are both represented explicitly in the filter state.

To correctly account for the correlations between consecutive relative-state measurements, the state vector is additionally augmented to include the errors of the latest exteroceptive measurement [11]. Thus, if the most recent exteroceptive measurement was recorded at t_k, the filter's error-state vector at t_{k+i} is:

$$\breve{X}_{k+i|k} = \begin{bmatrix} \widetilde{X}_{k|k}^T & \widetilde{X}_{k+i|k}^T & \tilde{z}_{k+i|k}^{k\,T} \end{bmatrix}^T \quad (20)$$

where the subscript $\ell|j$ denotes the value of a quantity at time t_ℓ, after exteroceptive measurements up to time t_j, and proprioceptive measurements up to time $t_{\ell-1}$, have been processed. It is important to note that when odometry and displacement measurements are combined for pose estimation, it is possible to apply *corrections* to the exteroceptive measurements (cf. Section IV-C). Therefore, in the SC-KF we also maintain an estimate, $\widehat{z}_{k+i|k}^k$, of the most recent measurement[4]. In this notation, the superscript denotes the time instant at which the measurement was received, while the double subscript has the meaning explained above. The errors $\tilde{z}_{k+i|k}^k$ are defined accordingly.

By including the measurement error in the system's state vector, the dependency of the relative-state measurement $z_{k,k+i}$ on the exteroceptive measurement z_k is transformed into a dependency on the *current state of the filter*, and the problem can now be treated in the standard EKF framework. It should be noted that since the measurement error is the source of the correlation between the current and previous displacement estimates, this is the "minimum-length" vector that must be appended to the state vector in order to incorporate the existing dependencies. Thus, this approach yields the minimal computational overhead needed to account for these correlations.

A. Filter Initialization

Consider the case where the first exteroceptive measurement, z_0, is taken at time $t_0 = 0$ and let the robot's state estimate and covariance be denoted by $\widehat{X}_{0|0}$ and $P_{0|0}$, respectively. The initial error-state vector for the SC-KF contains the robot state and its clone, as well as the errors of the exteroceptive measurements at time t_0 (cf. (20)):

$$\breve{X}_{0|0} = \begin{bmatrix} \widetilde{X}_{0|0}^{s\,T} & \widetilde{X}_{0|0}^T & \tilde{z}_{0|0}^{0\,T} \end{bmatrix}^T \quad (21)$$

The superscript s in (21) refers to the static copy of the state, which will remain unchanged during propagation.

Cloning of the robot state creates two identical random variables that convey the same information, and are thus fully correlated. Moreover, since z_0 is not used to estimate the initial robot state, the latter is independent of the measurement errors at time t_0. Thus, the initial covariance matrix of the SC-KF state vector has the form:

$$\breve{P}_{0|0} = \begin{bmatrix} P_{0|0} & P_{0|0} & 0 \\ P_{0|0} & P_{0|0} & 0 \\ 0 & 0 & R_0 \end{bmatrix} \quad (22)$$

where 0 denotes a zero matrix of appropriate dimensions.

B. State Propagation

During regular operation, the filter's state covariance matrix, immediately after the relative-state measurement $z_{k-\ell,k} = \xi(z_{k-\ell}, z_k)$ has been processed, takes the form:

$$\breve{P}_{k|k} = \begin{bmatrix} P_{k|k} & P_{k|k} & P_{X_k z_k} \\ P_{k|k} & P_{k|k} & P_{X_k z_k} \\ P_{X_k z_k}^T & P_{X_k z_k}^T & P_{z_k z_k} \end{bmatrix} \quad (23)$$

where $P_{k|k}$ is the covariance of the robot state at t_k, $P_{z_k z_k}$ is the covariance matrix of the error $\tilde{z}_{k|k}^k$, and $P_{X_k z_k} = E\{\widetilde{X}_k \tilde{z}_{k|k}^{k\,T}\}$ is the cross-correlation between the robot's state and the measurement error at t_k (closed-form expressions for $P_{z_k z_k}$ and $P_{X_k z_k}$ are derived in Section IV-C). Between two consecutive updates, proprioceptive measurements are employed to propagate the filter's state and its covariance. The robot's state estimate is propagated in time by the, generally non-linear, equation:

$$\widehat{X}_{k+1|k} = f(\widehat{X}_{k|k}, v_k) \quad (24)$$

[3] Since the extended form of the Kalman filter is employed for estimation, the state vector comprises the *errors* in the estimated quantities, rather than the estimates. Therefore, cloning has to be applied to both the error states, and the state estimates.

[4] To be more precise, this is an estimate of the *physical quantities* measured by the sensor, such as the distance and bearing to a set of features.

where v_k denotes the proprioceptive (e.g., linear and rotational velocity) measurement at t_k. Linearization of (24) yields the error-propagation equation for the (evolving) second copy of the robot state:

$$\widetilde{X}_{k+1|k} \simeq F_k \widetilde{X}_{k|k} + G_k \widetilde{v}_k \tag{25}$$

where F_k and G_k are the Jacobians of $f(\widehat{X}_{k|k}, v_k)$ with respect to $\widehat{X}_{k|k}$ and v_k, respectively. Since the cloned state, $X^s_{k|k}$, as well as the estimate for the measurement z_k, do not change with the integration of a new proprioceptive measurement, the error propagation equation for the augmented state vector is:

$$\breve{X}_{k+1|k} = \breve{F}_k \breve{X}_{k|k} + \breve{G}_k \widetilde{v}_k \tag{26}$$

$$\text{with} \quad \breve{F}_k = \begin{bmatrix} I & 0 & 0 \\ 0 & F_k & 0 \\ 0 & 0 & I \end{bmatrix} \quad \text{and} \quad \breve{G}_k = \begin{bmatrix} 0 \\ G_k \\ 0 \end{bmatrix} \tag{27}$$

where I denotes an identity matrix of appropriate dimensions. Thus the covariance matrix of the propagated filter state is:

$$\begin{aligned} \breve{P}_{k+1|k} &= \breve{F}_k \breve{P}_{k|k} \breve{F}_k^T + \breve{G}_k Q_k \breve{G}_k^T \\ &= \begin{bmatrix} P_{k|k} & P_{k|k}F_k^T & P_{X_k z_k} \\ F_k P_{k|k} & F_k P_{k|k} F_k^T + G_k Q_k G_k^T & F_k P_{X_k z_k} \\ P_{X_k z_k}^T & P_{X_k z_k}^T F_k^T & P_{z_k z_k} \end{bmatrix} \end{aligned} \tag{28}$$

where $Q_k = E\{\widetilde{v}_k \widetilde{v}_k^T\}$ is the covariance of the proprioceptive measurement v_k.

By straightforward calculation, if m propagation steps occur between two consecutive relative-state updates, the covariance matrix $\breve{P}_{k+m|k}$ is determined as

$$\breve{P}_{k+m|k} = \begin{bmatrix} P_{k|k} & P_{k|k}\mathcal{F}_{k+m,k}^T & P_{X_k z_k} \\ \mathcal{F}_{k+m,k} P_{k|k} & P_{k+m|k} & \mathcal{F}_{k+m,k} P_{X_k z_k} \\ P_{X_k z_k}^T & P_{X_k z_k}^T \mathcal{F}_{k+m,k}^T & P_{z_k z_k} \end{bmatrix} \tag{29}$$

where $\mathcal{F}_{k+m,k} = \prod_{i=0}^{m-1} F_{k+i}$, and $P_{k+m|k}$ is the propagated covariance of the robot state at t_{k+m}. The form of (29) shows that the covariance matrix of the filter can be propagated with minimal computation. In an implementation where efficiency is of utmost importance, the product $\mathcal{F}_{k+m,k}$ can be accumulated, and the matrix multiplications necessary to compute $\breve{P}_{k+m|k}$ can be delayed and carried out only when a new exteroceptive measurement is processed.

C. State Update

We next consider the state-update step of the SC-KF. Assume that a new exteroceptive measurement, z_{k+m}, is recorded at t_{k+m}, and along with $\widehat{z}^k_{k+m|k}$ it is used to produce a relative-state measurement, $z_{k,k+m} = \xi(\widehat{z}^k_{k+m|k}, z_{k+m})$, relating robot poses X_k and X_{k+m}. Note that $z_{k,k+m}$ may not fully determine all the degrees of freedom of the pose change between t_k and t_{k+m}. For example, the scale is unobservable when using a single camera to estimate displacement via point-feature correspondences [8]. Thus, the relative-state measurement is equal to a nonlinear function of the robot poses at t_k and t_{k+m}, with the addition of error:

$$z_{k,k+m} = h(X_k, X_{k+m}) + \widetilde{z}_{k,k+m} \ . \tag{30}$$

The expected value of $z_{k,k+m}$ is computed from the state estimates at t_k and t_{k+m}, as

$$\widehat{z}_{k,k+m} = h(\widehat{X}_{k|k}, \widehat{X}_{k+m|k}) \tag{31}$$

and therefore, based on (2), the innovation is given by:

$$\begin{aligned} r_{k+m} &= z_{k,k+m} - \widehat{z}_{k,k+m} \\ &\simeq H_k \widetilde{X}_{k|k} + H_{k+m} \widetilde{X}_{k+m|k} \\ &\quad + J^k_{k,k+m} \widetilde{z}^k_{k+m|k} + J^{k+m}_{k,k+m} \widetilde{z}^{k+m}_{k+m|k} + n_{k,k+m} \end{aligned} \tag{32}$$

where H_k and H_{k+m} are the Jacobians of $h(X_k, X_{k+m})$ with respect to X_k and X_{k+m}, correspondingly. We note that the quantity $\widetilde{z}^{k+m}_{k+m|k}$ appearing in the last equation is equal to the sensor noise in the measurement z_{k+m}, i.e., $\widetilde{z}^{k+m}_{k+m|k} = \widetilde{z}_{k+m}$.

In order to simplify the presentation of the state update equations, it is helpful to think of the displacement measurement $z_{k,k+m}$ as a constraint relating the robot poses X_k, X_{k+m} and the measurements z_k and z_{k+m}. If we consider the "temporary" variable:

$$X^* = [\breve{X}^T_{k+m|k} \quad \widetilde{z}^{k+m\,T}_{k+m|k}]^T$$

then we can write (32) as

$$\begin{aligned} r_{k+m} &\simeq \begin{bmatrix} H_k & H_{k+m} & J^k_{k,k+m} & J^{k+m}_{k,k+m} \end{bmatrix} X^* + n_{k,k+m} \\ &= H X^* + n_{k,k+m} \end{aligned} \tag{33}$$

This linearized residual expression can be used for carrying out an update on X^* (and thus on its constituent variables), using the standard EKF methodology. The covariance of the residual is

$$\breve{S}_{k+m} = H P H^T + R_{n_{k,k+m}} \tag{34}$$

where $R_{n_{k,k+m}}$ is the covariance of the noise term $n_{k,k+m}$ and

$$P = \begin{bmatrix} \breve{P}_{k+m|k} & 0 \\ 0 & R_{k+m} \end{bmatrix} \tag{35}$$

The Kalman gain for updating X^* is given by:

$$K = P H^T \breve{S}^{-1}_{k+m} = \begin{bmatrix} K_k^T & K_{k+m}^T & K_{z_k}^T & K_{z_{k+m}}^T \end{bmatrix}^T$$

where K_k, K_{k+m}, K_{z_k}, and $K_{z_{k+m}}$ are the block elements of K corresponding to X_k, X_{k+m}, z_k, and z_{k+m}, respectively. We note that although the measurement z_{k+m} can be used to update the robot's pose at t_k and the previous measurement, z_k, these variables will no longer be needed, so we can omit computation of K_k and K_{z_k}. Only the block elements K_{k+m} and $K_{z_{k+m}}$ need to be evaluated. Taking into consideration the special structure of H and P, we obtain:

$$\begin{aligned} K_{k+m} &= \big(\mathcal{F}_{k+m,k} P_{k|k} H_k^T + P_{k+m|k} H_{k+m}^T \\ &\qquad + \mathcal{F}_{k+m,k} P_{X_k z_k} J^{k\,T}_{k,k+m}\big)\, \breve{S}^{-1}_{k+m} \ , \end{aligned} \tag{36}$$

$$K_{z_{k+m}} = R_{k+m} J^{k+m\,T}_{k,k+m} \breve{S}^{-1}_{k+m} \ . \tag{37}$$

Using these results, the equations for updating the *current* robot state and the measurements z_{k+m} are

$$\widehat{X}_{k+m|k+m} = \widehat{X}_{k+m|k} + K_{k+m} r_{k+m} \tag{38}$$

$$\widehat{z}^{k+m}_{k+m|k+m} = z_{k+m} + K_{z_{k+m}} r_{k+m} \tag{39}$$

Algorithm 2 Stochastic Cloning Kalman filter

Initialization: When the first exteroceptive measurement is received:

- clone the state estimate $\widehat{X}_{0|0}$
- initialize the filter state covariance matrix using (22)

Propagation: For each proprioceptive measurement:

- propagate the evolving copy of the robot state via (24)
- propagate the filter covariance using (28), or equivalently (29)

Update: For each exteroceptive measurement:

- compute the relative-state measurement using (1), and its Jacobians with respect to the current and previous exteroceptive measurement, using (3).
- update the current robot state using equations (31), (32), (34), (36), and (38)
- update the current measurement using (37) and (39)
- remove the previous robot state and exteroceptive measurement
- create a cloned copy of the current robot state
- compute the covariance of the new augmented state vector (cf. (40)) using (41)-(44)

After $z_{k,k+m}$ is processed, the clone of the previous state error, $\widetilde{X}_{k|k}$, and the previous measurement error, $\widetilde{z}^k_{k+m|k}$, are discarded. The robot's current state, $X_{k+m|k+m}$, is cloned, and the updated exteroceptive measurement errors, $\widetilde{z}^{k+m}_{k+m|k+m}$, are appended to the new filter state.

Thus, the filter error-state vector becomes

$$\breve{X}_{k+m|k+m} = \begin{bmatrix} \widetilde{X}^T_{k+m|k+m} & \widetilde{X}^T_{k+m|k+m} & \widetilde{z}^{k+m\,T}_{k+m|k+m} \end{bmatrix}^T \quad (40)$$

The state update process is completed by computing the covariance matrix of $\breve{X}_{k+m|k+m}$. To this end, we note that the covariance matrix of X^* is updated as $P \leftarrow P - K\breve{S}_{k+m}K^T$. Using the structure of the matrices involved in this equation, we obtain

$$\breve{P}_{k+m|k+m} = \begin{bmatrix} P_{k+m|k+m} & P_{k+m|k+m} & P_{X_{k+m}z_{k+m}} \\ P_{k+m|k+m} & P_{k+m|k+m} & P_{X_{k+m}z_{k+m}} \\ P^T_{X_{k+m}z_{k+m}} & P^T_{X_{k+m}z_{k+m}} & P_{z_{k+m}z_{k+m}} \end{bmatrix} \quad (41)$$

where

$$P_{k+m|k+m} = P_{k+m|k} - K_{k+m}\breve{S}_{k+m}K^T_{k+m}\ , \quad (42)$$

$$P_{z_{k+m}z_{k+m}} = R_{k+m} - R_{k+m}J^{k+m\ T}_{k,k+m}\breve{S}^{-1}_{k+m}J^{k+m}_{k,k+m}R_{k+m} \quad (43)$$

$$P_{X_{k+m}z_{k+m}} = -K_{k+m}J^{k+m}_{k,k+m}R_{k+m}\ . \quad (44)$$

For clarity, the steps of the SC-KF algorithm are outlined in Algorithm 1.

D. Computational Complexity

While our proposed state-augmentation approach does account for the correlations that have been neglected in previous work, its use imposes a small additional cost in terms of computation and memory requirements. We now show that these algorithmic requirements are *linear* in the number of features observed at a single time-step.

If N and M_k respectively denote the dimensions of the robot's state and the size of the measurement vector at t_k, then the covariance matrix $\breve{P}_{k+m|k}$ has size $(2N + M_k) \times (2N + M_k)$. If $M_k \gg N$, the overhead of state augmentation is mostly due to the inclusion of the measurements in the filter state vector, which leads to the correct treatment of the temporal correlations in the relative-pose measurements. If these correlations are ignored, the size of the filter state vector is twice the size of the robot's state vector. In this case, the computational complexity and memory requirements are $\mathcal{O}(N^2)$. In the algorithm proposed in this paper, the most computationally expensive operation, for $M_k \gg N$, is the evaluation of the covariance of the residual, $\breve{S}_{k+m}$ (cf. (34)). The covariance matrix $\breve{P}_{k+m|k}$ is of dimensions $(2N+M_k)\times(2N+M_k)$, and thus the computational complexity of obtaining $\breve{S}_{k+m}$ is generally $\mathcal{O}((2N+M_k)^2) \approx \mathcal{O}(M_k^2)$. However, from (43) we see that that the submatrix $P_{z_kz_k}$ of $\breve{P}_{k+m|k}$, which corresponds to the updated measurement covariance matrix, has the following structure:

$$P_{z_kz_k} = \underbrace{R_k}_{M_k\times M_k} - \underbrace{R_kJ^{k\ T}_{k-m,k}}_{M_k\times N}\underbrace{\breve{S}^{-1}_k}_{N\times N}\underbrace{J^k_{k-m,k}R_k}_{N\times M_k}$$

As explained in Section III, the measurement noise covariance matrix R_k is commonly block diagonal. Therefore, $P_{z_kz_k}$ has the special structure of a block-diagonal matrix minus a rank-N update. By exploiting this structure when evaluating (34), the operations needed reduce to $\mathcal{O}(N^2M_k)$. Moreover, the submatrix $P_{z_kz_k}$ does not need to be explicitly formed, which decreases the storage requirements of the algorithm to $\mathcal{O}(N^2 + NM_k) \approx \mathcal{O}(NM_k)$. For more details on this point, the interested reader is referred to [18].

Furthermore, for a number of applications, it is not necessary to maintain a clone of the entire robot state and its covariance. Close inspection of the filter update equations reveals that only the states that directly affect the relative-state measurement (i.e., those that are needed to compute the expected relative-state measurement and its Jacobians) are required for the update step. The remaining states and their covariance need not be cloned, thus further reducing the memory and computational requirements. For example, when measurements from an inertial measurement unit (IMU) are employed for localization, estimates for the bias of the IMU measurements are often included in the state vector [19]. These bias estimates clearly do not appear in (31), and therefore it is not necessary to maintain their clones in the filter.

V. Extensions

A. Treatment of Additional Measurements

To simplify the presentation, in the previous section it was assumed that only proprioceptive and relative-pose measurements are available. However, this assumption is not necessary, as additional measurements can be processed in the standard EKF methodology [20]. For example, let

$$z_{k+\ell} = \zeta(X_{k+\ell}) + n_{k+\ell}$$

be an exteroceptive measurement received at $t_{k+\ell}$. By linearizing, we obtain the measurement error equation:

$$\begin{aligned}\widetilde{z}_{k+\ell} &= H'_{k+\ell}\widetilde{X}_{k+\ell|k} + n_{k+\ell} \\ &= \begin{bmatrix}0 & H'_{k+\ell} & 0\end{bmatrix}\begin{bmatrix}\widetilde{X}_{k|k} \\ \widetilde{X}_{k+\ell|k} \\ \widetilde{z}^k_{k+\ell|k}\end{bmatrix} + n_{k+\ell} \end{aligned} \quad (45)$$

Since this expression adheres to the standard EKF model, the augmented filter state can be updated without any modifications to the algorithm. However, if additional measurements are processed, the compact special expressions of (29) and (36) are no longer valid, as update steps occur between consecutive displacement estimates. In this case, the general form of the SC-KF equations must be used.

Another practically important case occurs when more than one sensor provides relative-pose measurements, but at different rates. Such a situation would arise, for example, when a mobile robot is equipped with a camera and a laser range finder. In such a scenario, the state-augmentation approach of the SC-KF still applies. In particular, every time either of the sensors records a measurement, cloning is applied. Therefore, at any given time the filter state vector is comprised of i) three instances of the robot state, corresponding to the current state, and the state at the last time instants where each sensor received a measurement, and ii) the errors in the latest exteroceptive measurement of each sensor. Although the propagation and update equations must be modified to account for the change in dimension of the state vector, the basic principles of the approach still apply.

B. Extension to Multiple States

In the algorithm presented in Section IV, feature measurements are processed to construct displacement estimates, which subsequently define constraints between consecutive robot poses. By including two robot poses in the filter state vector, the SC-KF can optimally process successive exteroceptive measurements, while incurring a computational cost linear in the number of observed features. However, when a static feature is observed more than two times, the basic SC-KF must be modified. Intuitively, the observation of a static feature from multiple robot poses should impose a geometric constraint involving these measurements and *all* of the corresponding poses. We now briefly describe an extension to the SC-KF approach that correctly incorporates multiple observations of a single point feature while still maintaining computational complexity *linear* in the number of locally observed features [21].

Let Y_{f_j} be the position of a static feature, which is observed from $L \geq 2$ consecutive robot poses, $X_k, X_{k+1}, \ldots, X_{k+L-1}$. The measurement function, h_{f_j}, corresponding to these measurements is

$$z^{f_j}_{k+i} = h_{f_j}(X_{k+i}, Y_{f_i}) + n^{f_j}_{k+i}, \quad \text{for} \quad i = 0 \ldots L-1 \quad (46)$$

where $n^{f_j}_{k+i}$ is the measurement noise. Stacking these L equations results in a block measurement equation of the form:

$$\mathbf{z}_{f_j} = \mathbf{h}_{f_j}(X_k, X_{k+1}, \ldots, X_{k+L-1}, Y_{f_j}) + \mathbf{n}_{f_j} \ . \quad (47)$$

Eliminating the feature position, Y_{f_j}, from (47) yields a constraint vector that involves all of the robot poses:

$$\mathbf{c}_{f_j}(X_k, X_{k+1}, \ldots, X_{k+L-1}, \mathbf{z}_{f_j}, \mathbf{n}_{f_j}) = \mathbf{0}_q \quad (48)$$

where q is the dimension of the constraint vector $\mathbf{c}_{f_j}$. If the EKF state vector has been augmented to include the L copies of the robot pose, the above equation can be used to perform an EKF update, thus utilizing all the geometric information provided by the observations of this feature. Furthermore, if M_k features are observed from L robot poses, then a constraint vector, $\mathbf{c}_{f_j}$, $j = 1 \ldots M_k$, can be written for each of these features. Since the feature measurements are mutually uncorrelated, the resulting constraints will also be uncorrelated, and therefore, an EKF update that utilizes all M_k constraints can be performed in $O(M_k)$ time.

VI. Relation to SLAM

An alternative approach to processing the feature measurements obtained with an exteroceptive sensor is to jointly estimate the robot's pose and the feature positions. This is the well-known SLAM problem, which has been extensively studied (e.g., [22]–[25]). This section examines the relation of the SC-KF algorithm to SLAM.

1) Computational complexity: If an exact solution to SLAM was possible, the resulting pose estimates would be optimal, since all the positioning information would be used and all the inter-dependencies between the robot and the feature states would be accounted for. However, good localization performance comes at a considerable computational cost. It is well known that the computational complexity and memory requirements of the EKF solution to SLAM increase quadratically with the total number of features in the environment [22]. While several approximate solutions exist that possess lower computational complexity (e.g., [23], [25], [26]), many of them cannot guarantee the consistency of the estimates, nor is there a concrete measure of suboptimality.

Since the high computational burden of SLAM is due to the need to maintain a map of the environment, the amount of computational resources allocated for localization constantly increases as the robot navigates in an unknown environment. For continual operation over an extended period, this overhead can become unacceptably large. Even in an approximate SLAM algorithm, the largest portion of the computational resources is devoted to maintaining the constantly expanding feature map. However, there exist a number of applications where building a map is not necessary, while real-time performance is of utmost importance (e.g., in autonomous aircraft landing [27], or emergency response [28]). Such applications require high localization accuracy, but with minimal computational overhead.

The SC-KF uses pairs of consecutive exteroceptive measurements to produce displacement estimates, which are then fused with proprioceptive sensing information. As shown in Section IV-D, our algorithm's complexity is linear in the number of features observed *only at each time-step*. In most cases this number is orders of magnitude smaller than the total number of features in the environment. A reduced-complexity

SLAM approach that is similar in spirit to the SC-KF would consist of maintaining only the most recently acquired local features, i.e., those that are currently visible by the robot, in the state vector. However, the algorithmic complexity of such an EKF-SLAM would be *quadratic* in the number of local features. In contrast, the SC-KF is *linear* in the number of local features.

2) Feature position observability: SLAM algorithms require the states of the local features to be completely observable, in order to be included in the state vector. When a single measurement does not provide sufficient information to initialize a feature's position estimate with bounded uncertainty, feature initialization schemes must be implemented [29], [30]. In fact, state augmentation is an integral part of many methods for delayed feature initialization [31], [32]. In contrast, in the SC-KF framework, feature initialization is not required since the feature *measurements* are included in the augmented state vector, instead of the feature *positions*.

3) Data association: Since only pairs of exteroceptive measurements are used by the SC-KF algorithm, the data association problem is simplified. In contrast, SLAM requires a correspondence search over all map features in the robot's vicinity and its computational overhead is considerably higher [33]. To facilitate robust data association, it is common practice to employ a feature detection algorithm that extracts "high-level" features (e.g., landmarks such as corners, junctions, straight-line segments) from raw sensor data. Then, only these features are employed for SLAM.

4) Information loss: While the extraction of high-level features results in more robust and computationally tractable algorithms (e.g., laser scans consist of hundreds of range points, but may contain only a few corner features), this approach effectively *discards information* contained in the sensor data (cf. Fig. 1). Consequently, the resulting estimates of the robot's pose are suboptimal compared to those that use all the available information. Maintaining and processing the entire history of raw sensor input (e.g., [34]) can lead to excellent localization performance, but such an approach may be infeasible for real-time implementation on typical mobile robots. A benefit of the SC-KF approach is that it takes advantage of all the available information in two consecutive exteroceptive measurements (i.e., most laser points in two scans can be used to estimate displacement by scan matching).

5) SC-KF and SLAM: For longer robot traverses, the positioning accuracy obtained when only pairs of exteroceptive measurements are considered is inferior to that of SLAM, as *no loop closing occurs*. Essentially, the SC-KF approach offers an enhanced form of Dead Reckoning, in the sense that the uncertainty of the robot's state monotonically increases over time. The rate of uncertainty increase, though, is significantly lower than that attained when only proprioceptive measurements are used (cf. Section VII). However, as mentioned in Section IV-D, in the SC-KF approach the state vector X_k is not required to contain only the robot pose. If high-level, stable features (landmarks) are available, their positions can be included in the "robot" state vector X_k. Therefore, the SC-KF method for processing relative-state measurements can be expanded and integrated with the SLAM framework. This integration would further improve the attainable localization accuracy within areas with lengthy loops. Since this modification is beyond the scope of this work, in the following section we present experimental results applying the SC-KF algorithm to the case where only relative-state and proprioceptive measurements are considered.

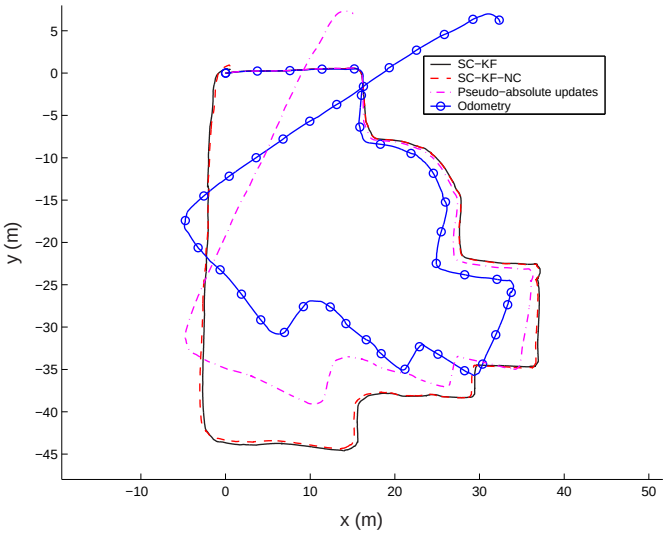

Fig. 4. The estimated trajectory of the robot using the SC-KF algorithm (solid line), the SC-KF-NC algorithm (dashed line), the method of [10] that uses absolute position pseudo-measurements (dash-dotted line), and odometry only (solid line with circles).

VII. Experimental results

This section presents experimental results that demonstrate the performance of the algorithms described in Sections IV and III-A. The experiments use a Pioneer II mobile robot equipped with a SICK LMS-200 laser rangefinder. The robot's pose consists of its planar position and orientation in a global frame:

$$X_k = \begin{bmatrix} {}^G x_k & {}^G y_k & {}^G \phi_k \end{bmatrix}^T = \begin{bmatrix} {}^G p_k^T & {}^G \phi_k \end{bmatrix}^T \quad (49)$$

We first present results from the application of the SC-KF, and then study the case where the robot's state is propagated based on displacement estimates exclusively (i.e., no proprioceptive measurements are processed).

A. Stochastic Cloning Kalman Filter

In this experiment, odometry measurements are fused with displacement measurements that are obtained by laser scan matching with the method presented in [6]. The SC-KF equations for the particular odometry and measurement model are presented in [18].

1) Experiment description: During the first experiment, the robot traversed a trajectory of approximately 165 m, while recording 378 laser scans. The robot processed a new laser scan approximately every 1.5 m, or every time its orientation changed by 10°. We here compare the performance of the SC-KF algorithm to that obtained by the approach of Hoffman *et al.* [10]. In [10], the displacement estimates and the previous pose estimates are combined to yield pseudo-measurements of

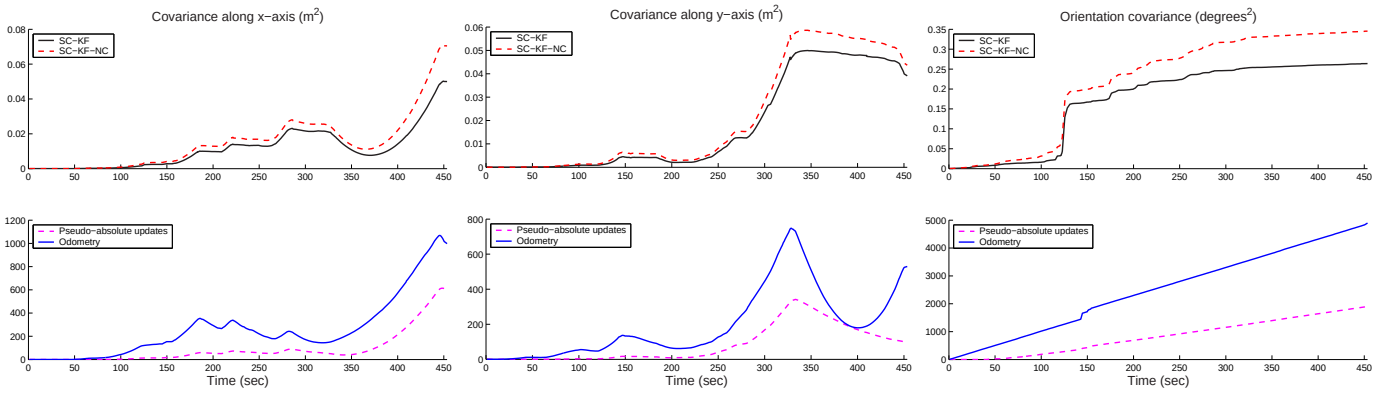

Fig. 5. The time evolution of the diagonal elements of the covariance matrix of the robot's pose. Note the difference in the vertical axes' scale. In these plots, the covariance values after filter updates are plotted.

the robot's absolute position. In order to guarantee consistent estimates for the latter case, we have employed the Covariance Intersection (CI) method [35] for fusing the pseudo-measurements of absolute position with the most current pose estimates. From here on we refer to this approach as "pseudo-absolute updates".

As discussed in Section IV-D, the SC-KF has computational complexity linear in the number of feature measurements taken at each pose. If even this computational complexity is deemed too high for a particular application, one can ignore the correlations between consecutive displacement measurements, at the expense of optimality. In that case, the augmented state only contains the two copies of the robot state [9]. Results for this approximate, though computationally simpler, variant of the SC-KF, referred to as SC-KF-NC (i.e., no correlations between the measurement errors are considered), are presented below and are compared with the performance of the SC-KF.

The robot trajectories estimated by the different algorithms are shown in Fig. 4. Fig. 5 presents the covariance estimates for the robot pose as a function of time. We observe that correctly accounting for the correlations between consecutive displacement estimates in the SC-KF, results in smaller covariance values. Even though ground truth for the entire trajectory is not known, the final robot pose is known to coincide with the initial one. The errors in the final robot pose are equal to $\widetilde{X} = [0.5\text{m} \quad 0.44\text{m} \quad -0.11^\text{o}]^T$ (0.4% of the trajectory length) for the SC-KF, $\widetilde{X} = [0.61\text{m} \quad 0.65\text{m} \quad -0.13^\text{o}]^T$ (0.54% of the trajectory length) for the SC-KF-NC, $\widetilde{X} = [15.03\text{m} \ 7.07\text{m} \ -32.3^\text{o}]^T$ (10.6% of the trajectory length) for the approach of [10], and $\widetilde{X} = [32.4\text{m} \quad 5.95\text{m} \quad -69.9^\text{o}]^T$ (19.9% of the trajectory length) for Dead Reckoning based on odometry. From these error values, as well as from visual inspection of the trajectory estimates in Fig. 4, we conclude that both the SC-KF and the SC-KF-NC yield very similar results. However, the approach based on creating pseudo-measurements of the absolute pose [10] performs significantly worse. It should be noted that the errors in the final robot pose are consistent with the estimated covariance in all cases considered.

2) Impact of correlations: Clearly, the lack of ground truth data along the entire trajectory for the real-world experiment does not allow for a detailed comparison of the performance of the SC-KF and SC-KF-NC algorithms, as both appear to attain comparable estimation accuracy. Simulations are used to perform a more thorough assessment of the impact of the measurement correlations on the position accuracy and the uncertainty estimates. The primary objective of these simulations is to contrast the magnitude of the estimation errors with the computed covariance values in the cases when the correlations between consecutive measurements are accounted for (SC-KF), vs. when they are ignored (SC-KF-NC).

For the simulation results shown here, a robot moves in a circular trajectory of radius 4 m, while observing a wall that lies 6 m from the center of its trajectory. The relative-pose measurements in this case are created by performing line-matching, instead of point matching between consecutive scans [36]. Since only one line is available, the motion of the robot along the line direction is unobservable. As a result, the singular value decomposition of the covariance matrix of the robot's displacement estimate can be written as

$$R_{k,k+m} \quad = \quad \begin{bmatrix} V_u & V_o \end{bmatrix} \begin{bmatrix} s_1 & 0 & 0 \\ 0 & s_2 & 0 \\ 0 & 0 & s_3 \end{bmatrix} \begin{bmatrix} V_u^T \\ V_o^T \end{bmatrix}, \quad s_1 \rightarrow \infty$$

where V_u is the basis vector of the unobservable direction (i.e., a unit vector along the direction of the wall, expressed with respect to the robot frame at time t_k) and V_o is a 3×2 matrix, whose column vectors form the basis of the observable subspace. To avoid numerical instability in the filter, the displacement measurements, $z_{k,k+m}$ computed by line-matching are projected onto the observable subspace, thus creating a relative-state measurement of dimension 2, given by $z'_{k,k+m} = V_o^T z_{k,k+m}$.

Fig. 6 shows the robot pose errors (solid lines), along with the corresponding 99.8th percentile of their distribution (dashed lines). The left column shows the results for the SC-KF algorithm presented in Section IV, while the right one for the SC-KF-NC algorithm. As evident from Fig. 6, the covariance estimates of the SC-KF-NC are not commensurate

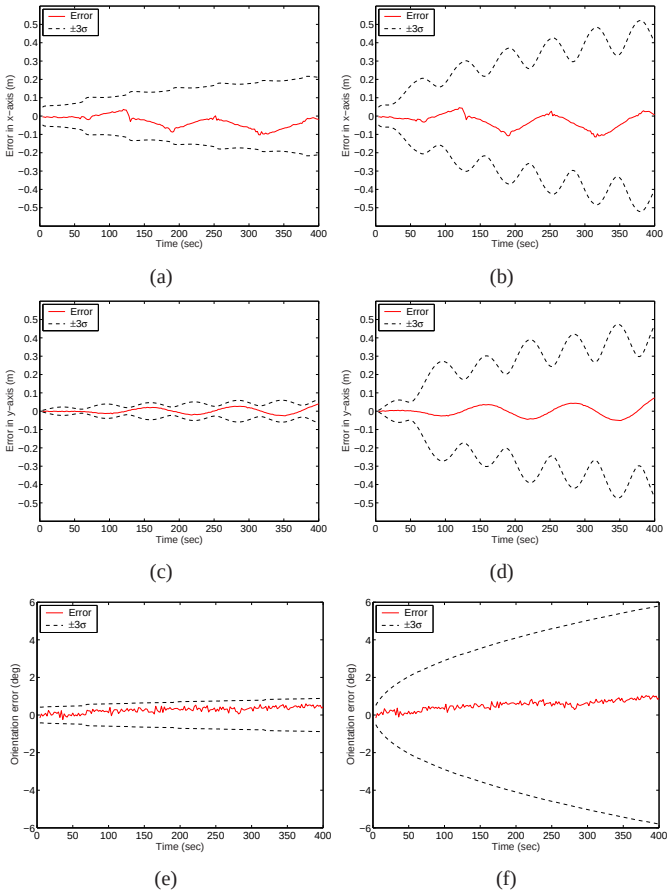

Fig. 6. The robot pose errors (solid lines) vs. the corresponding 99.8th percentile of their distribution, (dashed lines). The left column shows the results for the SC-KF algorithm proposed in this paper, while the right one demonstrates the results for the SC-KF-NC algorithm. In these plots, the covariance values after filter updates are plotted. (a - b) Errors and $\pm 3\sigma$ bounds along the x-axis (c - d) Errors and $\pm 3\sigma$ bounds along the y-axis (e - f) Orientation errors and $\pm 3\sigma$ bounds.

with the corresponding errors. When the temporal correlations of the measurements are properly treated, as is the case for the SC-KF, substantially more accurate covariance estimates, which reflect the true uncertainty of the robot's state, are computed. Moreover, evaluation of the rms value of the pose errors shows that the errors associated with the SC-KF algorithm (which accounts for correlations) are 25% smaller than those of the SC-KF-NC.

B. State Propagation based on Displacement Estimates

We now present results for the case in which the robot's pose is estimated using only displacement estimates computed from laser scan matching. Given a displacement estimate $z_{k,k+m} = [{}^k\widehat{p}_{k+m}^T \quad {}^k\widehat{\phi}_{k+m}]^T$, the global robot pose is propagated using the equations

$$\begin{aligned} \widehat{X}_{k+m} &= g(\widehat{X}_k, \widehat{z}_{k,k+m}) \Rightarrow \\ \begin{bmatrix} {}^G\widehat{p}_{k+m} \\ {}^G\widehat{\phi}_{k+m} \end{bmatrix} &= \begin{bmatrix} {}^G\widehat{p}_k \\ {}^G\widehat{\phi}_k \end{bmatrix} + \begin{bmatrix} C({}^G\widehat{\phi}_k){}^k\widehat{p}_{k+m} \\ {}^k\widehat{\phi}_{k+m} \end{bmatrix} \end{aligned} \quad (50)$$

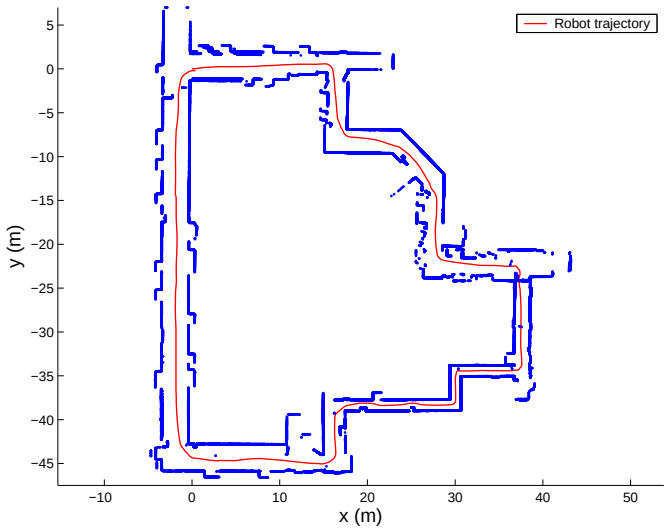

Fig. 7. The estimated trajectory of the robot based only on laser scan matching. The map is presented for visualization purposes only, by transforming all the laser points using the estimated robot pose. Some "spurious" points in the map are due to the presence of people.

where $C(\cdot)$ denotes the 2×2 rotation matrix. In this case, the Jacobian matrices Φ_k and Γ_k are given by

$$\Phi_k = \begin{bmatrix} I & -\Psi C({}^G\widehat{\phi}_k)\,{}^k\widehat{p}_{k+m} \\ 0 & 1 \end{bmatrix}, \quad \Psi = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$$

$$\Gamma_k = \begin{bmatrix} C({}^G\widehat{\phi}_k) & 0 \\ 0 & 1 \end{bmatrix}.$$

Fig. 7 presents the estimated robot trajectory, along with the map of the area that has been constructed by overlaying all the scan points, transformed using the estimates of the robot pose (we stress that the map is only plotted for visualization purposes, and is not estimated by the algorithm). This experiment used the same dataset from Section VII-A. Fig. 8 presents covariance estimates for the robot's pose, computed using (10) (SC-KF, solid lines) in contrast with those computed when the correlations between the consecutive displacement estimates are ignored (SC-KF-NC, dashed lines). As expected, the pose covariance is larger when only displacement measurements are used, compared to the case where odometry measurements are fused with displacement measurements (cf. Fig. 5). From Fig. 8 we also observe that accounting for the correlations results in significantly smaller values for the estimated covariance of the robot pose, thus corroborating the discussion of Section III-B.

VIII. Conclusions

In this paper, we have proposed an efficient EKF-based estimation algorithm, termed a *Stochastic Cloning*-Kalman Filtering (SC-KF), for the problem of fusing proprioceptive measurements with relative-state measurements that are inferred from exteroceptive sensory input. An analysis of the structure of the measurement equations demonstrated that when the same exteroceptive measurements are processed to estimate displacement in consecutive time intervals, the

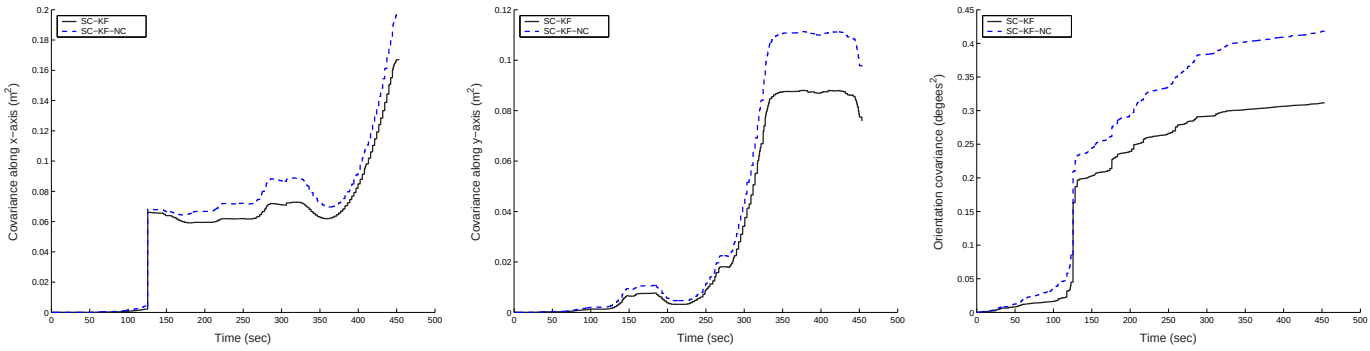

(a) (b) (c)

Fig. 8. The estimated covariance of the robot's pose when the correlation between consecutive measurements is properly accounted for (solid lines) vs. the covariance estimated when the correlations are ignored (dashed lines). (a) Covariance along the x-axis (b) Covariance along the y-axis (c) Orientation Covariance. At approximately 130 sec, a displacement estimate based on very few laser points was computed, resulting in a sudden increase in the covariance.

displacement errors are temporally correlated. The main contribution of this work is the introduction of a novel feature-marginalization process that allows for the processing of relative-pose measurements while also considering the correlations between these. This method is based on augmenting the state vector of the EKF to temporarily include the robot poses and the feature observations related through a local geometric constraint (i.e., a relative-state measurement). By employing state augmentation, the dependence of the relative-state measurement on previous states and measurements is transformed to a dependence on the *current* state of the filter, and this enables application of the standard EKF framework.

The experimental and simulation results demonstrate that the SC-KF method attains better localization performance compared to previous approaches [10], while the overhead imposed by the additional complexity is minimal. The method yields more accurate estimates, and most significantly, it provides a more precise description of the uncertainty in the robot's state estimates. Additionally, the method is versatile, since it is independent of the actual sensing modalities used to obtain the proprioceptive and exteroceptive measurements.

ACKNOWLEDGEMENTS

This work was supported by the University of Minnesota (DTC), the NASA Mars Technology Program (MTP-1263201), and the National Science Foundation (EIA-0324864, IIS-0643680).

REFERENCES

[1] A. Kelly, "General solution for linearized systematic error propagation in vehicle odometry," in *Proc. IEEE/RSJ Int. Conf. on Robots and Systems*, Maui, HI, Oct.29-Nov.3 2001, pp. 1938–45.

[2] L. Matthies, "Dynamic stereo vision," Ph.D. dissertation, Dept. of Computer Science, Carnegie Mellon University, 1989.

[3] C. Olson, L. Matthies, H. Schoppers, and M. Maimone, "Robust stereo ego-motion for long distance navigation," in *Proceedings of CVPR*, 2000, pp. 453–458.

[4] F. Lu and E. Milios, "Robot pose estimation in unknown environments by matching 2d range scans," *Journal of Intelligent and Robotic Systems: Theory and Applications*, vol. 18, no. 3, pp. 249–275, Mar. 1997.

[5] D. Silver, D. M. Bradley, and S. Thayer, "Scan matching for flooded subterranean voids," in *Proc. IEEE Conf. on Robotics, Automation and Mechatronics (RAM)*, Dec. 2004.

[6] S. T. Pfister, K. L. Kriechbaum, S. I. Roumeliotis, and J. W. Burdick, "Weighted range sensor matching algorithms for mobile robot displacement estimation," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington D.C., May 11-15 2002, pp. 1667–74.

[7] K. Konolige, "Large-scale map-making," in *AAAI National Conference on Artificial Intelligence*, San Jose, CA, July 2004, pp. 457–463.

[8] P. Torr and D. Murray, "The development and comparison of robust methods for estimating the fundamental matrix," *International Journal of Computer Vision*, vol. 24, no. 3, pp. 271–300, 1997.

[9] S. I. Roumeliotis and J. W. Burdick, "Stochastic cloning: A generalized framework for processing relative state measurements," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington D.C., 2002, pp. 1788–1795.

[10] B. D. Hoffman, E. T. Baumgartner, T. L. Huntsberger, and P. S. Schenker, "Improved state estimation in challenging terrain," *Autonomous Robots*, vol. 6, no. 2, pp. 113–130, April 1999.

[11] A. I. Mourikis and S. I. Roumeliotis, "On the treatment of relative-pose measurements for mobile robot localization," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Orlando, FL, May 15-19 2006, pp. 2277 – 2284.

[12] S. I. Roumeliotis, "A kalman filter for processing 3-d relative pose measurements," California Institute of Technology, Tech. Rep., Mar. 2002, http://robotics.caltech.edu/∼stergios/tech_reports/relative_3d_kf.pdf.

[13] S. Fleischer, "Bounded-error vision-based navigation of autonomous underwater vehicles," Ph.D. dissertation, Stanford University, 2000.

[14] R. Eustice, H. Singh, and J. Leonard, "Exactly sparse delayed-state filters," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Barcelona, Spain, April 2005, pp. 2428–2435.

[15] R. Eustice, O. Pizarro, and H. Singh, "Visually augmented navigation in an unstructured environment using a delayed state history," in *Proc. IEEE Int. Conf. on Robotics and Automation*, New Orleans, LA, April 2004, pp. 25–32.

[16] R. Garcia, J. Puig, O. Ridao, and X. Cufi, "Augmented state Kalman filtering for AUV navigation," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 2002, pp. 4010–4015.

[17] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping," *Autonomous Robots*, vol. 4, no. 4, 1997.

[18] A. I. Mourikis and S. I. Roumeliotis, "SC-KF mobile robot localization: A Stochastic Cloning-Kalman filter for processing relative-state measurements," University of Minnesota, Dept. of Computer Science and Engineering, Tech. Rep., October 2006.

[19] E. J. Lefferts, F. L. Markley, and M. D. Shuster, "Kalman filtering for spacecraft attitude estimation," *Journal of Guidance, Control, and Dynamics*, vol. 5, no. 5, pp. 417–429, Sept.-Oct. 1982.

[20] D. Bayard and P. B. Brugarolas, "An estimation algorithm for vision-based exploration of small bodies in space," in *Proceedings of the 2005 American Control Conference*, Portland, OR, June 2005, pp. 4589–4595.

[21] A. I. Mourikis and S. I. Roumeliotis, "A multi-state constraint Kalman filter for vision-aided inertial navigation," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Rome, Italy, April 10-14 2007, pp. 3565–3572.

[22] R. C. Smith, M. Self, and P. Cheeseman, *Autonomous Robot Vehicles*. Springer-Verlag, 1990, ch. Estimating Uncertain Spatial Relationships in Robotics, pp. 167–193.

[23] M. Montemerlo, "FastSLAM: A factored solution to the simultaneous localization and mapping problem with unknown data association," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, July 2003.
[24] S. Thrun, D. Koller, Z. Ghahramani, and H. Durrant-Whyte, "Simultaneous mapping and localization with sparse extended information filters: Theory and initial results," School of Computer Science, Carnegie Mellon University, Tech. Rep., 2002.
[25] P. Newman, J. Leonard, J. D. Tardos, and J. Neira, "Explore and return: experimental validation of real-time concurrent mapping and localization," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Washington, DC, May 11-15 2002, pp. 1802–9.
[26] S. J. Julier and J. K. Uhlmann, "Simultaneous localisation and map building using split covariance intersection," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Maui, HI, Oct. 29-Nov. 3 2001, pp. 1257–62.
[27] S. I. Roumeliotis, A. E. Johnson, and J. F. Montgomery, "Augmenting inertial navigation with image-based motion estimation," in *IEEE International Conference on Robotics and Automation*, Washington D.C., 2002, pp. 4326–33.
[28] P. Pinies and J. D. Tardos, "Fast localization of avalanche victims using sum of Gaussians," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, Orlando, FL, May 2006, pp. 3989 – 3994.
[29] T. Bailey, "Constrained initialisation for bearing-only SLAM," in *Proc. IEEE Int. Conf. on Robotics and Automation*, vol. 2, Sept. 2003, pp. 1966–1971.
[30] N. M. Kwok and G. Dissanayake, "An efficient multiple hypothesis filter for bearing-only SLAM," in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Sendai, Japan, Oct. 2004, pp. 736–741.
[31] J. Leonard, R. Rikoski, P. Newman, and M. Bosse, "Mapping partially observable features from multiple uncertain vantage points," *Int. J. Robotics Research*, vol. 21, no. 10-11, pp. 943–975, 2002.
[32] J. Leonard and R. Rikoski, "Incorporation of delayed decision making into stochastic mapping," *Experimental Robotics VII*, vol. 271, pp. 533–542, 2001.
[33] J. Neira and J. D. Tardos, "Data association in stochastic mapping using the joint compatibility test," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 890–897, 2001.
[34] A. Howard, "Multi-robot mapping using manifold representations," in *Proc. 2004 IEEE Int Conf. on Robotics and Automation*, New Orleans, LA, April 2004, pp. 4198–4203.
[35] S. Julier and J. Uhlman, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proceedings of the American Control Conference*, vol. 4, 1997, pp. 2369 – 2373.
[36] S. T. Pfister, S. I. Roumeliotis, and J. W. Burdick, "Weighted line fitting algorithms for mobile robot map building and efficient data representation," in *Proc. IEEE Int. Conf. on Robotics and Automation*, Taipei, Taiwan, Sep. 14-19 2003, pp. 1304–1311.

Anastasios I. Mourikis received the Diploma of Electrical and Computer Engineering with honors from the University of Patras, Greece in 2003. He is currently a PhD candidate at the Department of Computer Science and Engineering (CSE) at the University of Minnesota. His research interests lie in the areas of localization in single- and multi-robot systems, vision-aided inertial navigation, simultaneous localization and mapping, and structure from motion. He is the recipient of the 2005 and 2006 Excellence in Research Award Fellowships from the CSE Department of the University of Minnesota, and of the 2007-2008 Doctoral Dissertation Fellowship from the University of Minnesota.

Stergios I. Roumeliotis received his Diploma in Electrical Engineering from the National Technical University of Athens, Greece, in 1995, and the M.S. and Ph.D. degrees in Electrical Engineering from the University of Southern California, CA in 1997 and 2000 respectively. From 2000 to 2002 he was a postdoctoral fellow at the California Institute of Technology, CA. Since 2002 he has been an Assistant Professor in the Department of Computer Science and Engineering at the University of Minnesota, MN. His research interests include inertial navigation of aerial and ground autonomous vehicles, fault detection and identification, and sensor networks. Recently his research has focused on distributed estimation under communication and processing constraints and active sensing for reconfigurable networks of mobile sensors.

Dr. Roumeliotis is the recipient of the NSF CAREER award, the McKnight Land-Grant Professorship award, and he is the co-recipient of the NASA Tech Briefs award, the One NASA Peer award, and the One NASA Center Best award. He is currently serving as Associate Editor for the IEEE Transactions on Robotics.

Joel W. Burdick received his undergraduate degree in mechanical engineering from Duke University and M.S. and Ph.D. degrees in mechanical engineering from Stanford University. He has been with the department of Mechanical Engineering at the California Institute of Technology since May 1988, where he has been the recipient of the NSF Presidential Young Investigator award, the Office of Naval Research Young Investigator award, and the Feynman fellowship. He has been a finalist for the best paper award for the IEEE International Conference on Robotics and Automation in 1993, 1999, 2000, and 2005. He was appointed an IEEE Robotics Society Distinguished Lecturer in 2003. Prof. Burdick's research interests lie mainly in the areas of robotics, mechanical systems, and bioengineering. Current research interests include sensor based robot motion planning, multi-fingered grasping, coordination of multiple robots, neural prosthetics, and rehabilitation of spinal cord injuries.